

# Thompson Sampling Algorithms for Cascading Bandits

**Zixin Zhong**

*Department of Mathematics  
National University of Singapore  
119076, Singapore*

ZIXIN.ZHONG@U.NUS.EDU

**Wang Chi Chueng**

*Department of Industrial Systems and Management  
National University of Singapore  
117576, Singapore*

ISECWC@NUS.EDU.SG

**Vincent Y. F. Tan**

*Department of Electrical and Computer Engineering and  
Department of Mathematics  
National University of Singapore  
117583, Singapore*

VTAN@NUS.EDU.SG

**Editor:** Csaba Szepesvári

## Abstract

Motivated by the important and urgent need for efficient optimization in online recommender systems, we revisit the cascading bandit model proposed by Kveton et al. (2015a). While Thompson sampling (TS) algorithms have been shown to be empirically superior to Upper Confidence Bound (UCB) algorithms for cascading bandits, theoretical guarantees are only known for the latter. In this paper, we first provide a problem-dependent upper bound on the regret of a TS algorithm with Beta-Bernoulli updates; this upper bound is tighter than a recent derivation under a more general setting by Huyuk and Tekin (2019). Next, we design and analyze another TS algorithm with Gaussian updates, TS-CASCADE. TS-CASCADE achieves the state-of-the-art problem-independent regret bound for cascading bandits. Complementarily, we consider a linear generalization of the cascading bandit model, which allows efficient learning in large-scale cascading bandit problem instances. We introduce and analyze a TS algorithm, which enjoys a regret bound that depends on the dimension of the linear model but not the number of items. Finally, by using information-theoretic techniques and a judicious construction of cascading bandit instances, we derive a nearly-matching lower bound on the expected regret for the standard model. Our paper establishes the first theoretical guarantees on TS algorithms for a stochastic combinatorial bandit problem model with partial feedback. Numerical experiments demonstrate the superiority of the proposed TS algorithms compared to existing UCB-based ones.

**Keywords:** Multi-armed bandits, Thompson sampling, Cascading bandits, Linear bandits, Regret minimization

## 1. Introduction

Online recommender systems seek to recommend a small list of items (such as movies or hotels) to users from a larger ground set  $[L] := \{1, \dots, L\}$  of items. Optimizing the performance of these systems is of fundamental importance in the e-service industry, where

companies such as Yelp and Spotify strive to maximize users’ satisfaction by catering to their taste. The model we consider in this paper is the *cascading bandit* model (Kveton et al., 2015a), which models online learning in the standard *cascade* model by Craswell et al. (2008). The latter model (Craswell et al., 2008) is widely used in information retrieval and online advertising.

In the cascade model, a learning agent offers a list of items to a user. The user then scans through it in a sequential manner. The user looks at the first item, and if she is *attracted* by it, she *clicks* on it. If not, she skips to the next item and clicks on it if she finds it attractive. This process stops when she clicks on one item in the list, or when she comes to the end of the list, in which case she is *not attracted* by *any* of the items. The items that are in the ground set but not in the chosen list and those in the list that come after the attractive one are *unobserved*. Each item  $i$  is associated with a *click probability*  $w(i) \in [0, 1]$ , which in general varies across different items. The event that a user is attracted to an item is independent of the events that she is attracted to other items. The learning agent’s objective is to maximize the cumulative rewards.

In the cascading bandit model, which is a *multi-armed bandit* version of the cascade model by (Craswell et al., 2008), the click probabilities  $\mathbf{w} := \{w(i)\}_{i=1}^L$  are *not known* to the learning agent, and should be learnt over time. Based on the lists previously chosen and the observations obtained thus far, the agent tries to learn the click probabilities (exploration) in order to adaptively and judiciously recommend other lists of items (exploitation) to maximize his overall reward over  $T$  time steps.

Apart from the cascading bandit model where each item’s click probability is learnt individually (we call it the *standard cascading bandit problem*), we also consider a *linear generalization of the cascading bandit model*, which is called the linear cascading bandit problem, proposed by (Zong et al., 2016). In the linear model, each click probability  $w(i)$  is known to be equal to  $x(i)^\top \beta \in [0, 1]$ , where the feature vector  $x(i)$  is *known* for each item  $i$ , but the latent  $\beta \in \mathbb{R}^d$  is *unknown* and should be learnt over time. The feature vectors  $\{x(i)\}_{i=1}^L$  represents the prior knowledge of the learning agent. When  $d < L$ , the learning agent can potentially harness the prior knowledge  $\{x(i)\}_{i=1}^L$  to learn  $\{w(i)\}_{i=1}^L$  by estimating  $\beta$ , which is more efficient than by estimating each  $w(i)$  individually.

## 1.1 Main Contributions

First, we analyze the Thompson sampling algorithm with the Beta-Bernoulli update, which is an application of the Combinatorial Thompson Sampling (CTS) algorithm (Russo et al., 2018; Huyuk and Tekin, 2019) in the cascading bandit setting. We provide a concise analysis which leads to a tighter upper bound compared to the original one reported in Huyuk and Tekin (2019).

Our next contribution is the design and analysis of TS-CASCADE, a Thompson sampling algorithm (Thompson, 1933) for the standard cascading bandit problem. Our motivation is to design an algorithm that can be generalized and applied to the linear bandit setting, since it is not as natural to design a linear generalization of the CTS algorithm. The variation of TS-CASCADE for the linear setting is presented in Section 5. Our design involves two novel features. First, the Bayesian estimates on the vector of latent click probabilities  $\mathbf{w}$  are constructed by a univariate Gaussian distribution. Consequently, in each time step,

TS-CASCADE conducts exploration in a suitably defined one-dimensional space. Second, inspired by Audibert et al. (2009), we judiciously incorporate the empirical variance of each item’s click probability in the Bayesian update. This allows efficient exploration on item  $i$  when  $w(i)$  is close to 0 or 1.

We establish a problem-independent regret bound for our proposed algorithm TS-CASCADE. Our regret bound matches the state-of-the-art regret bound for UCB algorithms on the standard cascading bandit model (Wang and Chen, 2017), up to a multiplicative logarithmic factor in the number of time steps  $T$ , when  $T \geq L$ . Our regret bound is the first theoretical guarantee on a Thompson sampling algorithm for the standard cascading bandit problem model, or for any stochastic combinatorial bandit problem model with partial feedback (see literature review in Section 1.2).

Our consideration of Gaussian Thompson sampling is primarily motivated by Zong et al. (2016), who report the empirical effectiveness of Gaussian Thompson sampling on cascading bandits, and raise its theoretical analysis as an open question. In this paper, we answer this open question under both the standard setting and its linear generalization, and overcome numerous analytical challenges. We carefully design estimates on the latent mean reward (see inequality (4.2)) to handle the subtle statistical dependencies between partial monitoring and Thompson sampling. We reconcile the statistical inconsistency in using Gaussians to model click probabilities by considering a certain truncated version of the Thompson samples (Lemma 4.5). Our framework provides useful tools for analyzing Thompson sampling on stochastic combinatorial bandits with partial feedback in other settings.

Subsequently, we extend TS-CASCADE to LINTS-CASCADE( $\lambda$ ), a Thompson sampling algorithm for the linear cascading bandit problem and derive a problem-independent upper bound on its regret. According to Kveton et al. (2015a) and our own results, the regret on standard cascading bandits grows linearly with  $\sqrt{L}$ . Hence, the standard cascading bandit problem is intractable when  $L$  is large, which is typical in real-world applications. This motivates us to consider linear cascading bandit problem (Zong et al., 2016). Zong et al. (2016) introduce CASCADELINTS, an implementation of Thompson sampling in the linear case, but present no analysis. In this paper, we propose LINTS-CASCADE( $\lambda$ ), a Thompson sampling algorithm different from CASCADELINTS, and upper bound its regret. We are the first to theoretically establish a regret bound for Thompson sampling algorithm on the linear cascading bandit problem.

Finally, we derive a problem-independent lower bound on the regret incurred by any online algorithm for the standard cascading bandit problem. Note that a problem-dependent one is provided in Kveton et al. (2015a).

## 1.2 Literature Review

Our work is closely related to existing works on the class of stochastic combinatorial bandit (SCB) problems and Thompson sampling. In an SCB model, an arm corresponds to a subset of a ground set of items, each associated with a latent random variable. The corresponding reward depends on the constituent items’ realized random variables. SCB models with *semi-bandit feedback*, where a learning agent observes all random variables of the items in a pulled arm, are extensively studied in existing works. Assuming semi-bandit feedback,

Anantharam et al. (1987) study the case when the arms constitute a uniform matroid, Kveton et al. (2014) study the case of general matroids, Gai et al. (2010) study the case of permutations, and Gai, Chen et al. (2013), Combes et al. (2015), and Kveton et al. (2015b) investigate various general SCB problem settings. More general settings with contextual information (Li et al., 2010; Qin et al., 2014) and linear generalization (Wen et al., 2015; Agrawal and Goyal, 2013b) are also studied. When Wen et al. (2015) consider both UCB and TS algorithms, Agrawal and Goyal (2013b) focus on TS algorithms; all other works above hinge on UCBs.

Motivated by numerous applications in recommender systems and online advertisement, SCB models have been studied under a more challenging setting of *partial feedback*, where a learning agent only observes the random variables for a subset of the items in the pulled arm. A prime example of SCB model with partial feedback is the cascading bandit model, which is first introduced by Kveton et al. (2015a). Subsequently, Kveton et al. (2015c), Katariya et al. (2016), Lagrée et al. (2016) and Zoghi et al. (2017) study the cascading bandit model in various general settings. Cascading bandits with contextual information (Li et al., 2016) and linear generalization (Zong et al., 2016) are also studied. Wang and Chen (2017) provide a general algorithmic framework on SCB models with partial feedback. All of the works listed above only involve the analysis of UCB algorithms.

On the one hand, UCB has been extensively applied for solving various SCB problems. On the other hand, *Thompson sampling* (Thompson, 1933; Chapelle and Li, 2011; Russo et al., 2018), an online algorithm based on Bayesian updates, has been shown to be empirically superior compared to UCB and  $\epsilon$ -greedy algorithms in various bandit models. The empirical success has motivated a series of research works on the theoretical performance guarantees of Thompson sampling on multi-armed bandits (Agrawal and Goyal, 2012; Kaufmann et al., 2012; Agrawal and Goyal, 2017b,a), linear bandits (Agrawal and Goyal, 2013a), generalized linear bandits (Abeille and Lazaric, 2017), etc. Thompson sampling has also been studied for SCB problems with semi-bandit feedback. Komiyama et al. (2015) study the case when the combinatorial arms constitute a uniform matroid; Wang and Chen (2018) investigate the case of general matroids, and Gopalan et al. (2014) and Huyuk and Tekin (2019) consider settings with general reward functions. In addition, SCB problems with semi-bandit feedback are also studied in the Bayesian setting (Russo and Van Roy, 2014), where the latent model parameters are assumed to be drawn from a known prior distribution. Despite existing works, an analysis of Thompson sampling for an SCB problem in the more challenging case of partial feedback is yet to be done. Our work fills in this gap in the literature, and our analysis provides tools for handling the statistical dependence between Thompson sampling and partial feedback in the cascading bandit models.

This paper serves as a significant extension of Cheung et al. (2019), which was presented at the 22nd International Conference on Artificial Intelligence and Statistics in April 2019. The additional materials contained herein include the analysis of the CTS algorithm, LINTS-CASCADE( $\lambda$ ), an algorithm designed for the linear cascading bandit setting, a lower bound on the regret over all cascading bandit algorithms, and extensive experimental results.

### 1.3 Outline

In Section 2, we formally describe the setup of the standard and the linear cascading bandit problems. In Section 3, we provide an upper bound on the regret of CTS algorithm and a proof sketch. In Section 4, we propose our first algorithm TS-CASCADE for the standard problem, and present an asymptotic upper bound on its regret. Additionally, we compare our algorithm to the state-of-the-art algorithms in terms of their regret bounds. We also provide an outline to prove a regret bound of TS-CASCADE. In section 5, we design LINTS-CASCADE( $\lambda$ ) for the linear setting, provide an upper bound on its regret and the proof sketch. In Section 6, we provide a lower bound on the regret of any algorithm in the cascading bandits with its proof sketch. In Section 7, we present comprehensive experimental results that compare the performances of all the algorithms including new ones such as CTS and LINTS-CASCADE( $\lambda$ ). We discuss future work and conclude in Section 8. Details of proofs as given in Appendix C.

## 2. Problem Setup

Let there be  $L \in \mathbb{N}$  ground items, denoted as  $[L] := \{1, \dots, L\}$ . Each item  $i \in [L]$  is associated with a weight  $w(i) \in [0, 1]$ , signifying the item’s click probability. In the standard cascading bandit problem, the click probabilities  $w(1), \dots, w(L)$  are not known to the agent. The agent needs to learn these probabilities and construct an accurate estimate for each of them individually. In the linear generalization setting, the agent possesses certain linear contextual knowledge about each of the click probability terms  $w(1), \dots, w(L)$ . The agent knows that, for each item  $i$  it holds that

$$w(i) = \beta^\top x(i),$$

where  $x(i) \in \mathbb{R}^d$  is the feature vector of item  $i$ , and  $\beta \in \mathbb{R}^d$  is a common vector shared by all items. While the agent knows  $x(i)$  for each  $i$ , he does not know  $\beta$ . The problem now reduces to estimating  $\beta \in \mathbb{R}^d$ . The linear generalization setting represents a setting when the agent could harness certain prior knowledge (through the features  $x(1), \dots, x(L)$  and their linear relationship to the respective click probability terms) to learn the latent click probabilities efficiently. Indeed, in the linear generalization setting, the agent could jointly estimate  $w(1), \dots, w(L)$  by estimating  $\beta$ , which could allow more efficient learning than estimating  $w(1), \dots, w(L)$  individually when  $d < L$ . This linear parameterization is a generalization of the standard case in the sense that when  $\{x(i)\}_{i=1}^L$  is the set of  $(d = L)$ -dimensional standard basis vectors,  $\beta$  is a length- $L$  vector filled with click probabilities, i.e.,  $\beta = [w(1), w(2), \dots, w(L)]^\top$ , then the problem reverts to the standard case. This standard basis case actually represents the case when the  $x(i)$  does not carry useful information.

At each time step  $t \in [T]$ , the agent recommends a list of  $K \leq L$  items  $S_t := (i_1^t, \dots, i_K^t) \in [L]^{(K)}$  to the user, where  $[L]^{(K)}$  denotes the set of all  $K$ -permutations of  $[L]$ . The user examines the items from  $i_1^t$  to  $i_K^t$  by examining each item one at a time, until one item is clicked or all items are examined. For  $1 \leq k \leq K$ ,  $W_t(i_k^t) \sim \text{Bern}(w(i_k^t))$  are i.i.d. and  $W_t(i_k^t) = 1$  iff user clicks on  $i_k^t$  at time  $t$ .

The *instantaneous reward* of the agent at time  $t$  is

$$R(S_t|\mathbf{w}) := 1 - \prod_{k=1}^K (1 - W_t(i_k^t)) \in \{0, 1\}.$$

In other words, the agent gets a reward of  $R(S_t|\mathbf{w}) = 1$  if  $W_t(i_k^t) = 1$  for some  $1 \leq k \leq K$ , and a reward of  $R(S_t|\mathbf{w}) = 0$  if  $W_t(i_k^t) = 0$  for all  $1 \leq k \leq K$ .

The *feedback* from the user at time  $t$  is defined as

$$k_t := \min\{1 \leq k \leq K : W_t(i_k^t) = 1\},$$

where we assume that the minimum over an empty set is  $\infty$ . If  $k_t < \infty$ , then the agent observes  $W_t(i_k^t) = 0$  for  $1 \leq k < k_t$ , and also observes  $W_t(i_{k_t}^t) = 1$ , but does not observe  $W_t(i_k^t)$  for  $k > k_t$ ; otherwise,  $k_t = \infty$ , then the agent observes  $W_t(i_k^t) = 0$  for  $1 \leq k \leq K$ . The agent uses an *online algorithm*  $\pi$  to determine, based on the observation history, the arm  $S_t$  to pull at each time step  $t$ ; in other words, the random variable  $S_t$  is  $\mathcal{F}_{t-1}$ -measurable, where  $\mathcal{F}_t := \sigma(S_1, k_1, \dots, S_t, k_t)$ .

As the agent aims to maximize the sum of rewards over all steps, the expected cumulative regret is defined to evaluate the performance of an algorithm. First, the *expected instantaneous reward* is

$$r(S|\mathbf{w}) = \mathbb{E}[R(S|\mathbf{w})] = 1 - \prod_{i \in S} (1 - w(i)).$$

Note that the expected reward is permutation invariant, but the set of observed items is not. Without loss of generality, we assume that  $w(1) \geq w(2) \geq \dots \geq w(L)$ . Consequently, any permutation of  $\{1, \dots, K\}$  maximizes the expected reward. We let  $S^* = (1, \dots, K)$  be the optimal ordered  $K$ -subset that maximizes the expected reward. Additionally, we let items in  $S^*$  be optimal items and others be suboptimal items. In  $T$  steps, we aim to minimize the *expected cumulative regret*:

$$\text{Reg}(T) := T \cdot r(S^*|\mathbf{w}) - \sum_{t=1}^T r(S_t|\mathbf{w}),$$

while the vector of click probabilities  $\mathbf{w} \in [0, 1]^L$  is not known to the agent, and the recommendation list  $S_t$  is chosen online, i.e., dependent on the previous choices and previous rewards.

### 3. The COMBINATORIAL THOMPSON SAMPLING (CTS) Algorithm

In this section, we analyze a Thompson sampling algorithm for the cascading bandit problem and consider the posterior sampling based on the Beta-Bernoulli conjugate pair. The algorithm, dubbed CTS algorithm, was first proposed as Algorithm 6 in Section 7.3 of Russo et al. (2018) and analyzed by Huyuk and Tekin (2019). We derive a tighter bound on the regret than theirs. The algorithm and the bound are as follows.

---

**Algorithm 1** CTS, Thompson Sampling for Cascading Bandits with Beta Update
 

---

- 1: Initialize  $\alpha_1^B(i) = 1$  and  $\beta_1^B(i) = 1$  for all  $i \in [L]$ .
  - 2: **for**  $t = 1, 2, \dots$  **do**
  - 3:     Generate Thompson sample  $\theta_t^B(i) \sim \text{Beta}(\alpha_t^B(i), \beta_t^B(i))$  for all  $i \in [L]$ .
  - 4:     **for**  $k \in [K]$  **do**
  - 5:         Extract  $i_k^t \in \text{argmax}_{i \in [L] \setminus \{i_1^t, \dots, i_{k-1}^t\}} \theta_t^B(i)$ .
  - 6:     **end for**
  - 7:     Pull arm  $S_t = (i_1^t, i_2^t, \dots, i_K^t)$ .
  - 8:     Observe click  $k_t \in \{1, \dots, K, \infty\}$ .
  - 9:     For each  $i \in [L]$ , if  $W_t(i)$  is observed, define  $\alpha_{t+1}^B(i) = \alpha_t^B(i) + W_t(i)$ ,  $\beta_{t+1}^B(i) = \beta_t^B(i) + 1 - W_t(i)$ . Otherwise,  $\alpha_{t+1}^B(i) = \alpha_t^B(i)$ ,  $\beta_{t+1}^B(i) = \beta_t^B(i)$ .
  - 10: **end for**
- 

**Theorem 3.1 (Problem-dependent bound for CTS)** *Assume  $w(1) \geq \dots \geq w(K) > w(K+1) \geq \dots \geq w(L)$  and  $\Delta_{i,j} = w(i) - w(j)$  for  $i \in S^*$ ,  $j \notin S^*$ , the expected cumulative regret of CTS is bounded as:*

$$\begin{aligned} \text{Reg}(T) \leq & L - K + 16 \cdot \sum_{i=K+1}^L \left( \frac{1}{\Delta_{K,i}^2} + \frac{\log T}{\Delta_{K,i}} \right) + \sum_{i=K+1}^L \sum_{j=1}^K \Delta_{j,i} \\ & + \alpha_2 \cdot \sum_{j=1}^K \Delta_{j,L} \left( \frac{1[\Delta_{j,K+1} < \frac{1}{4}]}{\Delta_{j,K+1}^3} + \frac{\log T}{\Delta_{j,K+1}^2} \right), \end{aligned}$$

where  $\alpha_2$  is a constant independent of the problem instance.

We remark that it is common to involve a problem-independent constant such as  $\alpha_2$  in Theorem 3.1, in the upper bound (Wang and Chen, 2018). In fact, the bounds in Theorems 3.4, 4.1, and 5.1 also involve problem-independent constants.

Table 1: Problem-dependent upper bounds on the  $T$ -regret of CTS, CASCADEUCB1 and CASCADEKL-UCB; the problem-dependent lower bound apply to *all* algorithms for the cascading bandit model. In the table, we use the notation  $\text{KL}(a, b)$  to denote the Kullback-Leibler divergence between two Bernoulli distributions with parameters  $a$  and  $b$ .

| Algorithm/Setting | Reference              | Bounds   |
|-------------------|------------------------|--|
| CTS               | Theorem 3.1            | $O(L(\log T)/\Delta + L/\Delta^2)$                             |
| CTS               | Huyuk and Tekin (2019) | $O(KL(\log T)/[(1 - w(L))^{K-1}w(L)\Delta])$                   |
| CASCADEUCB1       | Kveton et al. (2015a)  | $O((L - K)(\log T)/\Delta + L)$                                |
| CASCADEKL-UCB     | Kveton et al. (2015a)  | $O((L - K)\Delta(\log T)\log(1/\Delta)/\text{KL}(w(L), w(1)))$ |
| Cascading Bandits | Kveton et al. (2015a)  | $\Omega((L - K)(\log T)/\Delta)$ (Lower Bd)                    |

Next, in Table 1, we compare our regret bound for the CTS algorithm to those in the existing literature. Note that the upper bound for CTS in Huyuk and Tekin (2019) can be specified to the cascading bandits as above, and the two other algorithms are based on the UCB idea (Kveton et al., 2015a). To present the results succinctly, we consider the problem where  $w(1) = w(2) = \dots = w(K)$ ,  $w(K+1) = w(K+2) = \dots = w(L)$ , i.e.,  $\Delta := \Delta_{j,i}$  for any optimal item  $j$  and any suboptimal item  $i$ . In other words, the gap between any optimal item  $j$  and any suboptimal item  $i$  is  $\Delta$ , which does not depend on  $i$  and  $j$ . The parameter  $\Delta$  measures the difficulty of the problem.

Table 1 implies that our upper bound grows like  $\log T$ , similarly to the other bounds for the other algorithms. When  $\log T \geq 1/\Delta$ , since  $L(\log T)/\Delta$  is the dominant term in our bound, we can see that  $L(\log T)/\Delta$  is smaller than the term  $KL(\log T)/[(1 - w(L))^{K-1}w(L)\Delta]$  in the bound by Huyuk and Tekin (2019). In addition, the dominant term  $L(\log T)/\Delta$  in our bound matches that in CASCADEUCB1 in the term involving  $\log T$ . However, we note that it is incomparable to the one for CASCADEKL-UCB (Kveton et al., 2015a) due to the following two reasons:

(i) on the one hand,

$$\frac{1}{\Delta \cdot \min\{w(1), 1 - w(1)\}} \leq \frac{\Delta}{\text{KL}(w(L), w(1))} \leq \frac{2}{\Delta}$$

by Pinsker's inequality (see Section 6 for the definition of the KL divergence and Pinsker's inequality);

(ii) on the other hand, CASCADEKL-UCB involves an additional  $\log(1/\Delta)$  term.

Furthermore, the dominant term involving  $\log T$  in our bound matches the lower bound provided in Kveton et al. (2015a) when  $L \gg K$ . Therefore, our derived bound on CTS, the Thompson sampling algorithm with Beta-Bernoulli update, is tight.

Moreover, apart from the problem-dependent bound, we also notice that it is valuable to derive a problem-independent bound for CTS. However, in the cascading bandit model, it cannot be trivially reduced from a problem-dependent one, and hence we leave it as an open question. For instance, Kveton et al. (2015a) (the first paper on cascading bandits) provided an problem-dependent regret bound, but left the problem-independent bound as an open question.

To derive an upper bound, the analytical challenge is to carefully handle the partial feedback from the user in the cascading bandits. Though our analysis has similarity to those in Wang and Chen (2018) (semi-bandit feedback) and Huyuk and Tekin (2019) (partial feedback), the difference is manifested in the following aspects:

- (i) To handle the partial feedback from the user, we handcraft the design of events  $\mathcal{A}_t(\ell, j, k)$ ,  $\tilde{\mathcal{A}}_t(\ell, j, k)$ ,  $\mathcal{A}_t(\ell, j)$  in (3.1), and apply the property of the reward to decompose the regret in terms of the observation of a suboptimal items in Lemma 3.2 and 3.3.
- (ii) Additionally, our algorithm crucially requires the items in  $S_t$  to be in the descending order of  $\theta_t^B(i)$  (see Line 4–7 of Algorithm 1), which is an important algorithmic property that enables our analysis (see Lemma 3.3).



Our analytical framework, which judiciously takes into the impact of partial feedback, is substantially different from Wang and Chen (2018), which assume semi-feedback and do not face the challenge of partial feedback. Besides, since Huyuk and Tekin (2019) consider a less specific problem setting, they define more complicated events than in (3.1) to decompose the regret. Moreover, we carefully apply the mechanism of CTS algorithm to derive Lemma 3.3, which leads to a tight bound on the regret. Overall, our concise analysis fully utilizes the properties of cascading bandits and Algorithm 1.

Now we outline the proof of Theorem 3.1; some details are deferred to Appendix C. Note that a pull of any suboptimal item contributes to the loss of expected instantaneous reward. Meanwhile, we choose  $S_t$  according to the Thompson samples  $\theta_t^{\text{B}}(i)$  instead of the unknown click probability  $w(i)$ . We expect that as time progresses, the empirical mean  $\hat{w}_t(i)$  will approach  $w(i)$  and the Thompson sample  $\theta_t^{\text{B}}(i)$  which is drawn from  $\text{Bern}(\theta_t^{\text{B}}(i))$  will concentrate around  $w(i)$  with high probability. In addition, this probability depends heavily on the number of observations of the item. Therefore, the number of observations of each item affects the distribution of the Thompson samples and hence, also affects the instantaneous rewards. Inspired by this observation, we decompose the the regret in terms of the number of observations of the suboptimal items in the chain of equations including (3.5) to follow.

We first define  $\pi_t$  as follows. For any  $k$ , if the  $k$ -th item in  $S_t$  is optimal, we place this item at position  $k$ ,  $\pi_t(k) = i_k^t$ . The remaining optimal items are positioned arbitrarily. Since  $S^*$  is optimal with respect to  $w$ ,  $w(i_k^t) \leq w(\pi_t(k))$  for all  $k$ . Similarly, since  $S_t$  is optimal with respect to  $\theta_t^{\text{B}}$ ,  $\theta_t^{\text{B}}(i_k^t) \geq \theta_t^{\text{B}}(\pi_t(k))$  for all  $k$ .

Next, for any optimal item  $j$  and suboptimal item  $\ell$ , index  $1 \leq k \leq K$ , time step  $t$ , we say  $\mathcal{A}_t(\ell, j, k)$  holds when suboptimal item  $\ell$  is the  $k$ -th item to pull, optimal item  $j$  is not to be pulled, and  $\pi_t$  places  $j$  at position  $k$ . We say  $\tilde{\mathcal{A}}_t(\ell, j, k)$  holds when  $\mathcal{A}_t(\ell, j, k)$  holds and item  $\ell$  is observed. Additionally,  $\mathcal{A}_t(\ell, j)$  is the union set of events  $\tilde{\mathcal{A}}_t(\ell, j, k)$  for  $k = 1, \dots, K$ . We can also write these events as follows:

$$\begin{aligned} \mathcal{A}_t(\ell, j, k) &= \{i_k^t = \ell, \pi_t(k) = j\}, \\ \tilde{\mathcal{A}}_t(\ell, j, k) &= \{i_k^t = \ell, \pi_t(k) = j, W_t(i_1^t) = \dots = W_t(i_{k-1}^t) = 0\} \\ &= \{i_k^t = \ell, \pi_t(k) = j, \text{Observe } W_t(\ell) \text{ at } t\}, \\ \mathcal{A}_t(\ell, j) &= \bigcup_{k=1}^K \tilde{\mathcal{A}}_t(\ell, j, k) = \{\exists 1 \leq k \leq K \text{ s.t. } i_k^t = \ell, \pi_t(k) = j, \text{Observe } W_t(\ell) \text{ at } t\}. \end{aligned} \tag{3.1}$$

Then we have

$$\begin{aligned} r(S^*|\mathbf{w}) - r(S_t|\mathbf{w}) \\ = r(\pi_t(1), \dots, \pi_t(K)|\mathbf{w}) - r(S_t|\mathbf{w}) \end{aligned} \tag{3.2}$$

$$\begin{aligned} &= \sum_{k=1}^K \left[ \prod_{m=1}^{k-1} [1 - w(i_m^t)] \right] \cdot [w(\pi_t(k)) - w(i_k^t)] \cdot \left[ \prod_{n=k+1}^K (1 - w(\pi_t(n))) \right] \\ &\leq \sum_{k=1}^K \left[ \prod_{m=1}^{k-1} (1 - w(i_m^t)) \right] \cdot [w(\pi_t(k)) - w(i_k^t)] \end{aligned} \tag{3.3}$$

$$\begin{aligned}
 &\leq \sum_{k=1}^K \sum_{j=1}^K \sum_{\ell=K+1}^L \mathbb{E} \left[ \left[ \prod_{m=1}^{k-1} (1 - w(i_m^t)) \right] \cdot [w(\pi_t(k)) - w(i_k^t)] \mid \mathcal{A}_t(\ell, j, k) \right] \cdot \mathbb{E} [1 \{ \mathcal{A}_t(\ell, j, k) \}] \\
 &= \sum_{j=1}^K \sum_{\ell=K+1}^L \sum_{k=1}^K \Delta_{j,\ell} \cdot \mathbb{E} [1 \{ W_t(i_1^t) = \dots = W_t(i_{k-1}^t) = 0 \} \mid \mathcal{A}_t(\ell, j, k)] \cdot \mathbb{E} [1 \{ \mathcal{A}_t(\ell, j, k) \}] \\
 &= \sum_{j=1}^K \sum_{\ell=K+1}^L \Delta_{j,\ell} \cdot \sum_{k=1}^K \mathbb{E} [1 \{ \tilde{\mathcal{A}}_t(\ell, j, k) \}] \\
 &= \sum_{j=1}^K \sum_{\ell=K+1}^L \Delta_{j,\ell} \cdot \mathbb{E} [1 \{ \mathcal{A}_t(\ell, j) \}].
 \end{aligned}$$

Line (3.2) results from the fact that any permutation of  $\{1, \dots, K\}$  maximizes the expected reward and  $(\pi_t(1), \dots, \pi_t(K))$  is one permutation of  $\{1, \dots, K\}$ . Line (3.3) applies the property of the reward (Lemma C.4). The subsequent ones follow from the various events defined in (3.1).

Let  $N^{\text{B}}(i, j)$  denote the number of time steps that  $\mathcal{A}_t(i, j)$  happens for optimal item  $j$  and suboptimal item  $i$  within the whole horizon, then

$$N^{\text{B}}(i, j) = \sum_{t=1}^T \mathbb{E} [1 \{ \mathcal{A}_t(i, j) \}], \quad (3.4)$$

$$\text{Reg}(T) = T \cdot r(S^* | \mathbf{w}) - \sum_{t=1}^T r(S_t | \mathbf{w}) \leq \sum_{j=1}^K \sum_{i=K+1}^L \Delta_{j,i} \cdot \mathbb{E} [N^{\text{B}}(i, j)]. \quad (3.5)$$

We use

$$\hat{w}_t(i) = \frac{\alpha_t^{\text{B}}(i) - 1}{\alpha_t^{\text{B}}(i) + \beta_t^{\text{B}}(i) - 2} = \frac{\sum_{\tau=1}^{t-1} 1 \{ \text{item } i \text{ is observed at time step } \tau \}}{\sum_{\tau=1}^{t-1} 1 \{ \text{item } i \text{ is pulled at time step } \tau \}} \quad (3.6)$$

to denote the empirical mean of item  $i$  up to and including time step  $t - 1$ . To bound the expectation of  $N^{\text{B}}(i, j)$ , we define two more events taking into account the concentration of Thompson sample  $\theta_t^{\text{B}}(i)$  and empirical mean  $\hat{w}_t(i)$ :

$$\mathcal{B}_t(i) = \{ \hat{w}_t(i) > w(i) + \varepsilon_i \}, \quad \mathcal{C}_t(i, j) = \{ \theta_t^{\text{B}}(i) > w(j) - \varepsilon_j \}, \quad 1 \leq j \leq K < i \leq L.$$

with

$$\varepsilon_i = \frac{1}{2} \cdot \left( \Delta_{K+1,i} + \frac{\Delta_{K,K+1}}{2} \right), \quad \varepsilon_j = \frac{1}{2} \cdot \left( \Delta_{j,K} + \frac{\Delta_{K,K+1}}{2} \right), \quad 1 \leq j \leq K < i \leq L.$$

Then  $\varepsilon_i \geq \Delta_{K,i}/4$ ,  $\varepsilon_j \geq \Delta_{j,K+1}/4$ ,  $\varepsilon_i + \varepsilon_j = \Delta_{j,i}/2$ , and we have

$$\begin{aligned}
 \mathbb{E} [N^{\text{B}}(i, j)] &= \mathbb{E} \left[ \sum_{t=1}^T 1 [\mathcal{A}_t(i, j)] \right] \\
 &\leq \mathbb{E} \left[ \sum_{t=1}^T 1 [\mathcal{A}_t(i, j) \wedge \mathcal{B}_t(i)] \right] + \mathbb{E} \left[ \sum_{t=1}^T 1 [\mathcal{A}_t(i, j) \wedge \neg \mathcal{B}_t(i) \wedge \mathcal{C}_t(i, j)] \right] \\
 &\quad + \mathbb{E} \left[ \sum_{t=1}^T 1 [\mathcal{A}_t(i, j) \wedge \neg \mathcal{C}_t(i, j)] \right].
 \end{aligned} \quad (3.7)$$

Subsequently, we bound the three terms in (3.7) separately. Since we need to count the number of observations of items in the subsequent analysis, we also use  $N_t(i)$  to denote one plus the number of observations of item  $i$  before time  $t$  (consistent with other sections).

**The first term.** We first derive a lemma which facilitates us to bound the probability the event  $\mathcal{A}_t(i, j) \wedge \mathcal{B}_t(i)$ .

**Lemma 3.2** *For any suboptimal item  $K < i \leq L$ ,*

$$\mathbb{E} \left[ \sum_{t=1}^T \mathbf{1}[i \in S_t, |\hat{w}_t(i) - w(i)| > \varepsilon_i, \text{Observe } W_t(i) \text{ at } t] \right] \leq 1 + \frac{1}{\varepsilon_i^2}.$$

Next, summing over all possible pairs  $(i, j)$  and then applying Lemma 3.2, the regret incurred by the first term can be bounded as follows:

$$\begin{aligned} & \sum_{j=1}^K \sum_{i=K+1}^L \Delta_{j,i} \cdot \mathbb{E} \left[ \sum_{t=1}^T \mathbf{1}[\mathcal{A}_t(i, j) \wedge \mathcal{B}_t(i)] \right] \\ & \leq \sum_{i=K+1}^L \mathbb{E} \left[ \sum_{t=1}^T \sum_{j=1}^K \mathbf{1}[\mathcal{A}_t(i, j) \wedge \mathcal{B}_t(i)] \right] \\ & \leq \sum_{i=K+1}^L \mathbb{E} \left[ \sum_{t=1}^T \mathbf{1}[\mathcal{B}_t(i), \text{Observe } W_t(i) \text{ at } t] \right] \\ & \leq \sum_{i=K+1}^L \mathbb{E} \left[ \sum_{t=1}^T \mathbf{1}[i \in S_t, |\hat{w}_t(i) - w(i)| > \varepsilon_i, \text{Observe } W_t(i) \text{ at } t] \right] \\ & \leq L - K + \sum_{i=K+1}^L \frac{1}{\varepsilon_i^2} \leq L - K + \sum_{i=K+1}^L \frac{16}{\Delta_{K,i}^2}. \end{aligned} \tag{3.8}$$

**The second term.** Considering the occurrence of  $\mathcal{A}_t(i, j) \wedge \neg \mathcal{B}_t(i) \wedge \mathcal{C}_t(i, j)$ , we bound the regret incurred by the second term as follows (see Appendix C.2 for details):

$$\sum_{i=K+1}^L \sum_{j=1}^K \Delta_{j,i} \cdot \mathbb{E} \left[ \sum_{t=1}^T \mathbf{1}[\mathcal{A}_t(i, j) \wedge \neg \mathcal{B}_t(i) \wedge \mathcal{C}_t(i, j)] \right] \leq \sum_{i=K+1}^L \frac{16 \log T}{\Delta_{K,i}} + \sum_{i=K+1}^L \sum_{j=1}^K \Delta_{j,i}. \tag{3.9}$$

**The third term.** We use  $\theta^{\text{B}}(-i)$  to denote the vector  $\theta^{\text{B}}$  without the  $i$ -th element  $\theta^{\text{B}}(i)$ . For any optimal item  $j$ , let  $W_{i,j}$  be the set of all possible values of  $\theta^{\text{B}}$  satisfying that  $\bar{\mathcal{A}}_{i,j}(t) \wedge \neg \mathcal{C}_t(i, j)$  happens for some  $i$ , and  $W_{i,-j} = \{\theta^{\text{B}}(-j) : \theta^{\text{B}} \in W_{i,j}\}$ . We now state Lemma 3.3 whose proof depends on a delicate design of Algorithm 1: to pull items in their descending order of Thompson samples  $\theta_t^{\text{B}}(i)$ .

**Lemma 3.3** *For any optimal item  $1 \leq j \leq K$ , we have*

$$\mathbb{E} \left[ \sum_{t=1}^T \Pr \left[ \exists i \text{ s.t. } \theta_t^{\text{B}} \in W_{i,j}, \text{Observe } W_t(i) \text{ at } t \mid \mathcal{F}_{t-1} \right] \right] \leq c \cdot \left( \frac{\mathbf{1}[\varepsilon_j < 1/16]}{\varepsilon_j^3} + \frac{\log T}{\varepsilon_j^2} \right),$$

where  $c$  is a constant which does not depend on the problem instance.

Recalling the definition of  $W_{i,j}$ , we bound the regret incurred by the third term

$$\begin{aligned}
 & \sum_{i=K+1}^L \sum_{j=1}^K \Delta_{j,i} \cdot \mathbb{E} \left[ \sum_{t=1}^T 1[\mathcal{A}_t(i,j) \wedge \neg \mathcal{C}_t(i,j)] \right] \\
 & \leq \sum_{j=1}^K \Delta_{j,L} \cdot \mathbb{E} \left[ \sum_{t=1}^T \sum_{i=K+1}^L 1[\mathcal{A}_t(i,j) \wedge \neg \mathcal{C}_t(i,j)] \right] \\
 & = \sum_{j=1}^K \Delta_{j,L} \cdot \mathbb{E} \left[ \sum_{t=1}^T \sum_{i=K+1}^L 1[\boldsymbol{\theta}_t^B \in W_{i,j}, \text{Observe } W_t(i) \text{ at } t] \right] \\
 & = \sum_{j=1}^K \Delta_{j,L} \cdot \mathbb{E} \left[ \sum_{t=1}^T 1[\exists i \text{ s.t. } \boldsymbol{\theta}_t^B \in W_{i,j}, \text{Observe } W_t(i) \text{ at } t] \right] \\
 & \leq c \cdot \sum_{j=1}^K \Delta_{j,L} \cdot \left( \frac{1[\varepsilon_j < 1/16]}{\varepsilon_j^3} + \frac{\log T}{\varepsilon_j^2} \right) \leq \alpha_2 \cdot \sum_{j=1}^K \Delta_{j,L} \cdot \left( \frac{1[\Delta_{j,K+1} < 1/4]}{\Delta_{j,K+1}^3} + \frac{\log T}{\Delta_{j,K+1}^2} \right)
 \end{aligned} \tag{3.10}$$

with  $\alpha_2 = 64c$  (a constant which does not depend on the problem instance).

**Summation of all the terms.** Overall, we derive an upper bound on the total regret:

$$\begin{aligned}
 \text{Reg}(T) & \leq L - K + \sum_{i=K+1}^L \frac{16}{\Delta_{K,i}^2} + \sum_{i=K+1}^L \frac{16 \log T}{\Delta_{K,i}} + \sum_{i=K+1}^L \sum_{j=1}^K \Delta_{j,i} \\
 & \quad + \alpha_2 \cdot \sum_{j=1}^K \Delta_{j,L} \left( \frac{1[\Delta_{j,K+1} < \frac{1}{4}]}{\Delta_{j,K+1}^3} + \frac{\log T}{\Delta_{j,K+1}^2} \right) \\
 & \leq L - K + 16 \cdot \sum_{i=K+1}^L \left( \frac{1}{\Delta_{K,i}^2} + \frac{\log T}{\Delta_{K,i}} \right) + \sum_{i=K+1}^L \sum_{j=1}^K \Delta_{j,i} \\
 & \quad + \alpha_2 \cdot \sum_{j=1}^K \Delta_{j,L} \left( \frac{1[\Delta_{j,K+1} < \frac{1}{4}]}{\Delta_{j,K+1}^3} + \frac{\log T}{\Delta_{j,K+1}^2} \right).
 \end{aligned}$$

The previous upper bound on the regret of the CTS algorithm in Theorem 3.1 is problem-dependent, i.e., it depends on the gaps  $\{\Delta_{i,j}\}_{i,j}$ . We now present a problem-independent bound on the regret of CTS.

**Theorem 3.4 (Problem-independent bound for CTS)** *Assume that  $w(1) = w(2) = \dots = w(K)$ . Algorithm CTS incurs an expected regret at most*

$$O(KT^{3/4} + L\sqrt{T} + K\sqrt{T} \log T),$$

where the big  $O$  notation hides a constant factor that is independent of  $K, L, T, \mathbf{w}$ .

We remark that if all variables other than  $T$  are kept fixed, the expected regret then scales as  $O(KT^{3/4})$ . The proof of Theorem 3.4 can be found in Appendix C.4.

Theorem 3.4 is established by following a case analysis on whether  $\Delta_{1,i} \leq \varepsilon$  or  $\Delta_{1,i} > \varepsilon$  for each suboptimal item  $i$ , and then tuning the parameter  $\varepsilon$  to optimize the bound. The case analysis is in a similar vein to the reduction from problem-dependent regret bounds to problem-independent regret bounds in the classical stochastic MAB literature (for example, see Theorem 6 in Kveton et al. (2015b)). While this reduction approach does not yield an optimal bound in terms of  $T$ , in the subsequent section we study a Thompson sampling-like algorithm TS-CASCADE which enjoys a regret bound that scales with  $T$  as  $\sqrt{T}$ . In addition, TS-CASCADE also naturally generalizes to a linear contextual setting.

#### 4. The TS-CASCADE Algorithm

The CTS algorithm cannot be applied directly to the linear generalization setting, since the Beta-Bernoulli update is incompatible with the linear generalization model on the click probabilities (see detailed discussions after Lemma 4.4 for an explanation). Therefore, we propose a Thompson sampling-like algorithm. In the rest of the paper, we call our algorithm TS-CASCADE with a slight abuse of terminology, as we acknowledge that it is not a *bona fide* Thompson sampling algorithm that is based on prior posterior updates. Instead of the Beta distribution, TS-CASCADE designs appropriate Thompson samples that are Gaussian distributed and uses these in the posterior sampling steps. We design a variation of TS-CASCADE for the linear setting in Section 5.

Our algorithm TS-CASCADE is presented in Algorithm 2. Intuitively, to minimize the expected cumulative regret, the agent aims to learn the true weight  $w(i)$  of each item  $i \in [L]$  by exploring the space to identify  $S^*$  (i.e., exploitation) after a hopefully small number of steps. In our algorithm, we approximate the true weight  $w(i)$  of each item  $i$  with a Bayesian statistic  $\theta_t(i)$  at each time step  $t$ . This statistic is known as the *Thompson sample*. To do so, first, we sample a one-dimensional standard Gaussian  $Z_t \sim \mathcal{N}(0, 1)$ , define the empirical variance  $\hat{v}_t(i) = \hat{w}_t(i)(1 - \hat{w}_t(i))$  of each item, and calculate  $\theta_t(i)$ . Secondly, we select  $S_t = (i_1^t, i_2^t, \dots, i_K^t)$  such that  $\theta_t(i_1^t) \geq \theta_t(i_2^t) \geq \dots \geq \theta_t(i_K^t) \geq \max_{j \notin S_t} \theta_t(j)$ ; this is reflected in Line 10 of Algorithm 2. Finally, we update the parameters for each observed item  $i$  in a standard manner by applying Bayes rule on the mean of the Gaussian (with conjugate prior being another Gaussian) in Line 14. The algorithm results in the following theoretical guarantee.

**Theorem 4.1 (Problem-independent)** *Consider the cascading bandit problem. Algorithm TS-CASCADE, presented in Algorithm 2, incurs an expected regret at most*

$$O(\sqrt{KLT} \log T + L \log^{5/2} T).$$

In practical applications,  $T \gg L$  and so the regret bound is essentially  $\tilde{O}(\sqrt{KLT})$ . Besides, we note that in problem-independent bounds for TS-based algorithms, the implied constants in the  $O(\cdot)$ -notation are usually large (see Agrawal and Goyal (2013b, 2017b)). These large constants result from the anti-concentration bounds used in the analysis. In this work, we mainly aim to show the dependence of regret on  $T$  similarly to other papers on Thompson Sampling in which constants are not explicitly shown. Hence, in the spirit of those papers on Thompson Sampling, we made no further attempt to optimize for the best constants in our regret bounds. Nevertheless we show via experiment that in practice,

---

**Algorithm 2** TS-CASCADE, Thompson Sampling for Cascading Bandits with Gaussian Update

---

- 1: Initialize  $\hat{w}_1(i) = 0$ ,  $N_1(i) = 0$  for all  $i \in [L]$ .
  - 2: **for**  $t = 1, 2, \dots$  **do**
  - 3:     Sample a 1-dimensional random variable  $Z_t \sim \mathcal{N}(0, 1)$ .
  - 4:     **for**  $i \in [L]$  **do**
  - 5:         Calculate the empirical variance  $\hat{\nu}_t(i) = \hat{w}_t(i)(1 - \hat{w}_t(i))$ .
  - 6:         Calculate standard deviation of the Thompson sample
 
$$\sigma_t(i) = \max \left\{ \sqrt{\frac{\hat{\nu}_t(i) \log(t+1)}{N_t(i)+1}}, \frac{\log(t+1)}{N_t(i)+1} \right\}.$$
  - 7:         Construct the Thompson sample  $\theta_t(i) = \hat{w}_t(i) + Z_t \sigma_t(i)$ .
  - 8:     **end for**
  - 9:     **for**  $k \in [K]$  **do**
  - 10:         Extract  $i_k^t \in \operatorname{argmax}_{i \in [L] \setminus \{i_1^t, \dots, i_{k-1}^t\}} \theta_t(i)$ .
  - 11:     **end for**
  - 12:     Pull arm  $S_t = (i_1^t, i_2^t, \dots, i_K^t)$ .
  - 13:     Observe click  $k_t \in \{1, \dots, K, \infty\}$ .
  - 14:     For each  $i \in [L]$ , if  $W_t(i)$  is observed, define  $\hat{w}_{t+1}(i) = \frac{N_t(i)\hat{w}_t(i) + W_t(i)}{N_t(i)+1}$  and  $N_{t+1}(i) = N_t(i) + 1$ . Otherwise (if  $W_t(i)$  is not observed),  $\hat{w}_{t+1}(i) = \hat{w}_t(i)$ ,  $N_{t+1}(i) = N_t(i)$ .
  - 15: **end for**
- 

TS-CASCADE’s empirical performance outperforms what the finite-time bounds (with large constants) suggest. We elaborate on the main features of the algorithm and the guarantee.

In a nutshell, TS-CASCADE is a Thompson sampling algorithm (Thompson, 1933), based on prior-posterior updates on Gaussian random variables with refined variances. The use of the Gaussians is useful, since it allows us to readily generalize the algorithm and analyses to the contextual setting (Li et al., 2010, 2016), as well as the linear bandit setting (Zong et al., 2016) for handling a large  $L$ . We consider the linear setting in Section 5 and plan to study the contextual setting in the future. To this end, we remark that the posterior update of TS can be done in a variety of ways. While the use of a Beta-Bernoulli update to maintain a Bayesian estimate on  $w(i)$  is a natural option (Russo et al., 2018), we use Gaussians instead, which is in part inspired by Zong et al. (2016); this is for ease of generalization to the linear bandit setting. Indeed, the conjugate prior-posterior update is not the only choice for TS algorithms for complex multi-armed bandit problems. For example, the posterior update in Algorithm 2 in Agrawal et al. (2017) for the multinomial logit bandit problem is not conjugate.

While the use of Gaussians is useful for generalizations, the analysis of Gaussian Thompson samples in the cascading setting comes with some difficulties, as  $\theta_t(i)$  is not in  $[0, 1]$  with probability one. We perform a truncation of the Gaussian Thompson sample in the proof of Lemma 4.5 to show that this replacement of the Beta by the Gaussian does not incur any significant loss in terms of the regret and the analysis is not affected significantly.

We elaborate on the refined variances of our Bayesian estimates. Lines 5–7 indicate that the Thompson sample  $\theta_t(i)$  is constructed to be a Gaussian random variable with mean  $\hat{w}_t(i)$  and variance being the maximum of  $\hat{v}_t(i) \log(t+1)/(N_t(i)+1)$  and  $[\log(t+1)/(N_t(i)+1)]^2$ . Note that  $\hat{v}_t(i)$  is the variance of a Bernoulli distribution with mean  $\hat{w}_t(i)$ . In Thompson sampling algorithms, the choice of the variance is of crucial importance. We considered a naïve TS implementation initially. However, the theoretical and empirical results were unsatisfactory, due in part to the large variance of the Thompson sample variance; this motivated us to improve on the algorithm leading to Algorithm 2. The reason why we choose the variance in this manner is to (i) make the Bayesian estimates behave like Bernoulli random variables and to (ii) ensure that it is tuned so that the regret bound has a dependence on  $\sqrt{K}$  (see Lemma 4.4) and does not depend on any pre-set parameters. We utilize a key result by Audibert et al. (2009) concerning the analysis of using the empirical variance in multi-arm bandit problems to achieve (i). In essence, in Lemma 4.4, the Thompson sample is shown to depend only on a *single* source of randomness, i.e., the Gaussian random variable  $Z_t$  (Line 3 of Algorithm 2). This shaves off a factor of  $\sqrt{K}$  vis-à-vis a more naïve analysis where the variance is pre-set in the relevant probability in Lemma 4.4 depends on  $K$  independent random variables. In detail, if instead of Line 3–7 in Algorithm 2, we sample independent Thompson samples as follows:

$$\theta_t(i) \sim \mathcal{N}\left(\hat{w}_t(i), \frac{K \log(t+1)}{N_t(i)+1}\right) \quad \forall i \in [L],$$

a similar analysis to our proof of Theorem 4.1 leads to an upper bound  $O(K\sqrt{LT} \log T)$ , which is  $\sqrt{K}$  worse than the upper bound for TS-CASCADE. Hence, we choose to design our Thompson samples based on a single one-dimensional Gaussian random variable  $Z_t$ .

Table 2: Problem-independent upper bounds on the  $T$ -regret of CTS, TS-CASCADE and CUCB; the problem-independent lower bound apply to *all* algorithms for the cascading bandit model.

| <b>Algorithm/Setting</b> | <b>Reference</b>     | <b>Bounds</b>                           |
|--------------------------|----------------------|---|
| CTS (*)                  | Theorem 3.4          | $O(KT^{3/4})$                           |
| TS-CASCADE               | Theorem 4.1          | $O(\sqrt{KLT} \log T + L \log^{5/2} T)$ |
| CUCB                     | Wang and Chen (2017) | $O(\sqrt{KLT} \log T)$                  |
| Cascading Bandits        | Theorem 6.1          | $\Omega(\sqrt{LT/K})$ (Lower Bd)        |

Next, in Table 2, we present various problem-independent bounds on the regret for cascading bandits. The (\*) in Table 2 indicates that our bound for CTS in Theorem 3.4 is derived for a special instance in which all the optimal items have the same click probability, i.e.,  $w(1) = w(2) = \dots = w(K)$ , while the remaining bounds are for *any* general cascading bandit instances. We observe that the bound for CTS is larger than the remaining bounds in the term involving  $T$  by a factor of approximately  $\sqrt[4]{T}$ . While we believe that it is possible to establish a tighter problem-independent bound for CTS, the derivation could

potentially involve analysis significant from the proof of CTS in Theorems 3.1 and 3.4. We leave the possible refinement as an interesting future direction. Subsequently, we focus on comparing the remaining bounds. Table 2 shows that our upper bound for TS-CASCADE grows like  $\sqrt{T}$  just like the other bounds. Our bound for TS-CASCADE also matches the state-of-the-art UCB bound (up to log factors) by Wang and Chen (2017), whose algorithm, when suitably specialized to the cascading bandit setting, is the same as CASCADEUCB1 in Kveton et al. (2015a). For the case in which  $T \geq L$ , our bound for TS-CASCADE is a  $\sqrt{\log T}$  factor worse than the problem-independent bound in Wang and Chen (2017), but we are the first to analyze Thompson sampling for the cascading bandit problem.

Moreover, Table 2 shows that the major difference between the upper bound in Theorem 4.1 and the lower bound in Theorem 6.1 is manifested in their dependence on  $K$  (the length of recommendation list). (i) On one hand, when  $K$  is larger, the agent needs to learn more optimal items. However, instead of the full feedback in the semi-bandit setting, the agent only gets partial feedback in the cascading bandit setting and therefore he is not guaranteed to observe more items at each time step. Therefore, a larger  $K$  may lead to a larger regret. (ii) On the other hand, when  $K$  is smaller, the maximum possible amount of observations that the agent can obtain at each time step diminishes. Hence, the agent may obtain less information from the user, which may also result in a larger regret. Intuitively, the effect of  $K$  on the regret is not clear. We leave the exact dependence of optimal algorithms on  $K$  as an open question.

Sequentially, we present a proof sketch of Theorem 4.1 with several lemmas in order, when some details are postponed to Appendix C.

During the iterations, we update  $\hat{w}_{t+1}(i)$  such that it approaches  $w(i)$  eventually. To do so, we select a set  $S_t$  according to the order of  $\theta_t(i)$ 's at each time step. Hence, if  $\hat{w}_{t+1}(i)$ ,  $\theta_t(i)$  and  $w(i)$  are close enough, then we are likely to select the optimal set. This motivates us to define two ‘‘nice events’’ as follows:

$$\mathcal{E}_{\hat{w},t} := \{\forall i \in [L] : |\hat{w}_t(i) - w(i)| \leq g_t(i)\}, \quad \mathcal{E}_{\theta,t} := \{\forall i \in [L] : |\theta_t(i) - \hat{w}_t(i)| \leq h_t(i)\},$$

where  $\hat{v}_t(i)$  is defined in Line 5 of Algorithm 2, and

$$g_t(i) := \sqrt{\frac{16\hat{v}_t(i) \log(t+1)}{N_t(i)+1}} + \frac{24 \log(t+1)}{N_t(i)+1}, \quad h_t(i) := \sqrt{\log(t+1)}g_t(i).$$

**Lemma 4.2** *For each  $t \in [T]$ ,  $H_t \in \mathcal{E}_{\hat{w},t}$ , we have*

$$\Pr[\mathcal{E}_{\hat{w},t}] \geq 1 - \frac{3L}{(t+1)^3}, \quad \Pr[\mathcal{E}_{\theta,t}|H_t] \geq 1 - \frac{1}{(t+1)^3}.$$

Demonstrating that  $\mathcal{E}_{\hat{w},t}$  has high probability requires the concentration inequality in Theorem B.1; this is a specialization of a result in Audibert et al. (2009) to Bernoulli random variables. Demonstrating that  $\mathcal{E}_{\theta,t}$  has high probability requires the concentration property of Gaussian random variables (cf. Theorem B.2).

To start our analysis, define

$$F(S, t) := \sum_{k=1}^K \left[ \prod_{j=1}^{k-1} (1-w(i_j)) \right] (g_t(i_k) + h_t(i_k)) \tag{4.1}$$



and the the set

$$\mathcal{S}_t := \left\{ S = (i_1, \dots, i_K) \in \pi_K(L) : F(S, t) \geq r(S^* | \mathbf{w}) - r(S | \mathbf{w}) \right\}.$$

Recall that  $w(1) \geq w(2) \geq \dots \geq w(L)$ . As such,  $\mathcal{S}_t$  is non-empty, since  $S^* = (1, 2, \dots, K) \in \mathcal{S}_t$ .

**Intuition behind the set  $\mathcal{S}_t$ :** Ideally, we expect the user to click an item in  $\mathcal{S}_t$  for every time step  $t$ . Recall that  $g_t(i)$  and  $h_t(i)$  are decreasing in  $N_t(i)$ , the number of time steps  $q$ 's in  $1, \dots, t-1$  when we get to *observe*  $W_q(i)$ . Naively, arms in  $\mathcal{S}_t$  can be thought of as arms that “lack observations”, while arms in  $\bar{\mathcal{S}}_t$  can be thought of as arms that are “observed enough”, and are believed to be suboptimal. Note that  $S^* \in \mathcal{S}_t$  is a prime example of an arm that is under-observed.

To further elaborate,  $g_t(i) + h_t(i)$  is the “statistical gap” between the Thompson sample  $\theta_t(i)$  and the latent mean  $w(i)$ . The gap shrinks with more observations of  $i$ . To balance exploration and exploitation, for any suboptimal item  $i \in [L] \setminus [K]$  and any optimal item  $k \in [K]$ , we should have  $g_t(i) + h_t(i) \geq w(k) - w(i)$ . However, this is too much to hope for, and it seems that hoping for  $S_t \in \mathcal{S}_t$  to happen would be more viable. (See the forthcoming Lemma 4.3.)

**Further notations.** In addition to set  $\mathcal{S}_t$ , we define  $H_t$  as the collection of observations of the agent, from the beginning until the end of time  $t-1$ . More precisely, we define  $H_t := \{S_q\}_{q=1}^{t-1} \cup \{(i_k^q, W_q(i_k^q))_{k=1}^{\min\{k_t, \infty\}}\}_{q=1}^{t-1}$ . Recall that  $S_q \in \pi_K(L)$  is the arm pulled during time step  $q$ , and  $(i_k^q, W_q(i_k^q))_{k=1}^{\min\{k_t, \infty\}}$  is the collection of observed items and their respective values during time step  $q$ . At the start of time step  $t$ , the agent has observed everything in  $H_t$ , and determine the arm  $S_t$  to pull accordingly (see Algorithm 2). Note that event  $\mathcal{E}_{\hat{w}, t}$  is  $\sigma(H_t)$ -measurable. For the convenience of discussion, we define  $\mathcal{H}_{\hat{w}, t} := \{H_t : \text{Event } \mathcal{E}_{\hat{w}, t} \text{ is true in } H_t\}$ . The first statement in Lemma 4.2 is thus  $\Pr[H_t \in \mathcal{H}_{\hat{w}, t}] \geq 1 - 3L/(t+1)^3$ .

The performance of Algorithm 2 is analyzed using the following four Lemmas. To begin with, Lemma 4.3 quantifies a set of conditions on  $\hat{\boldsymbol{\mu}}_t$  and  $\boldsymbol{\theta}_t$  so that the pulled arm  $S_t$  belongs to  $\mathcal{S}_t$ , the collection of arms that lack observations and should be explored. We recall from Lemma 4.2 that the events  $\mathcal{E}_{\hat{w}, t}$  and  $\mathcal{E}_{\theta, t}$  hold with high probability. Subsequently, we will crucially use our definition of the Thompson sample  $\boldsymbol{\theta}_t$  to argue that inequality (4.2) holds with non-vanishing probability when  $t$  is sufficiently large.

**Lemma 4.3** *Consider a time step  $t$ . Suppose that events  $\mathcal{E}_{\hat{w}, t}, \mathcal{E}_{\theta, t}$  and inequality*

$$\sum_{k=1}^K \left[ \prod_{j=1}^{k-1} (1 - w(j)) \right] \theta_t(k) \geq \sum_{k=1}^K \left[ \prod_{j=1}^{k-1} (1 - w(j)) \right] w(k) \quad (4.2)$$

*hold, then the event  $\{S_t \in \mathcal{S}_t\}$  also holds.*

In the following, we condition on  $H_t$  and show that  $\boldsymbol{\theta}_t$  is “typical” w.r.t.  $\mathbf{w}$  in the sense of (4.2). Due to the conditioning on  $H_t$ , the only source of the randomness of the pulled arm  $S_t$  is from the Thompson sample. Thus, by analyzing a suitably weighted version of the Thompson samples in inequality (4.2), we disentangle the statistical dependence between

partial monitoring and Thompson sampling. Recall that  $\theta_t$  is normal with  $\sigma(H_t)$ -measurable mean and variance (Lines 5–7 in Algorithm 2).

**Lemma 4.4** *There exists an absolute constant  $c \in (0, 1)$  independent of  $w, K, L, T$  such that, for any time step  $t$  and any historical observation  $H_t \in \mathcal{H}_{\hat{w}, t}$ , the following inequality holds:*

$$\Pr_{\theta_t} [\mathcal{E}_{\theta, t} \text{ and (4.2) hold} \mid H_t] \geq c - \frac{1}{(t+1)^3}.$$

**Proof** We prove the Lemma by setting the absolute constant  $c$  to be  $1/(192\sqrt{\pi}e^{1152}) > 0$ .

For brevity, we define  $\alpha(1) := 1$ , and  $\alpha(k) = \prod_{j=1}^{k-1} (1 - w(j))$  for  $2 \leq k \leq K$ . By the second part of Lemma 4.2, we know that  $\Pr[\mathcal{E}_{\theta, t} \mid H_t] \geq 1 - 1/(t+1)^3$ , so to complete this proof, it suffices to show that  $\Pr[(4.2) \text{ holds} \mid H_t] \geq c$ . For this purpose, consider

$$\begin{aligned} & \Pr_{\theta_t} \left[ \sum_{k=1}^K \alpha(k) \theta_t(k) \geq \sum_{k=1}^K \alpha(k) w(k) \mid H_t \right] \\ &= \Pr_{Z_t} \left[ \sum_{k=1}^K \alpha(k) [\hat{w}_t(k) + Z_t \sigma_t(k)] \geq \sum_{k=1}^K \alpha(k) w(k) \mid H_t \right] \end{aligned} \quad (4.3)$$

$$\geq \Pr_{Z_t} \left[ \sum_{k=1}^K \alpha(k) [w(k) - g_t(k)] + Z_t \cdot \sum_{k=1}^K \alpha(k) \sigma_t(k) \geq \sum_{k=1}^K \alpha(k) w(k) \mid H_t \right] \quad (4.4)$$

$$\begin{aligned} &= \Pr_{Z_t} \left[ Z_t \cdot \sum_{k=1}^K \alpha(k) \sigma_t(k) \geq \sum_{k=1}^K \alpha(k) g_t(k) \mid H_t \right] \\ &\geq \frac{\sum_{k=1}^K \alpha(k) \sigma_t(k)}{4\sqrt{\pi} \cdot \sum_{k=1}^K \alpha(k) g_t(k)} \cdot \exp \left\{ -\frac{1}{2} \left[ \frac{\sum_{k=1}^K \alpha(k) g_t(k)}{\sum_{k=1}^K \alpha(k) \sigma_t(k)} \right]^2 \right\} \end{aligned} \quad (4.5)$$

$$\geq \frac{1}{192\sqrt{\pi}e^{1152}} = c. \quad (4.6)$$

Step (4.3) is by the definition of  $\{\theta_t(i)\}_{i \in L}$  in Line 7 in Algorithm 2. It is important to note that these samples share the *same* random seed  $Z_t$ . Next, Step (4.4) is by the Lemma assumption that  $H_t \in \mathcal{H}_{\hat{w}, t}$ , which means that  $\hat{w}_t(k) \geq w_t(k) - g_t(k)$  for all  $k \in [K]$ . Step (4.5) is an application of the anti-concentration inequality of a normal random variable in Theorem B.2 as  $g_t(k) \geq \sigma_t(k)$ . Step (4.6) is by applying the definition of  $g_t(k)$  and  $\sigma_t(k)$ . ■

We would like to highlight an obstacle in deriving a problem-independent bound for the CTS algorithm. In the proof of Lemma 4.4, which is essential in deriving the bound for TS-CASCADE, step (4.4) results from the fact that the TS-CASCADE algorithm generates Thompson samples with the *same* random seed  $Z_t$  at time step  $t$  (Line 3 – 8 in Algorithm 2). However, CTS generate samples for each item *individually*, which hinders the derivation of a similar lemma. Therefore, we conjecture that a problem-independent bound for CTS cannot be derived using the same strategy as that for TS-CASCADE.

Combining Lemmas 4.3 and 4.4, we conclude that there exists an absolute constant  $c$  such that, for any time step  $t$  and any historical observation  $H_t \in \mathcal{H}_{\hat{w},t}$ ,

$$\Pr_{\boldsymbol{\theta}_t} [S_t \in \mathcal{S}_t \mid H_t] \geq c - \frac{1}{(t+1)^3}. \quad (4.7)$$

Equipped with (4.7), we are able to provide an upper bound on the regret of our Thompson sampling algorithm at every sufficiently large time step.

**Lemma 4.5** *Let  $c$  be an absolute constant such that Lemma 4.4 holds. Consider a time step  $t$  that satisfies  $c - 2/(t+1)^3 > 0$ . Conditional on an arbitrary but fixed historical observation  $H_t \in \mathcal{H}_{\hat{w},t}$ , we have*

$$\mathbb{E}_{\boldsymbol{\theta}_t} [r(S^*|\mathbf{w}) - r(S_t|\mathbf{w}) | H_t] \leq \left(1 + \frac{2}{c}\right) \mathbb{E}_{\boldsymbol{\theta}_t} [F(S_t, t) \mid H_t] + \frac{L}{(t+1)^3}.$$

The proof of Lemma 4.5 relies crucially on *truncating* the original Thompson sample  $\boldsymbol{\theta}_t \in \mathbb{R}$  to  $\tilde{\boldsymbol{\theta}}_t \in [0, 1]^L$ . Under this truncation operation,  $S_t$  remains optimal under  $\tilde{\boldsymbol{\theta}}_t$  (as it was under  $\boldsymbol{\theta}_t$ ) and  $|\tilde{\theta}_t(i) - w(i)| \leq |\theta_t(i) - w(i)|$ , i.e., the distance from the truncated Thompson sample to the ground truth is not increased.

For any  $t$  satisfying  $c - 2/(t+1)^3 > 0$ , define

$$\mathcal{F}_{i,t} := \{\text{Observe } W_t(i) \text{ at } t\}, \quad G(S_t, \mathbf{W}_t) := \sum_{k=1}^K \mathbf{1}(\mathcal{F}_{i_k^t, t}) \cdot (g_t(i_k^t) + h_t(i_k^t)),$$

we unravel the upper bound in Lemma 4.5 to establish the expected regret at time step  $t$ :

$$\begin{aligned} \mathbb{E} \{r(S^*|\mathbf{w}) - r(S_t|\mathbf{w})\} &\leq \mathbb{E} [\mathbb{E}_{\boldsymbol{\theta}_t} [r(S^*|\mathbf{w}) - r(S_t|\mathbf{w}) \mid H_t] \cdot \mathbf{1}(H_t \in \mathcal{H}_{\hat{w},t})] + \mathbb{E} [\mathbf{1}(H_t \notin \mathcal{H}_{\hat{w},t})] \\ &\leq \left(1 + \frac{2}{c}\right) \mathbb{E} [\mathbb{E}_{\boldsymbol{\theta}_t} [F(S_t, t) \mid H_t] \mathbf{1}(H_t \in \mathcal{H}_{\hat{w},t})] + \frac{1}{(t+1)^3} + \frac{3L}{(t+1)^3} \\ &\leq \left(1 + \frac{2}{c}\right) \mathbb{E} [F(S_t, t)] + \frac{4L}{(t+1)^3} \end{aligned} \quad (4.8)$$

$$\begin{aligned} &= \left(1 + \frac{2}{c}\right) \mathbb{E} [\mathbb{E}_{\mathbf{W}_t} [G(S_t, \mathbf{W}_t) | H_t, S_t]] + \frac{4L}{(t+1)^3} \\ &= \left(1 + \frac{2}{c}\right) \mathbb{E} [G(S_t, \mathbf{W}_t)] + \frac{4L}{(t+1)^3}, \end{aligned} \quad (4.9)$$

where (4.8) follows by assuming  $t$  is sufficiently large.

**Lemma 4.6** *For any realization of historical trajectory  $\mathcal{H}_{T+1}$ , we have*

$$\sum_{t=1}^T G(S_t, \mathbf{W}_t) \leq 6\sqrt{KLT} \log T + 144L \log^{5/2} T.$$

**Proof** Recall that for each  $i \in [L]$  and  $t \in [T+1]$ ,  $N_t(i) = \sum_{s=1}^{t-1} \mathbf{1}(\mathcal{F}_{i,s})$  is the number of rounds in  $[t-1]$  when we get to observe the outcome for item  $i$ . Since  $G(S_t, \mathbf{W}_t)$  involves  $g_t(i) + h_t(i)$ , we first bound this term. The definitions of  $g_t(i)$  and  $h_t(i)$  yield that

$$g_t(i) + h_t(i) \leq \frac{12 \log(t+1)}{\sqrt{N_t(i)+1}} + \frac{72 \log^{3/2}(t+1)}{N_t(i)+1}.$$

Subsequently, we decompose  $\sum_{t=1}^T G(S_t, \mathbf{W}_t)$  according to its definition. For a fixed but arbitrary item  $i$ , consider the sequence  $(U_t(i))_{t=1}^T = (\mathbf{1}(\mathcal{F}_{i,t}) \cdot (g_t(i) + h_t(i)))_{t=1}^T$ . Clearly,  $U_t(i) \neq 0$  if and only if the decision maker observes the realization  $W_t(i)$  of item  $i$  at  $t$ . Let  $t = \tau_1 < \tau_2 < \dots < \tau_{N_{T+1}(i)}$  be the time steps when  $U_t(i) \neq 0$ . We assert that  $N_{\tau_n}(i) = n - 1$  for each  $n$ . Indeed, prior to time steps  $\tau_n$ , item  $i$  is observed precisely in the time steps  $\tau_1, \dots, \tau_{n-1}$ . Thus, we have

$$\sum_{t=1}^T \mathbf{1}(\mathcal{F}_{i,t}) \cdot (g_t(i) + h_t(i)) = \sum_{n=1}^{N_{T+1}(i)} (g_{\tau_n}(i) + h_{\tau_n}(i)) \leq \sum_{n=1}^{N_{T+1}(i)} \frac{12 \log T}{\sqrt{n}} + \frac{72 \log^{3/2} T}{n}. \quad (4.10)$$

Now we complete the proof as follows:

$$\begin{aligned} \sum_{t=1}^T \sum_{k=1}^K \mathbf{1}(\mathcal{F}_{i_k^t, t}) \cdot (g_t(i_k^t) + h_t(i_k^t)) &= \sum_{i \in [L]} \sum_{t=1}^T \mathbf{1}(\mathcal{F}_{i,t}) \cdot (g_t(i) + h_t(i)) \\ &\leq \sum_{i \in [L]} \sum_{n=1}^{N_{T+1}(i)} \frac{12 \log T}{\sqrt{n}} + \frac{72 \log^{3/2} T}{n} \leq 6 \sum_{i \in [L]} \sqrt{N_{T+1}(i)} \log T + 72L \log^{3/2} T (\log T + 1) \end{aligned} \quad (4.11)$$

$$\leq 6 \sqrt{L \sum_{i \in [L]} N_{T+1}(i)} \log T + 72L \log^{3/2} T (\log T + 1) \leq 6\sqrt{KLT} \log T + 144L \log^{5/2} T, \quad (4.12)$$

where step (4.11) follows from step (4.10), the first inequality of step (4.12) follows from the Cauchy-Schwarz inequality, and the second one is because the decision maker can observe at most  $K$  items at each time step, hence  $\sum_{i \in [L]} N_{T+1}(i) \leq KT$ .  $\blacksquare$

Finally, we bound the total regret from above by separating the whole horizon into two parts with time step  $t_0 := \lceil \sqrt[3]{2/c} \rceil$ . We bound the regret for the time steps before  $t_0$  by 1, and the regret for time steps after  $t_0$  by (4.9), which holds for all  $t > t_0$ :

$$\begin{aligned} \text{Reg}(T) &\leq \left\lceil \sqrt[3]{\frac{2}{c}} \right\rceil + \sum_{t=t_0+1}^T \mathbb{E} \{r(S^* | \mathbf{w}) - r(S_t | \mathbf{w})\} \\ &\leq \left\lceil \sqrt[3]{\frac{2}{c}} \right\rceil + \left(1 + \frac{2}{c}\right) \mathbb{E} \left[ \sum_{t=1}^T G(S_t, \mathbf{W}_t) \right] + \sum_{t=t_0+1}^T \frac{4L}{(t+1)^3}. \end{aligned}$$

It is clear that the third term is  $O(L)$ , and by Lemma 4.6, the second term is  $O(\sqrt{KLT} \log T + L \log^{5/2} T)$ . Altogether, Theorem 4.1 is proved.

## 5. Linear generalization

In the standard cascading bandit problem, when the number of ground items  $L$  is large, the algorithm converges slowly. In detail, the agent learns each  $w(i)$  individually, and the

knowledge in  $w(i)$  does not help in the estimation of  $w(j)$  for  $j \neq i$ . To ameliorate this problem, we consider the linear generalization setting as described in Section 1. Recall that  $w(i) = x(i)^\top \beta$ , where  $x(i)$ 's are known feature vectors, but  $\beta$  is an unknown universal vector. Hence, the problem reduces to estimating  $\beta \in \mathbb{R}^d$ , which, when  $d < L$ , is easier than estimating the  $L$  distinct weights  $w(i)$ 's in the standard cascading bandit problem.

We construct a Thompson Sampling algorithm `LINTS-CASCADE`( $\lambda$ ) to solve the cascading bandit problem, and provide a theoretical guarantee for this algorithm.

---

**Algorithm 3** `LINTS-CASCADE`( $\lambda$ )
 

---

- 1: Input feature matrix  $X_{d \times L} = [x(1)^\top, \dots, x(L)^\top]$  and parameter  $\lambda$ . Set  $R = 1/2$ .
- 2: Initialize  $b_1 = \mathbf{0}_{d \times 1}$ ,  $M_1 = I_{d \times d}$ ,  $\hat{\psi}_1 = M_1^{-1} b_1$ .
- 3: **for**  $t = 1, 2, \dots$  **do**
- 4:     Sample a  $d$ -dimensional random variable  $\xi_t \sim \mathcal{N}(\mathbf{0}_{d \times 1}, I_{d \times d})$ .
- 5:     Construct the Thompson sample  $v_t = 3R\sqrt{d \log t}$ ,  $\rho_t = \hat{\psi}_t + \lambda v_t \sqrt{K} M_t^{-1/2} \xi_t$ .
- 6:     **for**  $k \in [K]$  **do**
- 7:         Extract  $i_k^t \in \operatorname{argmax}_{i \in [L] \setminus \{i_1^t, \dots, i_{k-1}^t\}} x(i)^\top \rho_t$ .
- 8:     **end for**
- 9:     Pull arm  $S_t = (i_1^t, i_2^t, \dots, i_K^t)$ .
- 10:    Observe click  $k_t \in \{1, \dots, K, \infty\}$ .
- 11:    Update  $M_t, B_t$  and  $\hat{\psi}_t$  with observed items

$$M_{t+1} = M_t + \sum_{k=1}^{\min\{k_t, K\}} x(i_k) x(i_k)^\top, \quad b_{t+1} = b_t + \sum_{k=1}^{\min\{k_t, K\}} x(i_k) W_t(i_k),$$

$$\hat{\psi}_{t+1} = M_{t+1}^{-1} b_{t+1}.$$

12: **end for**

---

**Theorem 5.1 ( Problem-independent bound for `LINTS-CASCADE`)** *Assuming the  $l_2$  norms of feature parameters are bounded by  $1/\sqrt{K}$ , i.e.,  $\|x(i)\| \leq 1/\sqrt{K}$  for all  $i \in [L]$ ,<sup>1</sup> Algorithm `LINTS-CASCADE`( $\lambda$ ), presented in Algorithm 3, incurs an expected regret at most*

$$O\left(\min\{\sqrt{d}, \sqrt{\log L}\} \cdot K d \sqrt{T} (\log T)^{3/2}\right).$$

Here the big  $O$  notation depends only on  $\lambda$ .

The upper bound in Theorem 5.1 grows with  $g(\lambda) = (1 + \lambda) \cdot (1 + 8\sqrt{\pi} e^{7/(2\lambda^2)})$ . Recall that for `CASCADELINUCB`, Zong et al. (2016) derived an upper bound in the form of  $\tilde{O}(K d \sqrt{T})$ ; and for Thompson sampling applied to linear contextual bandits, Agrawal and Goyal (2013a) proved a high probability regret bound of  $\tilde{O}(d^{3/2} \sqrt{T})$ . Our result is a factor  $\sqrt{d}$  worse than the bound on `CASCADELINUCB`. However, our bound matches that of

---

1. It is realistic to assume that  $\|x(i)\| \leq 1/\sqrt{K}$  for all  $i \in [L]$  since this can always be ensured by rescaling the feature vectors. If instead, we assume that  $\|x(i)\| \leq 1$  as in Gai; Wen et al. (2015); Zong et al. (2016), the result in Lemma 5.7 will be  $\sqrt{K}$  worse, which would lead to an upper bound larger than the bound stated in Theorem 5.1 by a factor of  $\sqrt{K}$ .

Thompson sampling for contextual bandits with  $K = 1$  up to log factors. While Zong et al. (2016) proposed CASCADELINTS, they did not provide accompanying performance analysis of this algorithm. Our work fills this gap by analyzing a Thompson sampling algorithm for the linear cascading bandit setting. The tunable parameter  $\sigma$  in CASCADELINTS (Zong et al., 2016) plays a role similar to that of  $\lambda v_t \sqrt{K}$  in LINTS-CASCADE( $\lambda$ ). However, the key difference is that  $v_t$  depends on the current time step  $t$ , which makes it more adaptive compared to  $\sigma$ , which is fixed throughout the horizon. As shown in Line 5 of Algorithm 3, we apply  $\lambda$  to scale the variance of Thompson samples. We see that when the given set of feature vectors are “well-conditioned”, i.e., the set of feature vectors of optimal items and suboptimal items are far away from each other in the feature space  $[0, 1]^d$ , then the agent should choose a small  $\lambda$  to stimulate exploitation. Otherwise, a large  $\lambda$  encourages more exploration.

Now we provide a proof sketch of Theorem 5.1, which is similar to that of Theorem 4.1. More details are given in Appendix C.

To begin with, we observe that a certain noise random variable possesses the sub-Gaussian property (same as Appendix A.1. of Filippi et al. (2010)), and we assume  $x(i)$  is bounded for each  $i$ :

**Remark 5.2 (Sub-Gaussian property)** *Define  $\eta_t(i) := W_t(i) - x(i)^\top \beta$  and let  $R = 1/2$ . The fact that  $\mathbb{E}W_t(i) = w(i) = x(i)^\top \beta$  and  $W_t(i) \in [0, 1]$  yields that  $\mathbb{E}\eta_t(i) = \mathbb{E}W_t(i) - w(i) = 0$  and  $\eta_t(i) \in [-w(i), 1 - w(i)]$ . Hence,  $\eta_t(i)$  is 0-mean and  $R$ -sub-Gaussian, i.e.,*

$$\mathbb{E}[\exp(\lambda \eta_t) | \mathcal{F}_{t-1}] \leq \exp\left(\frac{\lambda^2 R^2}{2}\right) \quad \forall \lambda \in \mathbb{R}.$$

During the iterations, we update  $\hat{\psi}_{t+1}$  such that  $x(i)^\top \hat{\psi}_{t+1}$  approaches  $x(i)^\top \beta$  eventually. To do so, we select a set  $S_t$  according to the order of  $x(i)^\top \rho_t$ 's at each time step. Hence, if  $x(i)^\top \hat{\psi}_{t+1}$ ,  $x(i)^\top \rho_t$  and  $x(i)^\top \beta$  are close enough, then we are likely to select the optimal set.

Following the same ideas as in the standard case, this motivates us to define two “nice events” as follows:

$$\begin{aligned} \mathcal{E}_{\hat{\psi},t} &:= \left\{ \forall i \in [L] : \left| x(i)^\top \hat{\psi}_t - x(i)^\top \beta \right| \leq g_t(i) \right\}, \\ \mathcal{E}_{\rho,t} &:= \left\{ \forall i \in [L] : \left| x(i)^\top \rho_t - x(i)^\top \hat{\psi}_t \right| \leq h_t(i) \right\}, \end{aligned}$$

where

$$\begin{aligned} v_t &:= 3R\sqrt{d \log t}, \quad l_t := R\sqrt{3d \log t} + 1, \quad u_t(i) := \sqrt{x(i)^\top M_t^{-1} x(i)}, \\ g_t(i) &:= l_t u_t(i), \quad h_t(i) := \min\{\sqrt{d}, \sqrt{\log L}\} \cdot \sqrt{4K \log t} \cdot \lambda v_t u_t(i). \end{aligned}$$

Analogously to Lemma 4.2, Lemma 5.3 illustrates that  $\mathcal{E}_{\hat{\psi},t}$  has high probability. This follows from the concentration inequality in Theorem B.3, and the analysis is similar to that of Lemma 1 in Agrawal and Goyal (2013a). It also illustrates that  $\mathcal{E}_{\rho,t}$  has high probability, which is a direct consequence of the concentration bound in Theorem B.2.

**Lemma 5.3** *For each  $t \in [T]$ , we have*

$$\Pr \left[ \mathcal{E}_{\hat{\psi}, t} \right] \geq 1 - \frac{1}{t^2}, \quad \Pr [\mathcal{E}_{\rho, t} | H_t] \geq 1 - \frac{d+1}{t^2}.$$

Next, we revise the definition of  $F(S, t)$  in equation (4.1) for the linear generalization setting:

$$F(S, t) := \sum_{k=1}^K \left[ \prod_{j=1}^{k-1} (1 - x(i_j)^\top \beta) \right] (g_t(i_k) + h_t(i_k)), \quad (5.1)$$

$$\mathcal{S}_t := \left\{ S = (i_1, \dots, i_K) \in \pi_K(L) : F(S, t) \geq r(S^* | w) - r(S | w) \right\}. \quad (5.2)$$

Further, following the intuition of Lemma 4.3 and Lemma 4.4, we state and prove Lemma 5.4 and Lemma 5.5 respectively. These lemmas are similar in spirit to Lemma 4.3 and Lemma 4.4 except that  $\theta_t(k)$  in Lemma 4.3 is replaced by  $x(k)^\top \rho_t$  in Lemma 5.4, when inequality (4.2) and  $c$  in Lemma 4.4 are replaced by inequality (5.3) and  $c(\lambda)$  in Lemma 5.5 respectively.

**Lemma 5.4** *Consider a time step  $t$ . Suppose that events  $\mathcal{E}_{\hat{\psi}, t}$ ,  $\mathcal{E}_{\rho, t}$  and inequality*

$$\sum_{k=1}^K \left[ \prod_{j=1}^{k-1} (1 - w(j)) \right] x(k)^\top \rho_t \geq \sum_{k=1}^K \left[ \prod_{j=1}^{k-1} (1 - w(j)) \right] w(k) \quad (5.3)$$

*hold, then the event  $\{S_t \in \mathcal{S}_t\}$  also holds.*

**Lemma 5.5** *There exists an absolute constant  $c(\lambda) = 1/(4\sqrt{\pi}e^{7/2\lambda^2}) \in (0, 1)$  independent of  $w, K, L, T$  such that, for any time step  $t \geq 4$  and any historical observation  $H_t \in \mathcal{H}_{\hat{\psi}, t}$ , the following inequality holds:*

$$\Pr_{\rho_t} [\mathcal{E}_{\rho, t} \text{ and (5.3) hold} \mid H_t] \geq c(\lambda) - \frac{d+1}{t^2}.$$

Above yields that there exists an absolute constant  $c(\lambda)$  such that, for any time step  $t$  and any historical observation  $H_t \in \mathcal{H}_{\hat{\psi}, t}$ ,

$$\Pr_{\rho_t} [S_t \in \mathcal{S}_t \mid H_t] \geq c(\lambda) - \frac{d+1}{t^2}. \quad (5.4)$$

This allows us to upper bound the regret of LINTS-CASCADE( $\lambda$ ) at every sufficiently large time step. The proof of this lemma is similar to that of Lemma 4.5 except that  $\theta_t(k)$  in Lemma 4.5 is replaced by  $x(k)^\top \theta_t$ . We omit the proof for the sake of simplicity.

**Lemma 5.6** *Let  $c(\lambda)$  be an absolute constant such that Lemma 5.5 holds. Consider a time step  $t$  that satisfies  $c(\lambda) > 2(d+1)/t^2$ ,  $t > 4$ . Conditional on an arbitrary but fixed historical observation  $H_t \in \mathcal{H}_{\hat{\psi}, t}$ , we have*

$$\mathbb{E}_{\rho_t} [r(S^* | w) - r(S_t | w) | H_t] \leq \left( 1 + \frac{4}{c(\lambda)} \right) \mathbb{E}_{\rho_t} [F(S_t, t) \mid H_t] + \frac{d+1}{t^2}.$$

For any  $t$  satisfying  $c(\lambda) > 2(d+1)/t^2$ ,  $t > 4$ , recall that

$$\mathcal{F}_{i,t} := \{\text{Observe } W_t(i) \text{ at } t\}, \quad G(S_t, W_t) := \sum_{k=1}^K \mathbf{1}(\mathcal{F}_{i_k^t, t}) \cdot (g_t(i_k^t) + h_t(i_k^t)),$$

we unravel the upper bound in Lemma 5.6 to establish the expected regret at time step  $t$ :

$$\begin{aligned} \mathbb{E} \{r(S^*|w) - r(S_t|w)\} &\leq \mathbb{E} \left[ \mathbb{E}_{\rho_t} [r(S^*|w) - r(S_t|w) \mid H_t] \cdot \mathbf{1}(H_t \in \mathcal{H}_{\hat{\psi}, t}) \right] + \mathbb{E} \left[ \mathbf{1}(H_t \notin \mathcal{H}_{\hat{\psi}, t}) \right] \\ &\leq \left(1 + \frac{2}{c(\lambda)}\right) \mathbb{E} \left[ \mathbb{E}_{\rho_t} [F(S_t, t) \mid H_t] \mathbf{1}(H_t \in \mathcal{H}_{\hat{\psi}, t}) \right] + \frac{d+1}{t^2} + \frac{1}{t^2} \end{aligned} \quad (5.5)$$

$$= \left(1 + \frac{2}{c(\lambda)}\right) \mathbb{E} [G(S_t, W_t)] + \frac{d+2}{t^2}, \quad (5.6)$$

where step (5.5) follows by assuming  $t$  is sufficiently large.

**Lemma 5.7** *For any realization of historical trajectory  $\mathcal{H}_{T+1}$ , we have*

$$\sum_{t=1}^T G(S_t, W_t) \leq 40R(1+\lambda) \min\{\sqrt{d}, \sqrt{\log L}\} \cdot Kd\sqrt{T}(\log T)^{3/2}.$$

Note that here similarly to Lemma 4.6, we prove a worst case bound without needing the expectation operator.

**Proof** Recall that for each  $i \in [L]$  and  $t \in [T+1]$ ,  $N_t(i) = \sum_{s=1}^{t-1} \mathbf{1}(\mathcal{F}_{i,s})$  is the number of rounds in  $[t-1]$  when we get to observe the outcome for item  $i$ . Since  $G(S_t, W_t)$  involves  $g_t(i) + h_t(i)$ , we first bound this term. The definitions of  $g_t(i)$  and  $h_t(i)$  yield that

$$\begin{aligned} g_t(i) + h_t(i) &= u_t(i)[l_t + \min\{\sqrt{d}, \sqrt{\log L}\} \cdot \lambda\sqrt{4K \log t} \cdot v_t] \\ &= u_t(i)[R\sqrt{3d \log t} + 1 + \min\{\sqrt{d}, \sqrt{\log L}\} \cdot \lambda\sqrt{4K \log t} \cdot 3R\sqrt{d \log t}] \\ &\leq 8R(1+\lambda)u_t(i) \log t \sqrt{dK} \cdot \min\{\sqrt{d}, \sqrt{\log L}\}. \end{aligned}$$

Further, we have

$$\begin{aligned} \sum_{t=1}^T G(S_t, W_t) &= \sum_{t=1}^T \sum_{k=1}^K \mathbf{1}(\mathcal{F}_{i_k^t, t}) \cdot [g_t(i_k^t) + h_t(i_k^t)] = \sum_{t=1}^T \sum_{k=1}^{k_t} [g_t(i_k^t) + h_t(i_k^t)] \\ &\leq 8R(1+\lambda) \log T \sqrt{dK} \cdot \min\{\sqrt{d}, \sqrt{\log L}\} \sum_{t=1}^T \sum_{k=1}^{k_t} u_t(i_k^t). \end{aligned}$$

Let  $x(i) = (x_1(i), x_2(i), \dots, x_d(i))$ ,  $y_t = (y_{t,1}, y_{t,2}, \dots, y_{t,d})$  with

$$y_{t,l} = \sqrt{\sum_{k=1}^{k_t} x_l(i_k^t)^2}, \quad \forall l = 1, 2, \dots, d,$$

then  $M_{t+1} = M_t + y_t y_t^\top$ ,  $\|y_t\| \leq 1$ .



Now, we have

$$\sum_{t=1}^T \sqrt{y_t^\top M_t^{-1} y_t} \leq 5\sqrt{dT \log T}, \quad (5.7)$$

which can be derived along the lines of Lemma 3 in Chu et al. (2011) using Lemma C.5 (see Appendix C.10 for details). Let  $M(t)_{m,n}^{-1}$  denote the  $(m, n)$ -th-element of  $M(t)^{-1}$ . Note that

$$\sum_{k=1}^{k_t} u_t(i_k^t) = \sum_{k=1}^{k_t} \sqrt{x(i_k^t)^\top M(t)^{-1} x(i_k^t)} \quad (5.8)$$

$$\leq \sqrt{k_t \cdot \sum_{k=1}^{k_t} x(i_k^t)^\top M(t)^{-1} x(i_k^t)} \quad (5.9)$$

$$\begin{aligned} &\leq \sqrt{K} \sqrt{\sum_{1 \leq m, n \leq d} M(t)_{m,n}^{-1} \left[ \sum_{k=1}^{k_t} x_m(i_k^t) x_n(i_k^t) \right]} \\ &\leq \sqrt{K} \sqrt{\sum_{1 \leq m, n \leq d} M(t)_{m,n}^{-1} \sqrt{\left[ \sum_{k=1}^{k_t} x_m(i_k^t)^2 \right] \cdot \left[ \sum_{k=1}^{k_t} x_n(i_k^t)^2 \right]}} \quad (5.10) \\ &= \sqrt{K} \sqrt{\sum_{1 \leq m, n \leq d} M(t)_{m,n}^{-1} y_{t,m} y_{t,n}} = \sqrt{K} \sqrt{y_t^\top M_t^{-1} y_t}. \end{aligned}$$

where step (5.8) follows from the definition of  $u_t(i)$ , (5.9) and (5.10) follow from Cauchy-Schwartz inequality.

Now we complete the proof as follows:

$$\begin{aligned} \sum_{t=1}^T G(S_t, W_t) &\leq 8R(1 + \lambda) \log T \sqrt{dK} \cdot \min\{\sqrt{d}, \sqrt{\log L}\} \sum_{t=1}^T \sum_{k=1}^{k_t} u_t(i_k^t) \\ &\leq 8R(1 + \lambda) \log T \sqrt{dK} \cdot \min\{\sqrt{d}, \sqrt{\log L}\} \cdot \sum_{t=1}^T \sqrt{K} \cdot \sqrt{y_t^\top M_t^{-1} y_t} \\ &\leq 40R(1 + \lambda) \min\{\sqrt{d}, \sqrt{\log L}\} \cdot Kd\sqrt{T}(\log T)^{3/2}. \end{aligned}$$

■

Finally, we bound the total regret from above by separating the whole horizon into two parts with time step  $t(\lambda) := \lceil \sqrt{8(d+1)/c(\lambda)} \rceil$  (noting  $c(\lambda) > 2(d+1)/t^2$ ,  $t > 4$  when  $t \geq t(\lambda)$ ). We bound the regret for the time steps before  $t(\lambda)$  by 1 and the regret for time steps after  $t(\lambda)$  by (5.6), which holds for all  $t > t(\lambda)$ :

$$\begin{aligned} \text{Reg}(T) &\leq \left\lceil \frac{2}{c(\lambda)} \right\rceil + \sum_{t=t(\lambda)+1}^T \mathbb{E} \{r(S^*|w) - r(S_t|w)\} \\ &\leq \left\lceil \frac{2}{c(\lambda)} \right\rceil + \left(1 + \frac{2}{c(\lambda)}\right) \mathbb{E} \left[ \sum_{t=1}^T G(S_t, W_t) \right] + \sum_{t=t(\lambda)+1}^T \frac{d+2}{t^2}. \end{aligned}$$

It is clear that the first term is  $O(1)$  and the third term is  $O(d)$ , and by Lemma 5.7, the second term is  $O(\min\{\sqrt{d}, \sqrt{\log L}\} \cdot Kd\sqrt{T}(\log T)^{3/2})$ . Altogether, Theorem 5.1 is proved.

## 6. Lower bound on regret of the standard setting

We derive a problem-independent minimax lower bound on the regret of any algorithm  $\pi$  (see definition in Section 2) in the standard cascading bandit problem, following the exposition in Bubeck et al. (2012, Theorem 3.5).

**Theorem 6.1 (Problem-independent lower bound)** *The following regret lower bound holds for the cascading bandit problem:*

$$\inf_{\pi} \sup_{\mathcal{I}} \text{Reg}(T; \pi, \mathcal{I}) = \Omega\left(\sqrt{LT/K}\right),$$

where  $\pi$  denotes an algorithm, and  $\mathcal{I}$  denotes a cascading bandit instance. The  $\Omega$  notation hides a constant factor that is independent of  $K, L, T, \mathbf{w}$ .

Now we outline the proof of Theorem 6.1; the accompanying lemmas are proved in Appendix C. The theorem is derived probabilistically: we construct several difficult instances such that the difference between click probabilities of optimal and suboptimal items is subtle in each of them, and hence the regret of any algorithm must be large when averaged over these distributions.

The key difference from the analysis in Bubeck et al. (2012) involves: (i) the difference between instance 0 and instance  $\ell$  ( $1 \leq \ell \leq L$ ) involves different distributions of  $k$  items; (ii) the partial feedback from the user results in a more delicate analysis of the KL divergence between the probability measures of stochastic process  $\{S_t^\pi, O_t^\pi\}_{t=1}^T$  under instances 0 and instance  $\ell$  ( $1 \leq \ell \leq L$ ) (see Lemmas 6.4 and 6.5).

**Notations.** Before presenting the proof, we remind the reader of the definition of the KL divergence (Cover and Thomas, 2012). For two probability mass functions  $P_X$  and  $P_Y$  defined on the same discrete probability space  $\mathcal{A}$ ,

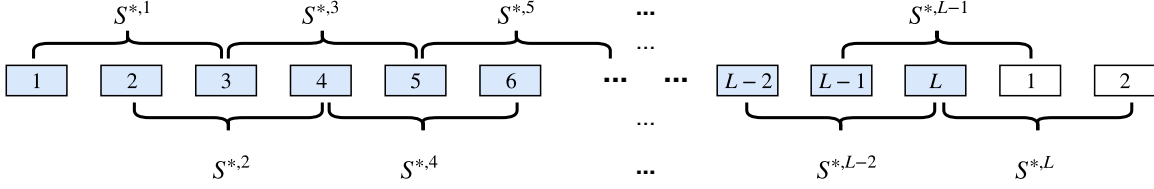
$$\text{KL}(P_X, P_Y) = \sum_{x \in \mathcal{A}} P_X(x) \log \frac{P_X(x)}{P_Y(x)}$$

denotes the KL divergence between the probability mass functions  $P_X, P_Y$ . (We abbreviate  $\log_e$  as  $\log$  for simplicity.) Lastly, when  $a$  and  $b$  are two real numbers between 0 and 1,  $\text{KL}(a, b) = \text{KL}(\text{Bern}(a) \parallel \text{Bern}(b))$ , i.e.,  $\text{KL}(a, b)$  denotes the KL divergence between  $\text{Bern}(a)$  and  $\text{Bern}(b)$ .

### First step: Construct instances

**Definition of instances.** To begin with, we define a class of  $L + 1$  instances, indexed by  $\ell = 0, 1, \dots, L$  as follows:

- Under instance 0, we have  $w^{(0)}(i) = \frac{1-\epsilon}{K}$  for all  $i \in [L]$ .
- Under instance  $\ell$ , where  $1 \leq \ell \leq L$ , we have  $w^{(\ell)}(i) = \frac{1+\epsilon}{K}$  if  $i \in \{\underline{\ell}, \underline{\ell} + 1, \dots, \underline{\ell} + K - 1\}$ , where  $\underline{a} = a$  if  $1 \leq a \leq L$ , and  $\underline{a} = a - L$  if  $L < a \leq 2L$ . We have  $w^{(\ell)}(i) = \frac{1-\epsilon}{K}$  if  $i \in [L] \setminus \{\underline{\ell}, \underline{\ell} + 1, \dots, \underline{\ell} + K - 1\}$ ,


 Figure 1: An example of instances with  $K = 3$ 

where  $\epsilon \in [0, 1]$  is a small and positive number such that  $0 < (1 - \epsilon)/K < (1 + \epsilon)/K < 1/4$ . The expected reward of a list  $S$  under instance  $\ell$  ( $0 \leq \ell \leq L$ ) is

$$r(S|w^{(\ell)}) = 1 - \prod_{i \in S} (1 - w^{(\ell)}(i)).$$

Under instance 0, every list  $S \in [L]^{(K)}$  is optimal, while under instance  $\ell$  ( $1 \leq \ell \leq L$ ), a list  $S$  is optimal iff  $S$  is an ordered  $K$ -set comprised of items  $\{\underline{\ell}, \underline{\ell} + 1, \dots, \underline{\ell} + K - 1\}$ . Let  $S^{*,\ell} = (\underline{\ell}, \underline{\ell} + 1, \dots, \underline{\ell} + K - 1)$  be an optimal list under instance  $\ell$  for  $\ell \in \{0, \dots, L\}$ . Evidently, it holds that  $r(S^{*,\ell}|w^{(\ell)}) = 1 - [1 - (1 + \epsilon)/K]^K$  if  $\ell \in \{1, \dots, L\}$ , and  $r(S^{*,0}|w^{(0)}) = 1 - [1 - (1 - \epsilon)/K]^K$ . Note that  $S^{*,1}, \dots, S^{*,L}$  ‘uniformly cover’ the ground set  $\{1, \dots, L\}$ , in the sense that each  $i \in \{1, \dots, L\}$  belongs to exactly  $K$  sets among  $S^{*,1}, \dots, S^{*,L}$ . In Figure 1, we give an example of  $S^{*,1}, S^{*,2}, \dots, S^{*,L}$  to illustrate how we construct these instances.

**Notations regarding to a generic online policy  $\pi$ .** A policy  $\pi$  is said to be deterministic and non-anticipatory if the list of items  $S_t^\pi$  recommended by policy  $\pi$  at time  $t$  is completely determined by the observation history  $\{S_s^\pi, O_s^\pi\}_{s=1}^{t-1}$ . That is,  $\pi$  is represented as a sequence of functions  $\{\pi_t\}_{t=1}^\infty$ , where  $\pi_t : ([L]^{(K)} \times \{0, 1, \star\}^K)^{t-1} \rightarrow [L]^{(K)}$ , and  $S_t^\pi = \pi_t(S_1^\pi, O_1^\pi, \dots, S_{t-1}^\pi, O_{t-1}^\pi)$ . We represent  $O_t^\pi$  as a vector in  $\{0, 1, \star\}^K$ , where 0, 1,  $\star$  represents observing no click, observing a click and no observation respectively. For example, when  $S_t^\pi = (2, 1, 5, 4)$  and  $O_t^\pi = (0, 0, 1, \star)$ , items 2, 1, 5, 4 are listed in the displayed order; items 2, 1 are not clicked, item 5 is clicked, and the response to item 4 is not observed. By the definition of the cascading model, the outcome  $O_t^\pi = (0, 0, 1, \star)$  is in general a (possibly empty) string of 0s, followed by a 1 (if the realized reward is 1), and then followed by a possibly empty string of  $\star$ s.

Next, for any instance  $\ell$ , we lower bound the instantaneous regret of any arm  $S$  with the number of suboptimal items within the arm:

**Lemma 6.2** *Let  $\epsilon \in [0, 1]$ , integer  $K$  satisfy  $0 < (1 - \epsilon)/K < (1 + \epsilon)/K < 1/4$ . Consider instance  $\ell$ , where  $1 \leq \ell \leq L$ . For any order  $K$ -set  $S$ , we have*

$$r(S^{*,\ell}|w^{(\ell)}) - r(S|w^{(\ell)}) \geq \frac{2|S \setminus S^{*,\ell}|\epsilon}{e^4 K}.$$

### Second step: Pinsker’s inequality

After constructing the family of instances, we consider a fixed but arbitrary non-anticipatory policy  $\pi$ , and analyze the stochastic process  $\{S_t^\pi, O_t^\pi\}_{t=1}^T$ . For instance  $\ell \in \{0, \dots, L\}$ , we

denote  $P_{\{S_t^\pi, O_t^\pi\}_{t=1}^T}^{(\ell)}$  as the probability mass function of  $\{S_t^\pi, O_t^\pi\}_{t=1}^T$  under instance  $\ell$ , and we denote  $\mathbb{E}^{(\ell)}$  as the expectation operator (over random variables related to  $\{S_t^\pi, O_t^\pi\}_{t=1}^T$ ) under instance  $\ell$ . Now, Lemma 6.2 directly leads to a fundamental lower bound on the expected cumulative regret of a policy under an instance  $\ell \in \{1, \dots, L\}$ :

$$\mathbb{E}^{(\ell)} \left[ \sum_{t=1}^T R(S^{*,\ell}|w^{(\ell)}) - R(S_t^\pi|w^{(\ell)}) \right] \geq \frac{2\epsilon}{e^4 K} \cdot \sum_{j \in [L] \setminus S^{*,\ell}} \mathbb{E}^{(\ell)} \left[ \sum_{t=1}^T \mathbf{1}(j \in S_t^\pi) \right]. \quad (6.1)$$

To proceed, we now turn to lower bounding the expectation of the number of times that suboptimal items appear in the chosen arms during the entire horizon. We apply Pinsker's inequality in this step.

**Lemma 6.3** *For any non-anticipatory policy  $\pi$ , it holds that*

$$\begin{aligned} & \frac{1}{L} \sum_{\ell=1}^L \sum_{j \in [L] \setminus S^{*,\ell}} \mathbb{E}^{(\ell)} \left[ \sum_{t=1}^T \mathbf{1}(j \in S_t^\pi) \right] \\ & \geq KT \left\{ 1 - \frac{K}{L} - \sqrt{\frac{1}{2L} \sum_{\ell=1}^L \text{KL} \left( P_{\{S_t^\pi, O_t^\pi\}_{t=1}^T}^{(0)}, P_{\{S_t^\pi, O_t^\pi\}_{t=1}^T}^{(\ell)} \right)} \right\}. \end{aligned}$$

Recall that under the null instance 0, each item has click probabilities  $\frac{1-\epsilon}{K}$ , and the only difference between instance 0 and instance  $\ell$  ( $1 \leq \ell \leq L$ ) is that the optimal items under instance  $\ell$  have click probabilities  $\frac{1-\epsilon}{K}$  under instance 0 but  $\frac{1+\epsilon}{K}$  under instance  $\ell$ . Since the difference between optimal and suboptimal items under instance  $\ell$  is small, it is difficult to distinguish the optimal items from suboptimal ones. Therefore, the set of chosen arms and outcomes under instance  $\ell$  would be similar to that under the null instance 0. More precisely, the distributions of the trajectory  $\{S_t^\pi, O_t^\pi\}_{t=1}^T$  and that of  $\{S_t^{\pi,\ell}, O_t^{\pi,\ell}\}_{t=1}^T$  would be similar, which implies that the KL divergence is small and hence the lower bound on the regret would be large. This intuitively implies that we have constructed instances that are difficult to discriminate and hence maximizing the expected regret.

### Third step: chain rule

To lower bound the regret, it suffices to bound the average of  $\text{KL}(\{S_t^{\pi,0}, O_t^{\pi,0}\}_{t=1}^T, \{S_t^{\pi,\ell}, O_t^{\pi,\ell}\}_{t=1}^T)$  (over  $\ell$ ) from above. Observe that the outcome  $O_t^{\pi,\ell}$  in each time step  $t$  must be one of the following  $K+1$  possibilities:

$$o^0 = (1, \star, \dots, \star), o^1 = (0, 1, \star, \dots, \star), \dots, o^{K-1} = (0, \dots, 0, 1), \text{ and } o^K = (0, \dots, 0).$$

We index each possible outcome by the number of zeros (no-clicks), which uniquely defines each of the possible outcomes, and denote  $\mathcal{O} = \{o^k\}_{k=0}^K$ . The set of all possible realization of each  $(S_t^{\pi,\ell}, O_t^{\pi,\ell})$  is simply  $[L]^{(K)} \times \mathcal{O}$ .

With these notations, we decompose the KL divergence according to how time progresses. In detail, we apply the chain rule for the KL divergence iteratively (Cover and Thomas, 2012) to decompose  $\text{KL}(\{S_t^{\pi,0}, O_t^{\pi,0}\}_{t=1}^T, \{S_t^{\pi,\ell}, O_t^{\pi,\ell}\}_{t=1}^T)$  into per-time-step KL divergences.

**Lemma 6.4** For any instance  $\ell$ , where  $1 \leq \ell \leq L$ ,

$$\begin{aligned} & \text{KL} \left( P_{\{S_t^\pi, O_t^\pi\}_{t=1}^T}^{(0)}, P_{\{S_t^\pi, O_t^\pi\}_{t=1}^T}^{(\ell)} \right) \\ &= \sum_{t=1}^T \sum_{s_t \in [L]^{(K)}} P^{(0)}[S_t^\pi = s_t] \cdot \text{KL} \left( P_{O_t^\pi | S_t^\pi}^{(0)}(\cdot | s_t) \parallel P_{O_t^\pi | S_t^\pi}^{(\ell)}(\cdot | s_t) \right). \end{aligned}$$

#### Fourth step: conclusion

The design of instances 1 to  $L$  implies that the optimal lists cover the whole ground set uniformly, and an arbitrary item in the ground set belongs to exactly  $K$  of the  $L$  optimal lists. This allows us to further upper bound the average of the KL divergence  $\text{KL} \left( P_{O_t^{\pi,0} | S_t^{\pi,0}}^{(0)}(\cdot | s_t) \parallel P_{O_t^{\pi,\ell} | S_t^{\pi,\ell}}^{(\ell)}(\cdot | s_t) \right)$  as follows:

**Lemma 6.5** For any  $t \in [T]$ ,

$$\begin{aligned} & \frac{1}{L} \sum_{\ell=1}^L \sum_{s_t \in [L]^{(K)}} P^{(0)}[S_t^\pi = s_t] \cdot \text{KL} \left( P_{O_t^\pi | S_t^\pi}^{(0)}(\cdot | s_t) \parallel P_{O_t^\pi | S_t^\pi}^{(\ell)}(\cdot | s_t) \right) \\ & \leq \frac{K}{L} \left[ (1 - \epsilon) \log \left( \frac{1 - \epsilon}{1 + \epsilon} \right) + (K - 1 + \epsilon) \log \left( \frac{K - 1 + \epsilon}{K - 1 - \epsilon} \right) \right]. \end{aligned}$$

Substituting the results of Lemma 6.4 and Lemma 6.5 into inequality (C.35), we have

$$\begin{aligned} & \frac{1}{L} \sum_{\ell=1}^L \mathbb{E} \left[ \sum_{t=1}^T R(S_t^{*,\ell} | w^{(\ell)}) - R(S_t^{\pi,\ell} | w^{(\ell)}) \right] \\ & \geq \frac{2\epsilon T}{e^4} \left[ \left( 1 - \frac{K}{L} \right) - \sqrt{\frac{1}{2L} \sum_{t=1}^T \sum_{\ell=1}^L \sum_{s_t \in [L]^{(K)}} P^{(0)}[S_t^\pi = s_t] \cdot \text{KL} \left( P_{O_t^\pi | S_t^\pi}^{(0)}(\cdot | s_t) \parallel P_{O_t^\pi | S_t^\pi}^{(\ell)}(\cdot | s_t) \right)} \right] \\ & = \Omega \left( \epsilon T \left( 1 - \frac{K}{L} - \epsilon \sqrt{\frac{TK}{2L}} \right) \right) = \Omega(\sqrt{LT/K}). \end{aligned}$$

The first step in the last line follows from the fact that

$$(1 - \epsilon) \log \left( \frac{1 - \epsilon}{1 + \epsilon} \right) + (K - 1 + \epsilon) \log \left( \frac{K - 1 + \epsilon}{K - 1 - \epsilon} \right) = O(\epsilon^2)$$

when  $\epsilon$  is a small positive number, and the second step follows from differentiation. Altogether, Theorem 6.1 is proved.

## 7. Numerical experiments

We evaluate the performance of our proposed algorithms TS-CASCADE, LINTS-CASCADE( $\lambda$ ), analyzed algorithm CTS (Huyuk and Tekin, 2019), as well as algorithms in some existing works such as CASCADEUCB1, CASCADEKL-UCB in Kveton et al. (2015a) and

CASCADELINUCB, CASCADELINTS, RANKEDLINTS in Zong et al. (2016). The purpose of this section is not to show that the proposed algorithms — which have strong theoretical guarantees — are uniformly the best possible (for all choices of the parameters), but rather to highlight some features of each algorithm and in which regime(s) which algorithm performs well. We believe this empirical study will be useful for practitioners in the future. The codes to reproduce all the experiments can be found at <https://github.com/zixinzhang/2021-JMLR.git>.

### 7.1 Comparison of Performances of TS-CASCADE and CTS to UCB-based Algorithms for Cascading Bandits

To demonstrate the effectiveness of the TS algorithms TS-CASCADE and CTS, we compare their expected cumulative regrets to CASCADEKL-UCB and CASCADEUCB1 (Kveton et al., 2015a). Since there are many numerical evaluations of these algorithms in the existing literature (Cheung et al., 2019; Huyuk and Tekin, 2019), here we only present a brief discussion, highlighting some salient features.

For the  $L$  items, we set the click probabilities of  $K$  of them to be  $w_1 = 0.2$ , the probabilities of another  $K$  of them to be  $w_2 = 0.1$ , and the probabilities of the remaining  $L - 2K$  items as  $w_3 = 0.05$ . We vary  $L \in \{16, 32, 64, 256, 512, 1024, 2048\}$  and  $K \in \{2, 4\}$ . We conduct 20 independent simulations with each algorithm under each setting of  $L, K$ . We calculate the averages and standard deviations of  $\text{Reg}(T)$ , and as well as the average running times of each experiment.

Our experiments indicate that the two TS algorithms always outperform the two UCB algorithms in all settings considered (see Table 4 in Appendix D) and Figure 2. Next, CASCADEKL-UCB outperforms CASCADEUCB1 when  $L \geq 512$  while CASCADEUCB1 outperforms CASCADEKL-UCB otherwise. Similarly, TS-CASCADE outperforms CTS when  $L \geq 2048$  while CTS outperforms TS-CASCADE otherwise. Thus, for a large number of items, TS-CASCADE is preferred to CTS, again demonstrating the utility of the Gaussian updates over the seemingly more natural Bernoulli updates. For simplicity, we only present a representative subset of the results in Figure 2. These observations imply that the best choice of algorithm depends on  $L$  the total number of items in the ground set.

### 7.2 Performance of LINTS-CASCADE( $\lambda$ ) compared to other algorithms for linear cascading bandits

In this subsection, we compare Algorithm 3 (LINTS-CASCADE( $\lambda$ )) to other competitors using numerical simulations. Note that the three competitors for cascading bandits with linear generalization CASCADELINUCB, CASCADELINTS, RANKEDLINTS (Zong et al., 2016) involve tuning parameters (called  $\sigma$ ). Thus to be fair to each algorithm, we also tune the parameters for each of the algorithms to ensure that the best performances are obtained so that the comparisons are fair.

We now describe how we set up these experiments. In real-world applications, good features that are extracted based on a large number of items are rarely obtained directly but rather learnt from historical data of users’ behaviors (Koren et al., 2009). One commonly used method to obtain the features of items is via low-rank matrix factorization. In detail, let  $A_{\text{train}} \in \{0, 1\}^{m \times L}$  be the matrix representing feedback of  $L$  items from  $m$  users. The

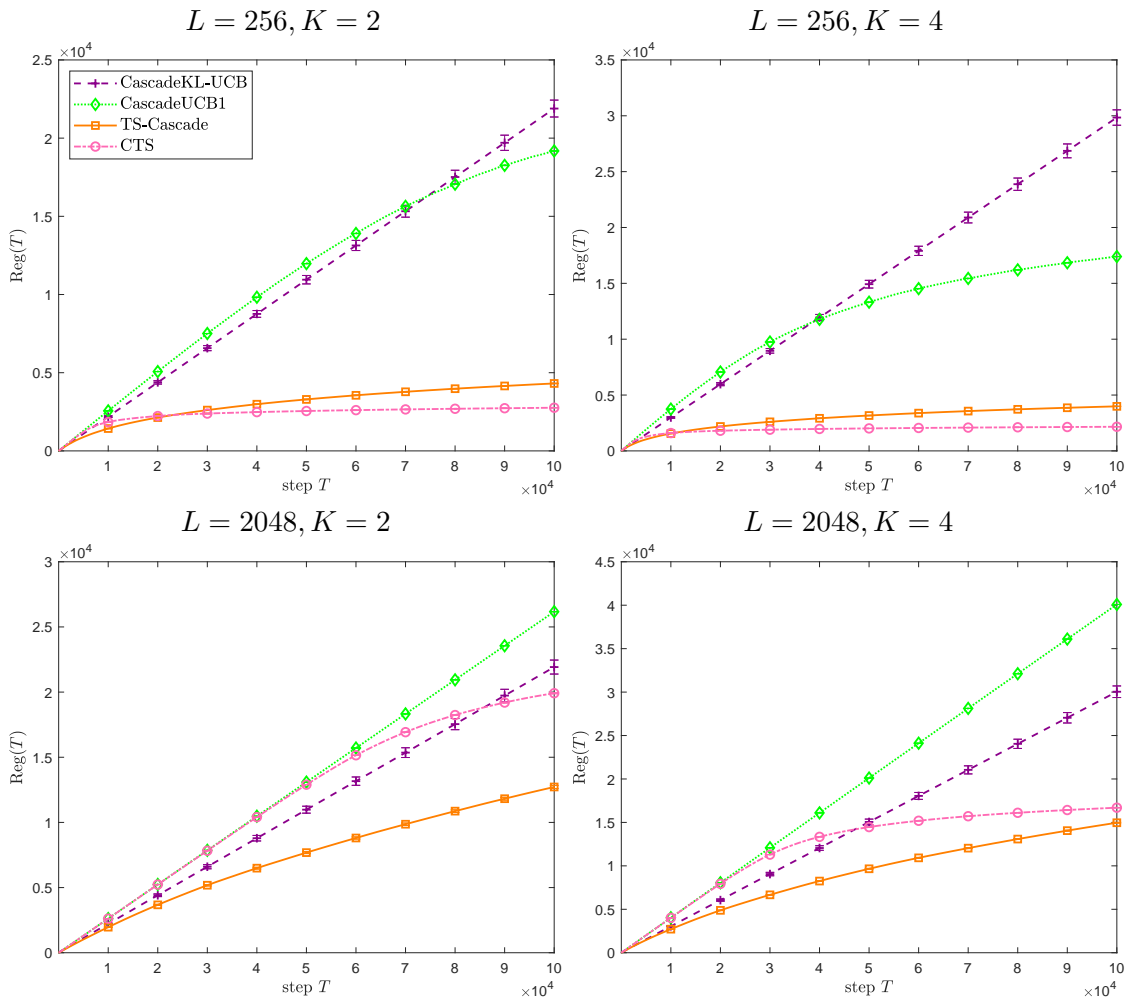


Figure 2:  $\text{Reg}(T)$  of algorithms when  $L \in \{256, 2048\}$ ,  $K \in \{2, 4\}$ . Each line indicates the average  $\text{Reg}(T)$  (over 20 runs).

$(i, j)^{\text{th}}$  element of feedback matrix  $A_{\text{train}}(i, j) = 1$  when user  $i \in [m]$  is attracted to item  $j \in [L]$ , and  $A_{\text{train}}(i, j) = 0$  when user  $i$  is not attracted to item  $j$ . We apply Algorithm 4 to learn the features of the items from  $A_{\text{train}}$ .

---

**Algorithm 4** Generate feature matrix with historical data (Zong et al., 2016)

---

- 1: Input historical data matrix  $A_{\text{train}} \in \{0, 1\}^{m \times L}$ , feature length  $d$ .
  - 2: Conduct rank- $d$  truncated SVD of  $A_{\text{train}}$ :  $A_{\text{train}} \approx U_{m \times d} S_{d \times d} V_{L \times d}^\top$ .
  - 3: Calculate the features of items  $X_{d \times L} = S V^\top$  to be used as an input in Algorithm 3.
- 

Our experiments reveal that, somewhat unsurprisingly, the amount of training data has a noticeable impact on the performances of all algorithms assuming the linear bandit model. In detail, when the amount of training data  $m$  is small, the feature matrix  $X_{d \times L}$  contains

Table 3: The performances of CASCADEUCB1, CASCADELINTS(100), CASCADEKL-UCB, CTS, TS-CASCADE, RANKEDLINTS(0.01), CASCADELINUCB(0.1) and LINTS-CASCADE(0.01) when  $L = 2048$ ,  $K = 4$ . The first column displays the mean and the standard deviation of  $\text{Reg}(T)$  and the second column displays the average running time in seconds. Algorithms which are analyzed in this work are highlighted in bold.

|                            | $\text{Reg}(T)$   | Running time                         |
|----------------------------|---|--------------------------------------|
| CASCADEUCB1                | $4.01 \times 10^4 \pm 4.18 \times 10^1$                   | $5.57 \times 10^1$                   |
| CASCADELINTS(100)          | $2.83 \times 10^4 \pm 2.06 \times 10^3$                   | $2.36 \times 10^1$                   |
| CASCADEKL-UCB              | $3.00 \times 10^4 \pm 1.00 \times 10^4$                   | $5.47 \times 10^1$                   |
| <b>CTS</b>                 | <b><math>1.67 \times 10^4 \pm 3.99 \times 10^2</math></b> | <b><math>1.25 \times 10^2</math></b> |
| <b>TS-Cascade</b>          | <b><math>1.50 \times 10^4 \pm 7.27 \times 10^2</math></b> | <b><math>4.55 \times 10^1</math></b> |
| RANKEDLINTS(0.01)          | $1.32 \times 10^4 \pm 5.03 \times 10^3$                   | $2.64 \times 10^2$                   |
| CASCADELINUCB(0.1)         | $7.71 \times 10^1 \pm 2.57 \times 10^2$                   | $4.19 \times 10^3$                   |
| <b>LinTS-Cascade(0.01)</b> | <b><math>7.12 \times 10^1 \pm 2.29 \times 10^2</math></b> | <b><math>5.08 \times 10^1</math></b> |

a small amount of noisy information of the latent click probabilities, and thus taking into account the paucity of information, any algorithm assuming the linear bandit model cannot reliably identify the optimal items and hence suffers from a large (and often linear) regret. Conversely, when  $m$  is large,  $A_{\text{train}}$  contains a significant amount of information about the latent click probabilities  $\{w(i)\}_{i=1}^L$ . Thus, Algorithm 4 is likely to successfully project the unknown feature vectors  $\{x(i)\}_{i=1}^L \subset \mathbb{R}^d$  onto an informative low-dimensional subspace, and to learn a linear separator in this subspace to classify the optimal and suboptimal items. This enables the algorithms assuming the linear bandit model with a small  $d$  to quickly *commit* to the optimal set  $S^*$  after small amount of exploration. However, in real-world applications, since the historical data is limited, we may not get sufficient training data but may require the algorithm itself to explore the items. Hence, for this problem setting, to illustrate the typical behavior of the algorithms under consideration, we set  $m = 200$  so that the amount of information contained in  $A_{\text{train}}$  is “appropriate”. However, obviously, other choices of  $m$  will lead to other experimental outcomes, but since our objective is to compare the algorithms for the linear bandit model and not to scrutinize their performances as functions of  $m$ , we do not investigate the effect of  $m$  any further.

To begin with, we generate independent Bernoulli random variables with click probabilities  $w_1$ ,  $w_2$  or  $w_3$  for each item to produce the matrix  $A_{\text{train}} \in \{0, 1\}^{m \times L}$  in which  $m = 200$ . Next, we vary the tuning parameters for each algorithm under consideration on a logarithmic scale. We present the best performance of each algorithm for cascading bandits with linear generalization with tuning parameters in the finite set  $\{0.01, 0.1, 1, 10, 100\}$ .

Table 3 displays the means and the standard deviations of  $\text{Reg}(T)$  as well as the average running times. The algorithms are listed in descending order of the average cumulative



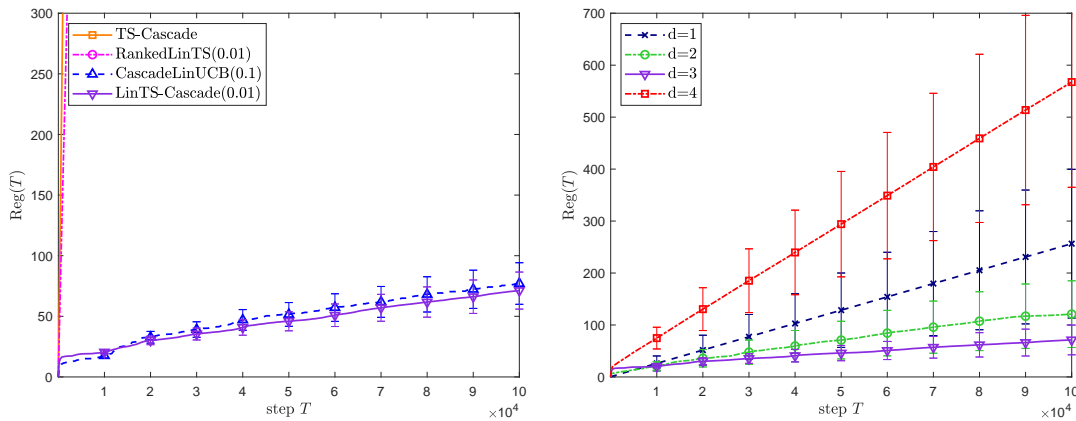


Figure 3: Left:  $\text{Reg}(T)$  of algorithms when  $L = 1024$ ,  $K = 2$ ,  $d = 3$ ; Right:  $\text{Reg}(T)$  of LINTS-CASCADE(0.04) when  $L = 1024$ ,  $K = 2$ ,  $d \in \{1, 2, 3, 4\}$ . Each line indicates the average  $\text{Reg}(T)$  (over 20 runs), and the length of each errorbar above and below each data point is the scaled standard deviation.

regrets. Note that we set  $d = 3$  for the algorithms based on the linear bandit model. We will explain this choice later.

In the left plot of Figure 3, we present the average cumulative regrets of the algorithms. Since CASCADEUCB1, CASCADEKL-UCB, CTS, CASCADELINTS(100) incur much larger regrets compared to the remaining algorithms, we omit their regrets in the plot for clarity. For the algorithms assuming the linear bandit model and for this dataset, our experiments indicate that  $d = 3$  is the most appropriate choice. For clarity, we only display the regrets of the proposed LINTS-CASCADE( $\lambda$ ) algorithm as  $d$  varies from 1 to 4 in the right plot of Figure 3. Figure 3 is a representative plot for all considered algorithms based on the linear bandits. It can be seen that  $d = 3$  results in the best performance in terms of the expected cumulative regret, i.e.,  $d = 3$  minimizes the average regret among the values of  $d$  considered here. Roughly speaking, this implies that the optimal and suboptimal items can be separated when the data is projected onto a three-dimensional subspace (compared to subspaces of other dimensions). When  $d > 3$ , the matrix that contains the extracted features  $X_{d \times L}$  overfits to the historical data and due to this overfitting, our algorithm—as well as other algorithms—incurs a larger regret. Hence, in our comparison to other algorithms for linear bandits, we only consider  $d = 3$  features.

Among the algorithms designed for linear bandits, CASCADELINUCB requires significantly more time to run (as compared to other algorithms) as suggested by Table 3 while RANKEDLINTS lacks theoretical analysis. The reason why CASCADELINUCB is less computationally efficient is that  $x(i)^\top M_t^{-1/2} x(i)$  needs to be calculated for each item  $i \in [L]$  (the definition of  $M_t$  in Zong et al. (2016) is slightly different from that in our paper; see Algorithm 2 therein), whereas we only calculate  $v_t \sqrt{K} M_t^{-1/2} \xi_t$  once for all items at each time step  $t$  (see Line 5 of Algorithm 3). Our proposed LINTS-CASCADE( $\lambda$ ) algorithm is the *only* efficient algorithm with *theoretical guarantees*.

For LINTS-CASCADE( $\lambda$ ), it is essential to pick an appropriate  $\lambda$  to balance exploration and exploitation. Varying  $\lambda$  is equivalent to tuning the variance of the Thompson samples (see Line 5 of Algorithm 3). We note that other algorithms such as CASCADELINUCB and CASCADELINTS (Zong et al., 2016) also contain tuning parameters (called  $\sigma$ ), and thus our algorithm is similar to these other competing algorithms in this aspect. However, in Figure 3, we have displayed results in which the hyperparameters of all algorithms are tuned on a logarithmic scale to be fair.

Lastly, we submit that our LINTS-CASCADE( $\lambda$ ) algorithm does not outperform other algorithms under all problem settings. Indeed, the performances of all algorithms based on the linear bandit model are sensitive to the amount of training data  $A_{\text{train}} \in \{0, 1\}^{m \times L}$  (i.e., the number of rows  $m$ ) and the choice of  $d$ . However, it is worth emphasizing that our proposed algorithm is computationally efficient and possesses appealing (asymptotic) *theoretical guarantees*; cf. Theorem 5.1. Indeed, it is the only Thompson sampling algorithm for linear cascading bandits that admits such theoretical guarantees.

## 8. Summary and Future work

This work presents a comprehensive theoretical analysis of various Thompson sampling algorithms for cascading bandits under the standard and linear models. Firstly, we provide a tighter upper bound on the regret of CTS compared to Huyuk and Tekin (2019). Secondly, the upper bound on the expected regret of TS-CASCADE we derived in Theorem 4.1 matches the state-of-the-art bound based on UCB by Wang and Chen (2017) (which is identical to CASCADEUCB1 in Kveton et al. (2015a)). Thirdly, we propose LINTS-CASCADE( $\lambda$ ) for the linear generalization of the problem and upper bound its regret in Theorem 5.1 by refining proof techniques used for TS-CASCADE. Finally, in Theorem 6.1, we also derive a problem-independent lower bound on the expected cumulative regret using the information-theoretic technique of Auer et al. (2002) as well as a judicious construction of bandit instance. Lastly, we also demonstrate the computational efficiency and efficacy in terms of regret of LINTS-CASCADE( $\lambda$ ).

In Section 1, we analyze the Thompson sampling algorithm with Beta-Bernoulli update and provide a tighter bound than Huyuk and Tekin (2019) in the cascading bandit setting. Our analysis can be readily generalized to the matroid bandit setting which is also widely considered in current literature as in Kveton et al. (2014). Although the CTS algorithm cannot be easily generalized in a straightforward way to a linear cascading bandit setting, it may be a worthy endeavor to extend the other TS-based algorithms (such as TS-CASCADE) introduced in this paper to other generalized linear bandit settings (such as logistic bandits). Next, since the attractiveness of items in real life may change with time, it is reasonable to consider the cascading bandit model with unknown change points of latent click probabilities (Garivier and Moulines, 2011). There are practical real-life implications in designing and analyzing algorithms that take into account the non-stationarity of the latent click probabilities, e.g., the attractiveness of hotels in a city depending on the season. Further, another interesting topic within the realm of multi-armed bandit problems concerns best arm identification (Audibert and Bubeck, 2010). In the cascading model since an *arm* is a *list of items*, the objective here is to minimize the total number of pulls of all arms to identify items that are “close to” or exactly equal to the optimal ones with high probability.

In the future, we plan to design and analyze an algorithm for this purpose assuming the cascading bandit model.

## Acknowledgements

We thank the reviewers for their meticulous reading and insightful comments. This work is partially funded by a National University of Singapore Start-Up Grant (R-266-000-136-133), a Singapore National Research Foundation (NRF) Fellowship (R-263-000-D02-281) and a Singapore Ministry of Education AcRF Tier 1 Grant (R-263-000-E80-114). This research is also supported by the National Research Foundation, Singapore under its AI Singapore Programme (AISG Award No: AISG2-RP-2020-018). Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of National Research Foundation, Singapore.

## References

- Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. In *Proceedings of the 24th Advances in Neural Information Processing Systems*, pages 2312–2320, 2011.
- Marc Abeille and Alessandro Lazaric. Linear Thompson sampling revisited. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, pages 176–184, 2017.
- Milton Abramowitz and Irene A. Stegun. *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. Dover, New York, ninth edition, 1964.
- Shipra Agrawal and Navin Goyal. Analysis of Thompson sampling for the multi-armed bandit problem. In *Proceedings of the 25th Annual Conference on Learning Theory*, pages 39.1–39.26, 2012.
- Shipra Agrawal and Navin Goyal. Thompson sampling for contextual bandits with linear payoffs. In *Proceedings of the 30th International Conference on Machine Learning*, volume 28, pages 127–135, 2013a.
- Shipra Agrawal and Navin Goyal. Further optimal regret bounds for Thompson sampling. In *Proceedings of the 16th International Conference on Artificial Intelligence and Statistics*, pages 99–107, 2013b.
- Shipra Agrawal and Navin Goyal. Near-optimal regret bounds for Thompson sampling. *J. ACM*, 64(5):30:1–30:24, September 2017a.
- Shipra Agrawal and Navin Goyal. Near-optimal regret bounds for Thompson sampling. *Journal of the ACM (JACM)*, 64(5):1–24, 2017b.
- Shipra Agrawal, Vashist Avadhanula, Vineet Goyal, and Assaf Zeevi. Thompson sampling for the MNL-bandit. *arXiv preprint arXiv:1706.00977*, 2017.

- Venkatachalam Anantharam, Pravin Varaiya, and Jean Walrand. Asymptotically efficient allocation rules for the multiarmed bandit problem with multiple plays-Part I: I.I.D. rewards. *IEEE Transactions on Automatic Control*, 32(11):968–976, November 1987.
- Jean-Yves Audibert and Sébastien Bubeck. Best arm identification in multi-armed bandits. In *Proceedings of the 23th Conference on Learning Theory*, pages 41–53, 2010.
- Jean-Yves Audibert, Rémi Munos, and Csaba Szepesvári. Exploration–exploitation tradeoff using variance estimates in multi-armed bandits. *Theoretical Computer Science*, 410(19):1876 – 1902, 2009. Algorithmic Learning Theory.
- Peter Auer, Nicolò Cesa-Bianchi, Yoav Freund, and Robert E. Schapire. The nonstochastic multiarmed bandit problem. *SIAM Journal of Computing*, 32(1):48–77, 2002.
- Sébastien Bubeck, Nicolo Cesa-Bianchi, et al. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends® in Machine Learning*, 5(1):1–122, 2012.
- Olivier Chapelle and Lihong Li. An empirical evaluation of Thompson sampling. In *Proceedings of the 24th Advances in Neural Information Processing Systems*, pages 2249–2257. 2011.
- Wei Chen, Yajun Wang, and Yang Yuan. Combinatorial multi-armed bandit: General framework and applications. In *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of *Proceedings of Machine Learning Research*, pages 151–159, 2013.
- Wang Chi Cheung, Vincent Tan, and Zixin Zhong. A Thompson sampling algorithm for cascading bandits. In *Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics*, pages 438–447, 2019.
- Wei Chu, Lihong Li, Lev Reyzin, and Robert Schapire. Contextual bandits with linear payoff functions. In *Proceedings of the 14th International Conference on Artificial Intelligence and Statistics*, pages 208–214, 2011.
- Richard Combes, M. Sadegh Talebi, Alexandre Proutiere, and Marc Lelarge. Combinatorial bandits revisited. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 2*, pages 2116–2124, 2015.
- Thomas M Cover and Joy A Thomas. *Elements of information theory*. John Wiley & Sons, 2012.
- Nick Craswell, Onno Zoeter, Michael Taylor, and Bill Ramsey. An experimental comparison of click position-bias models. In *Proceedings of the 1st ACM International Conference on Web Search and Data Mining*, pages 87–94, 2008.
- Sarah Filippi, Olivier Cappe, Aurélien Garivier, and Csaba Szepesvári. Parametric bandits: The generalized linear case. In *Proceedings of the 23rd Advances in Neural Information Processing Systems*, pages 586–594, 2010.

- Yi Gai, Bhaskar Krishnamachari, and Rahul Jain. Learning multiuser channel allocations in cognitive radio networks: A combinatorial multi-armed bandit formulation. In *2010 IEEE Symposium on New Frontiers in Dynamic Spectrum (DySPAN)*, pages 1–9, 2010.
- Aurélien Garivier and Eric Moulines. On upper-confidence bound policies for switching bandit problems. In *Proceedings of the 22nd International Conference on Algorithmic Learning Theory*, pages 174–188. Springer, 2011.
- Aditya Gopalan, Shie Mannor, and Yishay Mansour. Thompson sampling for complex online problems. In *Proceedings of the 31st International Conference on Machine Learning*, pages 100–108, 2014.
- Alihan Huyuk and Cem Tekin. Analysis of Thompson sampling for combinatorial multi-armed bandit with probabilistically triggered arms. In *Proceedings of Machine Learning Research*, pages 1322–1330, 2019.
- Sumeet Katariya, Branislav Kveton, Csaba Szepesvari, and Zheng Wen. DCM bandits: Learning to rank with multiple clicks. In *Proceedings of the 33rd International Conference on Machine Learning*, pages 1215–1224, 2016.
- Emilie Kaufmann, Nathaniel Korda, and Rémi Munos. Thompson sampling: An asymptotically optimal finite-time analysis. In *Proceedings of the 23rd International Conference on Algorithmic Learning Theory*, pages 199–213, 2012.
- Junpei Komiyama, Junya Honda, and Hiroshi Nakagawa. Optimal regret analysis of Thompson sampling in stochastic multi-armed bandit problem with multiple plays. In *Proceedings of The 32nd International Conference on Machine Learning*, pages 1152–1161, 2015.
- Yehuda Koren, Robert Bell, and Chris Volinsky. Matrix factorization techniques for recommender systems. *Computer*, (8):30–37, 2009.
- Branislav Kveton, Zheng Wen, Azin Ashkan, Hoda Eydgahi, and Brian Eriksson. Matroid bandits: Fast combinatorial optimization with learning. In *Proceedings of the 30th Conference on Uncertainty in Artificial Intelligence*, pages 420–429, 2014.
- Branislav Kveton, Csaba Szepesvari, Zheng Wen, and Azin Ashkan. Cascading bandits: Learning to rank in the cascade model. In *Proceedings of the 32nd International Conference on Machine Learning*, pages 767–776, 2015a.
- Branislav Kveton, Zheng Wen, Azin Ashkan, and Csaba Szepesvari. Tight regret bounds for stochastic combinatorial semi-bandits. In *Proceedings of the 18th International Conference on Artificial Intelligence and Statistics*, pages 535–543, 2015b.
- Branislav Kveton, Zheng Wen, Azin Ashkan, and Csaba Szepesvári. Combinatorial cascading bandits. In *Proceedings of the 28th International Conference on Neural Information Processing Systems*, pages 1450–1458, 2015c.
- Paul Lagrée, Claire Vernade, and Olivier Cappé. Multiple-play bandits in the position-based model. In *Proceedings of the 29th Advances in Neural Information Processing Systems*, pages 1597–1605. Curran Associates, Inc., 2016.

- Lihong Li, Wei Chu, John Langford, and Robert E. Schapire. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th International Conference on World Wide Web*, pages 661–670, 2010.
- Shuai Li, Baoxiang Wang, Shengyu Zhang, and Wei Chen. Contextual combinatorial cascading bandits. In *Proceedings of the 33rd International Conference on Machine Learning*, pages 1245–1253, 2016.
- Lijing Qin, Shouyuan Chen, and Xiaoyan Zhu. Contextual combinatorial bandit and its application on diversified online recommendation. In *Proceedings of the 14th SIAM International Conference on Data Mining*, pages 461–469. SIAM, 2014.
- Daniel Russo and Benjamin Van Roy. Learning to optimize via posterior sampling. *Mathematics of Operations Research*, 39(4):1221–1243, 2014.
- Daniel Russo, Benjamin Van Roy, Abbas Kazerouni, Ian Osband, and Zheng Wen. A tutorial on Thompson sampling. *Foundations and Trends in Machine Learning*, 11(1): 1–96, 2018.
- William R Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3–4):285–294, 1933.
- Qinshi Wang and Wei Chen. Improving regret bounds for combinatorial semi-bandits with probabilistically triggered arms and its applications. In *Proceedings of the 30th Advances in Neural Information Processing Systems*, pages 1161–1171, 2017.
- Siwei Wang and Wei Chen. Thompson sampling for combinatorial semi-bandits. In *Proceedings of the 35th International Conference on Machine Learning*, pages 5114–5122, 2018.
- Zheng Wen, Branislav Kveton, and Azin Ashkan. Efficient learning in large-scale combinatorial semi-bandits. In *Proceedings of the 32nd International Conference on International Conference on Machine Learning - Volume 37*, pages 1113–1122, 2015.
- Masrour Zoghi, Tomas Tunys, Mohammad Ghavamzadeh, Branislav Kveton, Csaba Szepesvari, and Zheng Wen. Online learning to rank in stochastic click models. In *Proceedings of the 34th International Conference on Machine Learning*, pages 4199–4208, 2017.
- Shi Zong, Hao Ni, Kenny Sung, Nan Rosemary Ke, Zheng Wen, and Branislav Kveton. Cascading bandits for large-scale recommendation problems. In *Proceedings of the 32nd Conference on Uncertainty in Artificial Intelligence*, pages 835–844, 2016.

**Appendix A. Notations**

|   |   |
|---|---|
| $[L]$   | ground set of size $L$  |
| $w(i)$  | click probability of item $i \in [L]$   |
| $x(i)$  | feature vector of item $i$ in the linear generalization                               |
| $\beta$   | common latent vector in the linear generalization                                     |
| $T$   | time horizon  |
| $d$   | number of features of any item in the linear generalization                           |
| $K$   | size of recommendation list   |
| $S_t$   | recommendation list at time $t$   |
| $i_t^i$   | $i$ -th recommended item at time $t$  |
| $[L]^{(K)}$                                     | set of all $K$ -permutations of $[L]$   |
| $W_t(i)$  | reflection of the interest of the user in item $t$ at time $t$                        |
| $R(S_t \mathbf{w})$                             | instantaneous reward of the agent at time $t$   |
| $k_t$   | feedback from the user at time $t$  |
| $r(S_t \mathbf{w})$                             | expected instantaneous reward of the agent at time $t$                                |
| $S^*$   | optimal ordered $K$ -subset for maximizing the expected instantaneous reward          |
| $\text{Reg}(T)$                                 | expected cumulative regret  |
| $\hat{w}_t(i)$                                  | empirical mean of item $i$ at time $t$  |
| $N_t(i)$  | number of observations of item $i$ before time $t$ plus 1                             |
| $Z_t$   | 1-dimensional random variable drawn from the standard normal distribution at time $t$ |
| $\hat{\nu}_t(i)$                                | empirical variance of item $i$ at time $t$  |
| $\sigma_t(i)$                                   | standard deviation of Thompson sample of item $i$ at time $t$ in TS-CASCADE           |
| $\theta_t(i)$                                   | Thompson sample of item $i$ at time $t$ in TS-CASCADE                                 |
| $\alpha_t^{\text{B}}(i), \beta_t^{\text{B}}(i)$ | statistics used in CTS  |
| $\theta_t^{\text{B}}(i)$                        | Thompson sample from Beta distribution at time $t$ in CTS                             |

|  |   |
|--|---|
| $\Delta_{i,j}$   | gap between click probabilities of optimal item $i$ and suboptimal item $j$   |
| $\Delta = \Delta_{K,K+1}$  | minimum gap between click probabilities of optimal and suboptimal items   |
| $\pi_t$  | a permutation of optimal items used in the analysis of CTS  |
| $\mathcal{A}_t(\ell, j, k), \tilde{\mathcal{A}}_t(\ell, j, k), \mathcal{A}_t(\ell, j)$ | events defined to help decompose the regret of CTS  |
| $N^B(i, j)$  | number of time steps that $\mathcal{A}_t(i, j)$ happens for optimal item $j$ and suboptimal item $i$ within the whole horizon   |
| $\mathcal{B}_t(i), \mathcal{C}_t(i, j)$  | events that measure the concentration of $\theta_t^B(i), \hat{w}_t(i)$ around $w(i)$ as CTS processes                           |
| $\theta^B(-i), W_{i,j}, W_{i,-j}$  | variables used in the analysis of CTS   |
| $\mathcal{E}_{\hat{w},t}, \mathcal{E}_{\theta,t}$                                      | “nice events” in TS-CASCADE   |
| $g_t(i), h_t(i)$   | functions used to define “nice events”  |
| $\mathcal{S}_t$  | “unsaturated” set at time $t$   |
| $F(S, t)$  | function used to define “unsaturated” set   |
| $\alpha$   | determined by click probabilities $w(i)$ ’s, used in the analysis   |
| $\mathcal{H}_{\hat{w},t}, \mathcal{F}_{i,t}, G(S_t, \mathbf{W}_t)$                     | used in the analysis of TS-CASCADE  |
| $\lambda$  | parameter in LINTS-CASCADE( $\lambda$ )   |
| $X$  | collection of feature vector $x(i)$ ’s  |
| $b_t, M_t, v_t$  | statistics used in LINTS-CASCADE( $\lambda$ )   |
| $\xi_t$  | $d$ -dimensional random variable drawn from standard multivariate normal distribution at time $t$ in LINTS-CASCADE( $\lambda$ ) |
| $\hat{\psi}_t$   | empirical mean of latent vector $\beta$ at time $t$   |
| $\rho_t$   | Thompson sample at time $t$ in TS-CASCADE   |
| $\eta_t(i)$  | random variable in the linear generalization  |
| $\mathcal{E}_{\hat{\psi},t}, \mathcal{E}_{\rho,t}$                                     | “nice events” in LINTS-CASCADE( $\lambda$ )   |
| $l_t, u_t(i)$  | functions used to define “nice events” in LINTS-CASCADE( $\lambda$ )  |
| $w^{(\ell)}(i)$  | click probability of item $i \in [L]$ under instance $\ell$ ( $0 \leq \ell \leq L$ )  |
| $r(S_t   \mathbf{w}^{(\ell)})$   | expected instantaneous reward of the agent at time $t$ under instance $\ell$  |



|   |   |
|---|---|
| $S^{*,\ell}$                            | optimal ordered $K$ -subset under instance $\ell$   |
| $\pi$                                   | deterministic and non-anticipatory policy   |
| $S_t^{\pi,(\ell)}$                      | recommendation list at time $t$ (under instance $\ell$ ) by policy $\pi$                              |
| $O_t^{\pi,(\ell)}$                      | stochastic outcome by pulling $S_t^{\pi,(\ell)}$ at time $t$ (under instance $\ell$ ) by policy $\pi$ |
| $A$                                     | feedback matrix extracted from real data  |
| $y_t, Q_t, \zeta_t, \hat{\sigma}_t$     | variables used in the analysis of LINTS-CASCADE( $\lambda$ )  |
| $Q, J_t^{\pi,\ell}, \Upsilon_t^\ell(S)$ | variables used in the analysis of the lower bound   |

## Appendix B. Useful theorems

Here are some basic facts from the literature that we will use:

### Theorem B.1 (Audibert et al. (2009), specialized to Bernoulli random variables)

Consider  $N$  independently and identically distributed Bernoulli random variables  $Y_1, \dots, Y_N \in \{0, 1\}$ , which have the common mean  $m = \mathbb{E}[Y_1]$ . In addition, consider their sample mean  $\hat{\xi}$  and their sample variance  $\hat{V}$ :

$$\hat{\xi} = \frac{1}{N} \sum_{i=1}^N Y_i, \quad \hat{V} = \frac{1}{N} \sum_{i=1}^N (Y_i - \hat{\xi})^2 = \hat{\xi}(1 - \hat{\xi}).$$

For any  $\delta \in (0, 1)$ , the following inequality holds:

$$\Pr \left( \left| \hat{\xi} - m \right| \leq \sqrt{\frac{2\hat{V} \log(1/\delta)}{N}} + \frac{3 \log(1/\delta)}{N} \right) \geq 1 - 3\delta.$$

### Theorem B.2 (Abramowitz and Stegun (1964), Formula 7.1.13) Let $Z \sim \mathcal{N}(\mu, \sigma^2)$ .

The following inequalities hold:

$$\begin{aligned} \frac{1}{2\sqrt{\pi}} \exp\left(-\frac{7z^2}{2}\right) &\leq \Pr(|Z - \mu| > z\sigma) \leq \exp\left(-\frac{z^2}{2}\right) \quad \forall z > 0, \\ \frac{1}{2\sqrt{\pi}z} \exp\left(-\frac{z^2}{2}\right) &\leq \Pr(|Z - \mu| > z\sigma) \leq \frac{\sqrt{2}}{\sqrt{\pi}z} \exp\left(-\frac{z^2}{2}\right) \quad \forall z \geq 1. \end{aligned}$$

**Theorem B.3** (Abbasi-Yadkori et al., 2011; Agrawal and Goyal, 2013a). Let  $(\mathcal{F}'_t; t \geq 0)$  be a filtration,  $(m_t; t \geq 1)$  be an  $\mathbb{R}^d$ -valued stochastic process such that  $m_t$  is  $(\mathcal{F}'_{t-1})$ -measurable,  $(\eta_t; t \geq 1)$  be a real-valued martingale difference process such that  $\eta_t$  is  $(\mathcal{F}'_t)$ -measurable. For  $t \geq 0$ , define  $\xi_t = \sum_{\tau=1}^t m_\tau \eta_\tau$  and  $M_t = I_d + \sum_{\tau=1}^t m_\tau m_\tau^\top$ , where  $I_d$  is the  $d$ -dimensional identity matrix. Assume  $\eta_t$  is conditionally  $R$ -sub-Gaussian.

Then, for any  $\delta' > 0$ ,  $t \geq 0$ , with probability at least  $1 - \delta'$ ,

$$\|\xi_t\|_{M_t^{-1}} \leq R \sqrt{d \log\left(\frac{t+1}{\delta'}\right)},$$

where  $\|\xi\|_{M_t^{-1}} = \sqrt{\xi_t^\top M_t^{-1} \xi_t}$ .

## Appendix C. Proofs of main results

In this Section, we provide proofs of Lemmas 3.2, 3.3, Theorem 3.4, Lemmas 4.2, 4.3, 4.5, 6.3, 6.4, 6.5 as well as derivations of inequalities (3.9) and (5.7).

### C.1 Proof of Lemma 3.2

**Proof** We let  $\tau_1, \tau_2, \dots$  be the time slots such that  $i \in S(t)$  and  $W_t(i)$  is observed, define  $\tau_0 = 0$ , and abbreviate  $\varepsilon_i$  as  $\varepsilon$  for simplicity. Then

$$\begin{aligned}
 & \mathbb{E} [ \exists t : 1 \leq t \leq T, i \in S_t, |\hat{w}_t(i) - w(i)| > \varepsilon, \text{Observe } W_t(i) \text{ at } t ] \\
 &= \mathbb{E} \left[ \sum_{t=1}^T \mathbf{1} [i \in S_t, |\hat{w}_t(i) - w(i)| > \varepsilon, \text{Observe } W_t(i) \text{ at } t] \right] \\
 &\leq \mathbb{E} \left[ \sum_{s=0}^T \mathbb{E} \left[ \sum_{t=\tau_s}^{\tau_{s+1}-1} \mathbf{1} [i \in S_t, |\hat{w}_t(i) - w(i)| > \varepsilon, \text{Observe } W_t(i) \text{ at } t] \right] \right] \\
 &\leq \mathbb{E} \left[ \sum_{s=0}^T \Pr [|\hat{w}_t(i) - w(i)| > \varepsilon, N_t(i) = s] \right] \\
 &\leq 1 + \sum_{s=1}^T \Pr [|\hat{w}_t(i) - w(i)| > \varepsilon, N_t(i) = s] \\
 &\leq 1 + \sum_{s=1}^T \exp[-w\text{KL}(w(i) + \varepsilon, w(i))] + \sum_{s=1}^T \exp[-w\text{KL}(w(i) - \varepsilon, w(i))] \\
 &\leq 1 + \sum_{s=1}^{\infty} \{\exp[-\text{KL}(w(i) + \varepsilon, w(i))]\}^s + \sum_{w=1}^{\infty} \{\exp[-\text{KL}(w(i) - \varepsilon, w(i))]\}^s \\
 &\leq 1 + \frac{\exp[-\text{KL}(w(i) + \varepsilon, w(i))]}{1 - \exp[-\text{KL}(w(i) + \varepsilon, w(i))]} + \frac{\exp[-\text{KL}(w(i) - \varepsilon, w(i))]}{1 - \exp[-\text{KL}(w(i) - \varepsilon, w(i))]} \\
 &\leq 1 + \frac{1}{\text{KL}(w(i) + \varepsilon, w(i))} + \frac{1}{\text{KL}(w(i) - \varepsilon, w(i))} \\
 &\leq 1 + \frac{1}{\varepsilon^2}.
 \end{aligned}$$

■

### C.2 Upper bound on the regret incurred by the second term of (3.7) in Section 3

To derive the inequality (3.9): for any  $0 < \varepsilon < \Delta_{K,K+1}/2$ ,

$$\begin{aligned} & \sum_{i=K+1}^L \sum_{j=1}^K \Delta_{j,i} \cdot \mathbb{E} \left[ \sum_{t=1}^T \mathbf{1} [\mathcal{A}_t(i, j) \wedge \neg \mathcal{B}_t(i) \wedge \mathcal{C}_t(i, j)] \right] \\ & \leq 4 \log T \cdot \sum_{i=K+1}^L \frac{\Delta_{K,i} - \varepsilon}{(\Delta_{K,i} - 2\varepsilon)^2} + \sum_{i=K+1}^L \sum_{j=1}^K \Delta_{j,i}, \end{aligned}$$

we use the following result:

**Lemma C.1 (Implied by Wang and Chen (2018), Lemma 4)** *In Algorithm 1, for any base item  $i$ , we have the following two inequalities:*

$$\Pr \left[ \theta_i^{\text{B}}(t) - \hat{w}_t(i) > \sqrt{\frac{2 \log T}{N_t(i)}} \right] \leq \frac{1}{T}, \quad \Pr \left[ \hat{w}_t(i) - \theta_i^{\text{B}}(t) > \sqrt{\frac{2 \log T}{N_t(i)}} \right] \leq \frac{1}{T}.$$

Note that

$$\varepsilon_i = \frac{1}{2} \cdot \left( \Delta_{K+1,i} + \frac{\Delta_{K,K+1}}{2} \right) \text{ and } \varepsilon_j = \frac{1}{2} \cdot \left( \Delta_{j,K} + \frac{\Delta_{K,K+1}}{2} \right)$$

implies  $\varepsilon_i + \varepsilon_j = \Delta_{j,i}/2$ . Let

$$M_{i,j} = \frac{2 \log T}{(\Delta_{j,i} - \varepsilon_i - \varepsilon_j)^2} = \frac{8 \log T}{\Delta_{j,i}^2},$$

then when  $N_t(i) > M_{i,j}$ ,  $\Delta_{j,i} - \varepsilon_i - \varepsilon_j > \sqrt{\frac{2 \log T}{N_t(i)}}$ . Thus in expectation, the time steps that  $\mathcal{A}_t(i, j) \wedge \neg \mathcal{B}_t(i) \wedge \mathcal{C}_t(i, j)$  occurs can be decomposed as follows:

$$\begin{aligned} & \mathbb{E} \left[ \sum_{t=1}^T \mathbf{1} [\mathcal{A}_t(i, j) \wedge \neg \mathcal{B}_t(i) \wedge \mathcal{C}_t(i, j)] \right] \\ & = \mathbb{E} \left[ \sum_{t=1}^T \mathbf{1} [\mathcal{A}_t(i, j) \wedge \neg \mathcal{B}_t(i) \wedge \mathcal{C}_t(i, j), N_t(i) \leq M_{i,j}] \right] \\ & \quad + \mathbb{E} \left[ \sum_{t=1}^T \mathbf{1} [\mathcal{A}_t(i, j) \wedge \neg \mathcal{B}_t(i) \wedge \mathcal{C}_t(i, j), N_t(i) > M_{i,j}] \right]. \end{aligned}$$

And further

$$\begin{aligned} & \mathbb{E} \left[ \sum_{t=1}^T \mathbf{1} [\mathcal{A}_t(i, j) \wedge \neg \mathcal{B}_t(i) \wedge \mathcal{C}_t(i, j), N_t(i) > M_{i,j}] \right] \\ & \leq \mathbb{E} \left[ \sum_{t=1}^T \mathbf{1} [\hat{w}_t(i) \leq w(i) + \varepsilon_i, \theta_t^{\text{B}}(i) > w(j) - \varepsilon_j, N_t(i) > M_{i,j}, \text{Observe } W_t(i) \text{ at } t] \right] \end{aligned}$$

$$\leq \mathbb{E} \left[ \sum_{t=1}^T \mathbf{1} \left[ \theta_t^{\text{B}}(i) > \hat{w}_t(i) + \Delta_{j,i} - \varepsilon_i - \varepsilon_j, N_t(i) > M_{i,j}, \text{Observe } W_t(i) \text{ at } t \right] \right] \quad (\text{C.1})$$

$$\leq \mathbb{E} \left[ \sum_{t=1}^T \mathbf{1} \left[ \theta_t^{\text{B}}(i) > \hat{w}_t(i) + \sqrt{\frac{2 \log T}{N_t(i)}} \right] \right] \quad (\text{C.2})$$

$\leq 1$ .

To derive line (C.1), we observe that when  $\hat{w}_t(i) \leq w(i) + \varepsilon_i$ ,  $\theta_t^{\text{B}}(i) > w(j) - \varepsilon_j$ , we have

$$\theta_t^{\text{B}}(i) > w(j) - \varepsilon_j = w(i) + \Delta_{j,i} - \varepsilon_j \geq \hat{w}_t(i) + \Delta_{j,i} - \varepsilon_i - \varepsilon_j.$$

Besides, line (C.2) follows from the definition of  $M_{i,j}$ .

Hence,

$$\begin{aligned} & \sum_{j=1}^K \Delta_{j,i} \cdot \mathbb{E} \left[ \sum_{t=1}^T \mathbf{1} [\mathcal{A}_t(i,j) \wedge \neg \mathcal{B}_t(i) \wedge \mathcal{C}_t(i,j)] \right] \\ & \leq \underbrace{\sum_{j=1}^K \Delta_{j,i} \cdot \mathbb{E} \left[ \sum_{t=1}^T \mathbf{1} [\mathcal{A}_t(i,j), N_t(i) \leq M_{i,j}] \right]}_{(*)} + \sum_{j=1}^K \Delta_{j,i}. \end{aligned}$$

Note that (i)  $N_t(i)$  of suboptimal item  $i$  increases by one when the event  $\mathcal{A}_{i,j,t}$  happens for any optimal item  $j$ , (ii) the event  $\mathcal{A}_t(i,j)$  happens for at most one optimal item  $i$  at any time  $t$ ; and (iii)  $M_{1,i} \leq \dots \leq M_{K,i}$ . Let

$$Y_{i,j} = \sum_{t=1}^T \mathbf{1} [\mathcal{A}_t(i,j), N_t(i) \leq M_{i,j}].$$

Based on these facts, it follows that  $Y_{i,j} \leq M_{i,j}$  almost surely, and moreover  $\sum_{j=1}^K Y_{i,j} \leq M_{K,i}$  almost surely. Therefore, (\*) can be bounded from above by:

$$\max \left\{ \sum_{j=1}^K \Delta_{j,i} y_{i,j} : 0 \leq y_{i,j} \leq M_{i,j}, \sum_{i=1}^K y_{i,j} \leq M_{K,j} \right\}.$$

Since the gaps are decreasing,  $\Delta_{1,i} \geq \Delta_{2,i} \geq \Delta_{K,i}$ , the solution to the above problem is  $y_{1,i}^* = M_{1,i}$ ,  $y_{2,i}^* = M_{2,i} - M_{1,i}$ ,  $\dots$ ,  $y_{K,i}^* = M_{K,i} - M_{K-1,i}$ . Therefore, the value of (\*) is bounded from above by:

$$8 \log T \left[ \Delta_{1,i} \frac{1}{\Delta_{1,i}^2} + \sum_{j=2}^K \Delta_{j,i} \left( \frac{1}{\Delta_{j,i}^2} - \frac{1}{\Delta_{j-1,i}^2} \right) \right].$$

Further, we use the following result:

**Lemma C.2 (Kveton et al. (2014), Lemma 3)** *Let  $\Delta_1 \geq \dots \geq \Delta_K$  be a sequence of  $K$  positive numbers, then*

$$\left[ \Delta_1 \frac{1}{\Delta_1^2} + \sum_{k=2}^K \Delta_k \left( \frac{1}{\Delta_k^2} - \frac{1}{\Delta_{k-1}^2} \right) \right] \leq \frac{2}{\Delta_K}.$$

Therefore,

$$\sum_{j=1}^K \Delta_{j,i} \cdot \mathbb{E} \left[ \sum_{t=1}^T \mathbf{1}[\mathcal{A}_t(i,j) \wedge \neg \mathcal{B}_t(i) \wedge \mathcal{C}_t(i,j)] \right] \leq \frac{16 \log T}{\Delta_{K,i}} + \sum_{j=1}^K \Delta_{j,i}.$$

Lastly, the total regret incurred by the second term in equation (3.7) can be bounded as follows:

$$\sum_{i=K+1}^L \sum_{j=1}^K \Delta_{j,i} \cdot \mathbb{E} \left[ \sum_{t=1}^T \mathbf{1}[\mathcal{A}_t(i,j) \wedge \neg \mathcal{B}_t(i) \wedge \mathcal{C}_t(i,j)] \right] \leq 16 \log T \cdot \sum_{i=K+1}^L \frac{1}{\Delta_{K,i}} + \sum_{i=K+1}^L \sum_{j=1}^K \Delta_{j,i}.$$

### C.3 Proof of Lemma 3.3

**Proof** We abbreviate  $\varepsilon_j$  as  $\varepsilon$  for simplicity. By the definition of  $W_{i,j}$  and  $W_{i,-j}$ ,  $\theta_t^B \in W_j$  yields that there exists  $i$  such that  $\{\theta_t^B(i) \leq w(j) - \varepsilon, \theta_t^B(-j) \in W_{i,-j}\}$ , thus we can use

$$\mathbb{E} \left[ \sum_{t=1}^T \Pr \left[ \exists i \text{ s.t. } \theta_t^B(i) \leq w(j) - \varepsilon, \theta_t^B(-j) \in W_{i,-j}, \text{ Observe } W_t(i) \text{ at } t \mid \mathcal{F}_{t-1} \right] \right]$$

as an upper bound. Since  $\theta_t^B \in W_j$  indicates  $\theta_t^B(i) \geq \theta_t^B(j)$ , we have

$$\begin{aligned} & \Pr \left[ \exists i \text{ s.t. } \theta_t^B(i) \leq w(j) - \varepsilon, \theta_t^B(-j) \in W_{i,-j}, \text{ Observe } W_t(i) \text{ at } t \mid \mathcal{F}_{t-1} \right] \\ & \leq \Pr \left[ \exists i \text{ s.t. } \theta_t^B(j) \leq w(j) - \varepsilon, \theta_t^B(-j) \in W_{i,-j}, \text{ Observe } W_t(i) \text{ at } t \mid \mathcal{F}_{t-1} \right]. \end{aligned}$$

Denote the value  $p_{j,t}$  as  $\Pr[\theta_t^B(j) > w(j) - \varepsilon \mid \mathcal{F}_{t-1}]$ . Notice that according to the cascading model, given  $\mathcal{F}_{t-1}$ , the value  $\theta_t^B(j)$  and the set of values  $\theta_t^B(-j)$  are independent. Besides, note that  $\neg \mathcal{C}_{i,j}(t) = \{\theta_t^B(i) \leq w(j) - \varepsilon\}$  holds when  $\theta_t^B(-j) \in W_{i,-j}$ . If we further have  $\theta_t^B(j) > w(j) - \varepsilon$ , then  $j \in S_t$ .

**Observation of fixed  $i$ .** Fix  $\theta_t^B(-j) \in W_{i,-j}$ , let

- $S_t = \{i_1^t, i_2^t, \dots, i_K^t\}$ ,  $i_{k_1}^t = i$  when  $\theta_t^B(j) \leq w(j) - \varepsilon$ ,
- $\bar{S}_t = \{\bar{i}_1^t, \bar{i}_2^t, \dots, \bar{i}_K^t\}$ ,  $\bar{i}_{k_2}^t = j$  when  $\theta_t^B(j)$  changes to  $\bar{\theta}_t^B(j)$  such that  $\bar{\theta}_t^B(j) > w(j) - \varepsilon$ .

Since  $\theta_t^B(i_1^t) \geq \theta_t^B(i_2^t) \geq \dots, \theta_t^B(\bar{i}_1^t) \geq \theta_t^B(\bar{i}_2^t) \geq \dots$  according to Algorithm 1,

$$\bar{\theta}_t^B(j) > w(j) - \varepsilon \geq \theta_t^B(i) = \theta_t^B(i_{k_1}^t) \geq \theta_t^B(i_{k_1+1}^t) \geq \dots$$

Therefore,  $k_2 \leq k_1$  and  $i_l^t = \bar{i}_l^t$  for  $1 \leq l < k_2$ . Besides, the observation of item  $i$  or  $j$  is determined by  $\theta_t^B(-j)$

With the observation and notation above, we have

$$\begin{aligned} & \Pr \left[ \theta_t^B(j) \leq w(j) - \varepsilon, \theta_t^B(-j) \in W_{i,-j}, \text{ Observe } W_t(i) \text{ at } t \mid \mathcal{F}_{t-1} \right] \\ & = (1 - p_{j,t}) \cdot \Pr \left[ \theta_t^B(-j) \in W_{i,-j} \mid \mathcal{F}_{t-1} \right] \cdot \Pr \left[ \text{Observe } W_t(i) \text{ at } t \mid \theta_t^B \right] \\ & = (1 - p_{j,t}) \cdot \Pr \left[ \theta_t^B(-j) \in W_{i,-j} \mid \mathcal{F}_{t-1} \right] \cdot \mathbb{E} \left[ \prod_{l=1}^{k_1-1} (1 - w(i_l^t)) \mid \theta_t^B(-j) \right], \end{aligned}$$

and

$$\begin{aligned}
 & \Pr [j \in S_t, \theta_t^{\text{B}}(-j) \in W_{i,-j}, \text{Observe } W_t(j) \text{ at } t | \mathcal{F}_{t-1}] \\
 & \geq \Pr [\theta_t^{\text{B}}(j) > w(j) - \varepsilon, \theta_t^{\text{B}}(-j) \in W_{i,-j}, \text{Observe } W_t(j) \text{ at } t | \mathcal{F}_{t-1}] \\
 & = p_{j,t} \cdot \Pr [\theta_t^{\text{B}}(-j) \in W_{i,-j} | \mathcal{F}_{t-1}] \cdot \Pr [\text{Observe } W_t(j) \text{ at } t | \theta_t^{\text{B}}(-j)] \\
 & = p_{j,t} \cdot \Pr [\theta_t^{\text{B}}(-j) \in W_{i,-j} | \mathcal{F}_{t-1}] \cdot \mathbb{E} \left[ \prod_{l=1}^{k_2-1} (1 - w(i_l^t)) \mid \theta_t^{\text{B}}(-j) \right] \\
 & \geq p_{j,t} \cdot \Pr [\theta_t^{\text{B}}(-j) \in W_{i,-j} | \mathcal{F}_{t-1}] \cdot \mathbb{E} \left[ \prod_{l=1}^{k_1-1} (1 - w(i_l^t)) \mid \theta_t^{\text{B}}(-j) \right].
 \end{aligned}$$

This leads to

$$\begin{aligned}
 & \Pr [\theta_t^{\text{B}} \in W_{i,j}, \text{Observe } W_t(i) \text{ at } t | \mathcal{F}_{t-1}] \\
 & \leq (1 - p_{j,t}) \cdot \Pr [\theta_t^{\text{B}}(-j) \in W_{i,-j} | \mathcal{F}_{t-1}] \cdot \mathbb{E} \left[ \prod_{l=1}^{k_1-1} (1 - w(i_l^t)) \mid \theta_t^{\text{B}}(-j) \right] \\
 & \leq \frac{1 - p_{j,t}}{p_{j,t}} \cdot \Pr [j \in S_t, \theta_t^{\text{B}}(-j) \in W_{i,-j}, \text{Observe } W_t(j) \text{ at } t | \mathcal{F}_{t-1}].
 \end{aligned}$$

**Observation in the whole horizon.** By definition,  $W_{i,j}$  are disjoint sets for different  $i$ . Hence,

$$\begin{aligned}
 & \mathbb{E} \left[ \sum_{t=1}^T \Pr [\exists i \text{ s.t. } \theta_t^{\text{B}} \in W_{i,j}, \text{Observe } W_t(i) \text{ at } t | \mathcal{F}_{t-1}] \right] \\
 & = \sum_{i=K+1}^L \mathbb{E} \left[ \sum_{t=1}^T \Pr [\theta_t^{\text{B}} \in W_{i,j}, \text{Observe } W_t(i) \text{ at } t | \mathcal{F}_{t-1}] \right] \\
 & \leq \sum_{t=1}^T \mathbb{E} \left[ \sum_{i=K+1}^L \mathbb{E} \left[ \frac{1 - p_{j,t}}{p_{j,t}} \cdot \Pr [j \in S_t, \theta_t^{\text{B}}(-j) \in W_{i,-j}, \text{Observe } W_t(j) \text{ at } t | \mathcal{F}_{t-1}] \right] \right] \\
 & = \sum_{t=1}^T \mathbb{E} \left[ \frac{1 - p_{j,t}}{p_{j,t}} \cdot \Pr [\exists i \text{ s.t. } j \in S_t, \theta_t^{\text{B}}(-j) \in W_{i,-j}, \text{Observe } W_t(j) \text{ at } t | \mathcal{F}_{t-1}] \right] \\
 & \leq \sum_{t=1}^T \mathbb{E} \left[ \frac{1 - p_{j,t}}{p_{j,t}} \cdot \Pr [j \in S_t, \text{Observe } W_t(j) \text{ at } t | \mathcal{F}_{t-1}] \right] \\
 & = \sum_{t=1}^T \mathbb{E} \left[ \frac{1 - p_{j,t}}{p_{j,t}} \cdot \mathbf{1}[j \in S_t, \text{Observe } W_t(j) \text{ at } t | \mathcal{F}_{t-1}] \right].
 \end{aligned}$$

Let  $\tau_{j,q}$  is the time step that item  $j$  is observed for the  $q$ -th time. Then

$$\begin{aligned} & \mathbb{E} \left[ \sum_{t=1}^T \Pr \left[ \exists i \text{ s.t. } \boldsymbol{\theta}_t^{\mathbf{B}} \in W_{i,j}, \text{ Observe } W_t(i) \text{ at } t \mid \mathcal{F}_{t-1} \right] \right] \\ & \leq \sum_{q=0}^{T-1} \mathbb{E} \left[ \frac{1 - p_{j,\tau_{j,q}+1}}{p_{j,\tau_{j,q}+1}} \sum_{t=\tau_{j,q}+1}^{\tau_{j,q}+1} \mathbf{1}[j \in S_t, \text{ Observe } W_t(j) \text{ at } t \mid \mathcal{F}_{t-1}] \right] \quad (\text{C.3}) \\ & = \sum_{q=0}^{T-1} \mathbb{E} \left[ \frac{1 - p_{j,\tau_{j,q}+1}}{p_{j,\tau_{j,q}+1}} \right]. \end{aligned}$$

Line (C.3) follows from the fact that the probability  $\Pr [\theta_t^{\mathbf{B}}(j) > w(j) - \varepsilon \mid \mathcal{F}_{t-1}]$  only changes when we get a feedback of item  $j$ , but during time steps in  $[\tau_{j,q} + 1, \tau_{j,q+1} - 1]$ , we do not have any such feedback.

Next, we apply a variation of Lemma 2.9 in Agrawal and Goyal (2017b):

**Lemma C.3 (Implied by Agrawal and Goyal (2017b), Lemma 2.9)** *For any  $j$ ,  $p_{j,t}$  is  $\Pr [\theta_t^{\mathbf{B}}(j) > w(j) - \varepsilon \mid \mathcal{F}_{t-1}]$ , and  $\tau_{j,q}$  is the time step that item  $j$  is observed for the  $q$ -th time. Then <sup>2</sup>*

$$\mathbb{E} \left[ \frac{1}{p_{j,\tau_{j,q}+1}} - 1 \right] \leq \begin{cases} \frac{3}{\varepsilon} & \text{if } q < \frac{8}{\varepsilon}, \\ \Theta \left( e^{-\varepsilon^2 q/2} + \frac{1}{(q+1)\varepsilon^2} e^{-D_j q} + \frac{1}{e^{\varepsilon^2 q/4} - 1} \right) & \text{else,} \end{cases}$$

where  $D_j = \text{KL}(w(j), w(j) - \varepsilon) = [w(j) - \varepsilon] \cdot \log \left[ \frac{w(j) - \varepsilon}{w(j)} \right] + [1 - w(j) + \varepsilon] \cdot \log \left[ \frac{1 - w(j) + \varepsilon}{1 - w(j)} \right]$ .

According to Pinsker's and reverse Pinsker's inequality for any two distributions  $P$  and  $Q$  defined in the same finite space  $X$ , we have

$$\delta(P, Q)^2 \leq \frac{1}{2} \text{KL}(P, Q) \leq \frac{1}{\alpha_Q} \delta(P, Q)^2$$

where  $\delta(P, Q) = \sup\{|P(A) - Q(A)| \mid A \subset X\}$  and  $\alpha_Q = \min_{x \in X: Q(x) > 0} Q(x)$ . Therefore,

$$D_j \geq 2\{w(j) - [w(j) - \varepsilon]\}^2 = 2\varepsilon^2.$$

Hence, there exists a constant  $0 < c_1 < \infty$  such that

$$\begin{aligned} & \sum_{q=0}^{T-1} \mathbb{E} \left[ \frac{1 - p_{j,\tau_{j,q}+1}}{p_{j,\tau_{j,q}+1}} \right] \\ & \leq \sum_{0 \leq q < 8/\varepsilon} \frac{3}{\varepsilon} + \sum_{T-1 \geq q \geq 8/\varepsilon} c_1 \cdot \left[ e^{-\varepsilon^2 q/2} + \frac{1}{(q+1)\varepsilon^2} e^{-2\varepsilon^2 q} + \frac{1}{e^{\varepsilon^2 q/4} - 1} \right] \\ & \leq \frac{24}{\varepsilon^2} + c_1 \cdot \int_{8/\varepsilon-1}^{T-1} e^{-\varepsilon^2 x/2} + \frac{1}{(x+1)\varepsilon^2} e^{-2\varepsilon^2 x} + \frac{1}{e^{\varepsilon^2 x/4} - 1} dx \end{aligned}$$

2. A (non-negative) function  $f(T) = O(T)$  if there exists a constant  $0 < c < \infty$  (dependent on  $w$  but not  $T$ ) such that  $f(T) \leq cT$  for sufficiently large  $T$ . Similarly  $f(T) = \Omega(T)$  if there exists  $c > 0$  such that  $f(T) \geq cT$  for sufficiently large  $T$ . Finally,  $f(T) = \Theta(T)$  if  $f(T) = O(T)$  and  $f(T) = \Omega(T)$ .

Note the Taylor series of  $e^x$ :  $e^x = \sum_{n=0}^{\infty} x^n/n!$  ( $0! = 1$ ). We have

$$\begin{aligned}
 \text{(i)} \quad & \int_{8/\varepsilon-1}^{T-1} e^{-\varepsilon^2 x/2} dx = -\frac{2}{\varepsilon^2} e^{-\varepsilon^2 x/2} \Big|_{x=8/\varepsilon-1}^{T-1} = \frac{2}{\varepsilon^2} e^{\varepsilon^2/2} \cdot \left[ e^{-\varepsilon^2 T/2} - e^{-\varepsilon^2(8/\varepsilon)/2} \right] \leq \frac{2\sqrt{e}}{\varepsilon^2}, \\
 \text{(ii)} \quad & \int_{8/\varepsilon-1}^{T-1} \frac{1}{(x+1)\varepsilon^2} e^{-2\varepsilon^2 x} dx = e^{\varepsilon^2/2} \cdot \int_{8/\varepsilon-1}^{T-1} \frac{1}{(x+1)\varepsilon^2} e^{-2\varepsilon^2(x+1)} dx \\
 & = e^{\varepsilon^2/2} \cdot \int_{8/\varepsilon}^T \frac{1}{x\varepsilon^2} e^{-2\varepsilon^2 x} dx = e^{\varepsilon^2/2} \cdot \int_{16\varepsilon}^{2T\varepsilon^2} \frac{e^{-x}}{x/2} d\left(\frac{x}{2\varepsilon^2}\right) = \frac{e^{\varepsilon^2/2}}{\varepsilon^2} \cdot \int_{16\varepsilon}^{2T\varepsilon^2} \frac{e^{-x}}{x} dx \\
 & \leq \frac{e^{\varepsilon^2/2}}{\varepsilon^2} \cdot \int_{16\varepsilon}^1 \frac{e^{-x}}{x} dx \cdot \mathbf{1}[16\varepsilon < 1] + \frac{e^{\varepsilon^2/2}}{\varepsilon^2} \cdot \int_1^{2T\varepsilon^2} \frac{e^{-x}}{x} dx \\
 & \leq \frac{e^{\varepsilon^2/2}}{\varepsilon^2} \cdot \int_{16\varepsilon}^1 \frac{1}{16\varepsilon} dx \cdot \mathbf{1}[16\varepsilon < 1] + \frac{e^{\varepsilon^2/2}}{\varepsilon^2} \cdot \int_1^{2T\varepsilon^2} e^{-x} dx \\
 & \leq \frac{e^{\varepsilon^2/2}}{\varepsilon^2} \cdot \frac{1}{16\varepsilon} \cdot \mathbf{1}[16\varepsilon < 1] + \frac{e^{\varepsilon^2/2}}{\varepsilon^2} \cdot \frac{1}{e}, \\
 \text{(iii)} \quad & \int_{8/\varepsilon-1}^{T-1} \frac{1}{e^{\varepsilon^2 x/4} - 1} dx \leq \int_{8/\varepsilon-1}^{T-1} \frac{1}{\varepsilon^2 x/4} dx = \frac{4}{\varepsilon^2} \cdot \log x \Big|_{8/\varepsilon-1}^{T-1} \leq \frac{4}{\varepsilon^2} \cdot \log T.
 \end{aligned}$$

Let  $c = c_1 \cdot (2\sqrt{e} + 5) + 24$ , then (i) if  $\varepsilon < 1/16$ ,

$$\sum_{q=0}^{T-1} \mathbb{E} \left[ \frac{1 - p_{j, \tau_{j,q}+1}}{p_{j, \tau_{j,q}+1}} \right] \leq \frac{24}{\varepsilon^2} + c_1 \cdot \left( \frac{2\sqrt{e}}{\varepsilon^2} + \frac{2}{\sqrt{e} \cdot \varepsilon^3} + \frac{4}{\varepsilon^2} \cdot \log T \right) \leq c \cdot \left( \frac{1}{\varepsilon^3} + \frac{\log T}{\varepsilon^2} \right);$$

(ii) else, we have  $\varepsilon \geq 1/16$  and

$$\sum_{q=0}^{T-1} \mathbb{E} \left[ \frac{1 - p_{j, \tau_{j,q}+1}}{p_{j, \tau_{j,q}+1}} \right] \leq \frac{24}{\varepsilon^2} + c_1 \cdot \left( \frac{2\sqrt{e}}{\varepsilon^2} + \frac{1}{\sqrt{e} \cdot \varepsilon^2} + \frac{4}{\varepsilon^2} \cdot \log T \right) \leq \frac{c \log T}{\varepsilon^2}.$$

■

#### C.4 Proof of Theorem 3.4

**Proof** We assume that  $w(1) = w(2) = \dots = w(K)$ . The problem-independent bound for CTS can be derived with a similar analysis as the problem-dependent bound in Theorem 3.1. The key difference is to divide the items into two sets according to the gaps  $\{\Delta_{1,i}\}_{i=K+1}^L$  which were defined in the statement of Theorem 3.1.

For any  $\varepsilon \in (0, 1)$ , we let  $\tilde{S}_1 = \{i \in [L] : 0 < \Delta_{1,i} \leq \varepsilon\}$  and  $\tilde{S}_2 = \{i \in [L] : \Delta_{1,i} > \varepsilon\}$ . Then  $\tilde{S}_1 \cap \tilde{S}_2 = \emptyset$  and  $\tilde{S}_1 \cup \tilde{S}_2 = [L] \setminus S^*$ . According to the definition of  $N^B(i, j)$  in (3.4),

$$\sum_{i \in \tilde{S}_1} \Delta_{j,i} \cdot \mathbb{E}[N^B(i, j)] \leq T \quad \forall 1 \leq j \leq K < i \leq L.$$



Recall from the decomposition of the expected regret in (3.5) that we have

$$\begin{aligned} \text{Reg}(T) &\leq \sum_{j=1}^K \sum_{i=K+1}^L \Delta_{j,i} \cdot \mathbb{E}[N^{\text{B}}(i, j)] = \sum_{j=1}^K \sum_{i \in \tilde{S}_1} \Delta_{j,i} \cdot \mathbb{E}[N^{\text{B}}(i, j)] + \sum_{j=1}^K \sum_{i \in \tilde{S}_2} \Delta_{j,i} \cdot \mathbb{E}[N^{\text{B}}(i, j)] \\ &\leq \varepsilon K T + \sum_{j=1}^K \sum_{i \in \tilde{S}_2} \Delta_{j,i} \cdot \mathbb{E}[N^{\text{B}}(i, j)]. \end{aligned}$$

Now we focus on the second part of the regret. Similarly to before, we define two more events taking into account the concentration of Thompson sample  $\theta_t^{\text{B}}(i)$  (see Line 3 of Algorithm 1) and empirical mean  $\hat{w}_t(i)$  (see (3.6) for definition):

$$\mathcal{B}'_t(i) = \{\hat{w}_t(i) > w(i) + \varepsilon'_i\}, \quad \mathcal{C}'_t(i, j) = \{\theta_t^{\text{B}}(i) > w(j) - \varepsilon'_j\}, \quad j \in S^*, i \in \tilde{S}_2.$$

with

$$\varepsilon'_j = \frac{\varepsilon}{4} \quad \forall j \in S^*, \quad \varepsilon'_i = \frac{\Delta_{1,i}}{2} - \frac{\varepsilon}{4} \geq \frac{\varepsilon}{4} \quad \forall i \in \tilde{S}_2.$$

Then  $\varepsilon'_i + \varepsilon'_j = \Delta_{j,i}/2$ . Subsequently, we separate the second part of the regret similarly to (3.7) and bound the three terms separately as follows:

$$\begin{aligned} \text{(i)} \quad &\sum_{j=1}^K \sum_{i \in \tilde{S}_2} \Delta_{j,i} \cdot \mathbb{E} \left[ \sum_{t=1}^T \mathbf{1}[\mathcal{A}_t(i, j) \wedge \mathcal{B}'_t(i)] \right] \leq L - K + \sum_{i \in \tilde{S}_2} \frac{1}{(\varepsilon'_i)^2} \leq (L - K) \cdot \left(1 + \frac{16}{\varepsilon^2}\right), \\ \text{(ii)} \quad &\sum_{i \in \tilde{S}_2} \sum_{j=1}^K \Delta_{j,i} \cdot \mathbb{E} \left[ \sum_{t=1}^T \mathbf{1}[\mathcal{A}_t(i, j) \wedge \neg \mathcal{B}'_t(i) \wedge \mathcal{C}'_t(i, j)] \right] \leq \sum_{i \in \tilde{S}_2} \frac{16 \log T}{\Delta_{K,i}} + \sum_{i \in \tilde{S}_2} \sum_{j=1}^K \Delta_{j,i} \\ &\leq \frac{16(L - K) \log T}{\varepsilon} + K(L - K), \\ \text{(iii)} \quad &\sum_{i \in \tilde{S}_2} \sum_{j=1}^K \Delta_{j,i} \cdot \mathbb{E} \left[ \sum_{t=1}^T \mathbf{1}[\mathcal{A}_t(i, j) \wedge \neg \mathcal{C}'_t(i, j)] \right] \\ &\leq c \cdot \sum_{j=1}^K \Delta_{j,L} \cdot \left( \frac{\mathbf{1}[\varepsilon'_j < 1/16]}{(\varepsilon'_j)^3} + \frac{\log T}{(\varepsilon'_j)^2} \right) \leq 2\alpha_2 K \cdot \left( \frac{4 \cdot \mathbf{1}[\varepsilon < 1/4]}{\varepsilon^3} + \frac{\log T}{\varepsilon^2} \right), \end{aligned}$$

with  $\alpha_2 = 8c$  (a constant that does not depend on the problem instance, which appeared in (3.10) for the first time). Similarly to (3.8), the first bound above results from Lemma 3.2. The second one is derived similarly to (3.9). We apply Lemma 3.3 to derive the final bound similarly to (3.10). Overall, we can derive an upper bound on the total regret as a function of  $\varepsilon$  as follows:

$$\begin{aligned} \text{Reg}(T) &\leq \varepsilon K T + (L - K) \cdot \left(1 + \frac{16}{\varepsilon^2}\right) + \frac{16(L - K) \log T}{\varepsilon} + K(L - K) \\ &\quad + 2\alpha_2 K \cdot \left( \frac{4 \cdot \mathbf{1}[\varepsilon < 1/4]}{\varepsilon^3} + \frac{\log T}{\varepsilon^2} \right) \\ &= \varepsilon K T + \frac{16(L - K) \log T}{\varepsilon} + \frac{16(L - K) + 2\alpha_2 K \log T}{\varepsilon^2} + \frac{8\alpha_2 K}{\varepsilon^3} + (K + 1)(L - K). \end{aligned}$$

If we let  $\varepsilon = T^{-1/4}$ , we have

$$\text{Reg}(T) \leq (1 + 8\alpha_2) \cdot KT^{3/4} + 32(L - K)(\sqrt{T} + K + 1) + 2\alpha_2 K \sqrt{T} \log T.$$

This bound is of order  $O(KT^{3/4} + L\sqrt{T} + K\sqrt{T} \log T)$ . ■

### C.5 Proof of Lemma 4.2

**Proof Bounding probability of event  $\mathcal{E}_{\hat{w},t}$**  : We first consider a fixed non-negative integer  $N$  and a fixed item  $i$ . Let  $Y_1, \dots, Y_N$  be i.i.d. Bernoulli random variables, with the common mean  $w(i)$ . Denote  $\hat{\xi}_N = \sum_{i=1}^N Y_i/N$  as the sample mean, and  $\hat{V}_N = \hat{\xi}_N(1 - \hat{\xi}_N)$  as the empirical variance. By applying Theorem B.1 with  $\delta = 1/(t+1)^4$ , we have

$$\Pr \left( \left| \hat{\xi}_N - w(i) \right| > \sqrt{\frac{8\hat{V}_N \log(t+1)}{N}} + \frac{12 \log(t+1)}{N} \right) \leq \frac{3}{(t+1)^4}. \quad (\text{C.4})$$

By an abuse of notation, let  $\hat{\xi}_N = 0$  if  $N = 0$ . Inequality (C.4) implies the following concentration bound when  $N$  is non-negative:

$$\Pr \left( \left| \hat{\xi}_N - w(i) \right| > \sqrt{\frac{16\hat{V}_N \log(t+1)}{N+1}} + \frac{24 \log(t+1)}{N+1} \right) \leq \frac{3}{(t+1)^4}. \quad (\text{C.5})$$

Subsequently, we can establish the concentration property of  $\hat{w}_t(i)$  by a union bound of  $N_t(i)$  over  $\{0, 1, \dots, t-1\}$ :

$$\begin{aligned} & \Pr \left( \left| \hat{w}_t(i) - w(i) \right| > \sqrt{\frac{16\hat{v}_t(i) \log(t+1)}{N_t(i)+1}} + \frac{24 \log(t+1)}{N_t(i)+1} \right) \\ & \leq \Pr \left( \left| \hat{\xi}_N - w(i) \right| > \sqrt{\frac{16\hat{V}_N \log(t+1)}{N+1}} + \frac{24 \log(t+1)}{N+1} \text{ for some } N \in \{0, 1, \dots, t-1\} \right) \\ & \leq \frac{3}{(t+1)^3}. \end{aligned}$$

Finally, taking union bound over all items  $i \in L$ , we know that event  $\mathcal{E}_{\hat{w},t}$  holds true with probability at least  $1 - 3L/(t+1)^3$ .

**Bounding probability of event  $\mathcal{E}_{\theta,t}$ , conditioned on event  $\mathcal{E}_{\hat{w},t}$** : Consider an observation trajectory  $H_t$  satisfying event  $\mathcal{E}_{\hat{w},t}$ . By the definition of the Thompson sample  $\theta_t(i)$  (see Line 7 in Algorithm 2), we have

$$\begin{aligned} & \Pr(|\theta_t(i) - \hat{w}_t(i)| > h_t(i) \text{ for all } i \in L | H_{\hat{w},t}) \\ & = \Pr \left( \left| Z_t \cdot \max \left\{ \sqrt{\frac{\hat{v}_t(i) \log(t+1)}{N_t(i)+1}}, \frac{\log(t+1)}{N_t(i)+1} \right\} \right| > \right. \\ & \quad \left. \sqrt{\log(t+1)} \left[ \sqrt{\frac{16\hat{v}_t(i) \log(t+1)}{N_t(i)+1}} + \frac{24 \log(t+1)}{N_t(i)+1} \right] \text{ for all } i \in [L] \mid \hat{w}_t(i), N_t(i) \right) \end{aligned}$$

$$\begin{aligned}
 &\leq \Pr \left( \left| Z_t \cdot \max \left\{ \sqrt{\frac{\hat{\nu}_t(i) \log(t+1)}{N_t(i)+1}}, \frac{\log(t+1)}{N_t(i)+1} \right\} \right| > \right. \\
 &\quad \left. \sqrt{16 \log(t+1)} \max \left\{ \sqrt{\frac{\hat{\nu}_t(i) \log(t+1)}{N_t(i)+1}}, \frac{\log(t+1)}{N_t(i)+1} \right\} \text{ for all } i \in [L] \mid \hat{w}_t(i), N_t(i) \right) \\
 &\leq \exp[-8 \log(t+1)] \leq \frac{1}{(t+1)^3}. \tag{C.6}
 \end{aligned}$$

The inequality in (C.6) is by the concentration property of a Gaussian random variable, see Theorem B.2. Altogether, the lemma is proved.  $\blacksquare$

### C.6 Proof of Lemma 4.3

**Proof** To start, we denote the shorthand  $\theta_t(i)^+ = \max\{\theta_t(i), 0\}$ . We demonstrate that, if events  $\mathcal{E}_{\hat{w}_t, t}$ ,  $\mathcal{E}_{\theta_t, t}$  and inequality (4.2) hold, then for all  $\bar{S} = (\bar{i}_1, \dots, \bar{i}_K) \in \bar{\mathcal{S}}_t$  we have:

$$\sum_{k=1}^K \left[ \prod_{j=1}^{k-1} (1 - w(\bar{i}_j)) \right] \cdot \theta_t(\bar{i}_k)^+ \stackrel{(\ddagger)}{<} \sum_{k=1}^K \left[ \prod_{j=1}^{k-1} (1 - w(j)) \right] \cdot \theta_t(k)^+ \stackrel{(\dagger)}{\leq} \sum_{k=1}^K \left[ \prod_{j=1}^{k-1} (1 - w(i_j^t)) \right] \cdot \theta_t(i_k^t)^+, \tag{C.7}$$

where we recall that  $S_t = (i_1^t, \dots, i_K^t)$  in an optimal arm for  $\theta_t$ , and  $\theta_t(i_1^t) \geq \theta_t(i_2^t) \geq \dots \geq \theta_t(i_K^t) \geq \max_{j \in [L] \setminus \{i_1^t, \dots, i_K^t\}} \theta_t(i)$ . The inequalities in (C.7) clearly implies that  $S_t \in \mathcal{S}_t$ . To justify these inequalities, we proceed as follows:

**Showing**  $(\dagger)$ : This inequality is true even without requiring events  $\mathcal{E}_{\hat{w}_t, t}$ ,  $\mathcal{E}_{\theta_t, t}$  and inequality (4.2) to be true. Indeed, we argue the following:

$$\sum_{k=1}^K \left[ \prod_{j=1}^{k-1} (1 - w(j)) \right] \cdot \theta_t(k)^+ \leq \sum_{k=1}^K \left[ \prod_{j=1}^{k-1} (1 - w(j)) \right] \cdot \theta_t(i_k^t)^+ \tag{C.8}$$

$$\leq \sum_{k=1}^K \left[ \prod_{j=1}^{k-1} (1 - w(i_j^t)) \right] \cdot \theta_t(i_k^t)^+. \tag{C.9}$$

To justify inequality (C.8), consider function  $f : \pi_K(L) \rightarrow \mathbb{R}$  defined as

$$f((i_k)_{k=1}^K) := \sum_{k=1}^K \left[ \prod_{j=1}^{k-1} (1 - w(j)) \right] \cdot \theta_t(i_k)^+.$$

We assert that  $S_t \in \operatorname{argmax}_{S \in \pi_K(L)} f(S)$ . The assertion can be justified by the following two properties. First, by the choice of  $S_t$ , we know that  $\theta_t(i_1^t)^+ \geq \theta_t(i_2^t)^+ \geq \dots \geq \theta_t(i_K^t)^+ \geq \max_{i \in [L] \setminus S_t} \theta_t(i)^+$ . Second, the linear coefficients in the function  $f$  are monotonic and non-negative, in the sense that  $1 \geq 1 - w(1) \geq (1 - w(1))(1 - w(2)) \geq \dots \geq \prod_{k=1}^{K-1} (1 - w(k)) \geq 0$ . Altogether, we have  $f(S_t) \geq f(S^*)$ , hence inequality (C.8) is shown.

Next, inequality (C.9) clearly holds, since for each  $k \in [K]$  we know that  $\theta_t(i_k^t)^+ \geq 0$ , and  $\prod_{j=1}^{k-1} (1 - w(j)) \leq \prod_{j=1}^{k-1} (1 - w(i_j^t))$ . The latter is due to the fact that  $1 \geq w(1) \geq w(2) \geq \dots \geq w(K) \geq \max_{i \in [L] \setminus [K]} w(i)$ . Altogether, inequality  $(\dagger)$  is established.

**Showing** (‡): The demonstration crucially hinges on events  $\mathcal{E}_{\hat{w},t}, \mathcal{E}_{\theta,t}$  and inequality (4.2) being held true. For any  $\bar{S} = (\bar{i}_1, \dots, \bar{i}_K) \in \bar{\mathcal{S}}_t$ , we have

$$\begin{aligned} & \sum_{k=1}^K \left[ \prod_{j=1}^{k-1} (1 - w(\bar{i}_j)) \right] \theta_t(\bar{i}_k)^+ \\ & \leq \sum_{k=1}^K \left[ \prod_{j=1}^{k-1} (1 - w(\bar{i}_j)) \right] (w(\bar{i}_k) + g_t(\bar{i}_k) + h_t(\bar{i}_k)) \end{aligned} \quad (\text{C.10})$$

$$< \left\{ \sum_{k=1}^K \left[ \prod_{j=1}^{k-1} (1 - w(\bar{i}_j)) \right] w(\bar{i}_k) \right\} + r(S^*|\mathbf{w}) - r(\bar{S}|\mathbf{w}) \quad (\text{C.11})$$

$$\begin{aligned} & = r(S^*|\mathbf{w}) = \sum_{k=1}^K \left[ \prod_{j=1}^{k-1} (1 - w(j)) \right] w(k) \\ & \leq \sum_{k=1}^K \left[ \prod_{j=1}^{k-1} (1 - w(j)) \right] \theta_t(k) \leq \sum_{k=1}^K \left[ \prod_{j=1}^{k-1} (1 - w(j)) \right] \theta_t(k)^+. \end{aligned} \quad (\text{C.12})$$

Inequality (C.10) is by the assumption that events  $\mathcal{E}_{\hat{w},t}, \mathcal{E}_{\theta,t}$  are true, which means that for all  $i \in [L]$  we have  $\theta_t(i)^+ \leq \mu(i) + g_t(i) + h_t(i)$ . Inequality (C.11) is by the fact that  $S \in \bar{\mathcal{S}}_t$ . Inequality (C.12) is by our assumption that inequality (4.2) holds.

Altogether, the inequalities (†, ‡) in (C.7) are shown, and the Lemma is established. ■

## C.7 Proof of Lemma 4.5

The proof of Lemma 4.5 crucially uses the following lemma on the expression of the difference in expected reward between two arms:

**Lemma C.4 (Implied by Zong et al. (2016), Lemma 1)** *Let  $S = (i_1, \dots, i_K)$ ,  $S' = (i'_1, \dots, i'_K)$  be two arbitrary ordered  $K$ -subsets of  $[L]$ . For any  $\mathbf{w}, \mathbf{w}' \in \mathbb{R}^L$ , the following equalities holds:*

$$\begin{aligned} r(S|\mathbf{w}) - r(S'|\mathbf{w}') & = \sum_{k=1}^K \left[ \prod_{j=1}^{k-1} (1 - w(i_j)) \right] \cdot (w(i_k) - w'(i'_k)) \cdot \left[ \prod_{j=k+1}^K (1 - w'(i'_j)) \right] \\ & = \sum_{k=1}^K \left[ \prod_{j=1}^{k-1} (1 - w'(i'_j)) \right] \cdot (w(i_k) - w'(i'_k)) \cdot \left[ \prod_{j=k+1}^K (1 - w(i_j)) \right]. \end{aligned}$$

While Lemma C.4 is folklore in the cascading bandit literature, we provide a proof in Appendix C.11 for the sake of completeness. Now, we proceed to the proof of Lemma 4.5: **Proof** In the proof, we always condition to the historical observation  $H_t$  stated in the Lemma. To proceed with the analysis, we define  $\tilde{S}_t = (\tilde{i}'_1, \dots, \tilde{i}'_K) \in \mathcal{S}_t$  as an ordered

$K$ -subset that satisfies the following minimization criterion:

$$\sum_{k=1}^K \left[ \prod_{j=1}^{k-1} (1 - w(\tilde{i}_j)) \right] (g_t(\tilde{i}_k) + h_t(\tilde{i}_k)) = \min_{S=(i_1, \dots, i_K) \in \mathcal{S}_t} \sum_{k=1}^K \left[ \prod_{j=1}^{k-1} (1 - w(i_j)) \right] (g_t(i_k) + h_t(i_k)). \quad (\text{C.13})$$

We emphasize that both  $\tilde{S}_t$  and the left hand side of (C.13) are deterministic in the current discussion, where we condition on  $H_t$ . To establish tight bounds on the regret, we consider the truncated version,  $\tilde{\theta}_t \in [0, 1]^L$ , of the Thompson sample  $\theta_t$ . For each  $i \in L$ , define

$$\tilde{\theta}_t(i) = \min\{1, \max\{0, \theta_t(i)\}\}.$$

The truncated version  $\tilde{\theta}_t(i)$  serves as a correction of  $\theta_t(i)$ , in the sense that the Thompson sample  $\theta_t(i)$ , which serves as a Bayesian estimate of click probability  $w(i)$ , should lie in  $[0, 1]$ . It is important to observe the following two properties hold under the truncated Thompson sample  $\tilde{\theta}_t$ :

**Property 1** Our pulled arm  $S_t$  is still optimal under the truncated estimate  $\tilde{\theta}_t$ , i.e.

$$S_t \in \operatorname{argmax}_{S \in \pi_K(L)} r(S | \tilde{\theta}_t).$$

Indeed, the truncated Thompson sample can be sorted in a descending order in the same way as for the original Thompson sample<sup>3</sup>, i.e.  $\theta_t(i_1^t) \geq \theta_t(i_2^t) \geq \dots \geq \theta_t(i_K^t) \geq \max_{i \in [L] \setminus \{i_1^t, \dots, i_K^t\}} \theta_t(i)$ . The optimality of  $S_t$  thus follows.

**Property 2** For any  $t, i$ , if it holds that  $|\theta_t(i) - w(i)| \leq g_t(i) + h_t(i)$ , then it also holds that  $|\tilde{\theta}_t(i) - w(i)| \leq g_t(i) + h_t(i)$ . Indeed, we know that  $|\tilde{\theta}_t(i) - w(i)| \leq |\theta_t(i) - w(i)|$ .

Now, we use the ordered  $K$ -subset  $\tilde{S}_t$  and the truncated Thompson sample  $\tilde{\theta}_t$  to decompose the conditionally expected round  $t$  regret as follows:

$$\begin{aligned} r(S^* | \mathbf{w}) - r(S_t | \mathbf{w}) &= \left[ r(S^* | \mathbf{w}) - r(\tilde{S}_t | \mathbf{w}) \right] + \left[ r(\tilde{S}_t | \mathbf{w}) - r(S_t | \mathbf{w}) \right] \\ &\leq \left[ r(S^* | \mathbf{w}) - r(\tilde{S}_t | \mathbf{w}) \right] + \left[ r(\tilde{S}_t | \mathbf{w}) - r(S_t | \mathbf{w}) \right] \mathbf{1}(\mathcal{E}_{\theta,t}) + \mathbf{1}(-\mathcal{E}_{\theta,t}) \\ &\leq \underbrace{\left[ r(S^* | \mathbf{w}) - r(\tilde{S}_t | \mathbf{w}) \right]}_{(\diamond)} + \underbrace{\left[ r(\tilde{S}_t | \tilde{\theta}_t) - r(S_t | \tilde{\theta}_t) \right]}_{(\clubsuit)} \\ &\quad + \underbrace{\left[ r(S_t | \tilde{\theta}_t) - r(S_t | \mathbf{w}) \right]}_{(\heartsuit)} \mathbf{1}(\mathcal{E}_{\theta,t}) + \underbrace{\left[ r(\tilde{S}_t | \mathbf{w}) - r(\tilde{S}_t | \tilde{\theta}_t) \right]}_{(\spadesuit)} \mathbf{1}(\mathcal{E}_{\theta,t}) + \mathbf{1}(-\mathcal{E}_{\theta,t}). \end{aligned} \quad (\text{C.14})$$

We bound  $(\diamond, \clubsuit, \heartsuit, \spadesuit)$  from above as follows:

**Bounding  $(\diamond)$ :** By the assumption that  $\tilde{S}_t = (\tilde{i}_1^t, \dots, \tilde{i}_K^t) \in \mathcal{S}_t$ , with certainty we have

$$(\diamond) \leq \sum_{k=1}^K \left[ \prod_{j=1}^{k-1} (1 - w(\tilde{i}_j^t)) \right] (g_t(\tilde{i}_k^t) + h_t(\tilde{i}_k^t)). \quad (\text{C.15})$$

3. Recall that  $\theta_t(i_1^t) \geq \theta_t(i_2^t) \geq \dots \geq \theta_t(i_K^t) \geq \max_{i \in [L] \setminus \{i_k^t\}_{k=1}^K} \theta_t(i)$  for the original Thompson sample  $\theta_t$ .

**Bounding (♣):** By **Property 1** of the truncated Thompson sample  $\tilde{\theta}_t$ , we know that  $r(S_t|\tilde{\theta}_t) = \max_{S \in \pi_K(L)} r(S|\tilde{\theta}_t) \geq r(\tilde{S}_t|\tilde{\theta}_t)$ . Therefore, with certainty we have

$$(\clubsuit) \leq 0. \quad (\text{C.16})$$

**Bounding (♡):** . We bound the term as follows:

$$\begin{aligned} & \mathbf{1}(\mathcal{E}_{\theta,t}) \left[ r(S_t|\tilde{\theta}_t) - r(S_t|\mathbf{w}) \right] \\ &= \mathbf{1}(\mathcal{E}_{\theta,t}) \sum_{k=1}^K \left[ \prod_{j=1}^{k-1} (1 - w(i_j^t)) \right] \cdot (\tilde{\theta}_t(i_k^t) - w(i_k^t)) \cdot \left[ \prod_{j=k+1}^K (1 - \tilde{\theta}_t(i_j^t)) \right] \end{aligned} \quad (\text{C.17})$$

$$\begin{aligned} &\leq \mathbf{1}(\mathcal{E}_{\theta,t}) \sum_{k=1}^K \left[ \prod_{j=1}^{k-1} (1 - w(i_j^t)) \right] \cdot \left| \tilde{\theta}_t(i_k^t) - w(i_k^t) \right| \cdot \left[ \prod_{j=k+1}^K \left| 1 - \tilde{\theta}_t(i_j^t) \right| \right] \\ &\leq \mathbf{1}(\mathcal{E}_{\theta,t}) \sum_{k=1}^K \left[ \prod_{j=1}^{k-1} (1 - w(i_j^t)) \right] \cdot [g_t(i_k^t) + h_t(i_k^t)] \end{aligned} \quad (\text{C.18})$$

$$\leq \sum_{k=1}^K \left[ \prod_{j=1}^{k-1} (1 - w(i_j^t)) \right] \cdot [g_t(i_k^t) + h_t(i_k^t)]. \quad (\text{C.19})$$

Equality (C.17) is by applying the second equality in Lemma C.4, with  $S = S' = S_t$ , as well as  $\mathbf{w}' \leftarrow \mathbf{w}$ ,  $\mathbf{w} \leftarrow \theta_t$ . Inequality (C.18) is by the following two facts: (1) By the definition of the truncated Thompson sample  $\tilde{\theta}$ , we know that  $|1 - \tilde{\theta}_t(i)| \leq 1$  for all  $i \in [L]$ ; (2) By assuming event  $\mathcal{E}_{\theta,t}$  and conditioning on  $H_t$  where event  $\mathcal{E}_{\hat{w},t}$  holds true, **Property 2** implies that that  $|\tilde{\theta}_t(i) - w(i)| \leq g_t(i) + h_t(i)$  for all  $i$ .

**Bounding (♠):** The analysis is similar to the analysis on (♡):

$$\begin{aligned} & \mathbf{1}(\mathcal{E}_{\theta,t}) \left[ r(\tilde{S}_t|\mathbf{w}) - r(\tilde{S}_t|\tilde{\theta}_t) \right] \\ &= \mathbf{1}(\mathcal{E}_{\theta,t}) \sum_{k=1}^K \left[ \prod_{j=1}^{k-1} (1 - w(\tilde{i}_j^t)) \right] \cdot (w(\tilde{i}_k^t) - \tilde{\theta}_t(\tilde{i}_k^t)) \cdot \left[ \prod_{j=k+1}^K (1 - \tilde{\theta}_t(\tilde{i}_j^t)) \right] \end{aligned} \quad (\text{C.20})$$

$$\leq \mathbf{1}(\mathcal{E}_{\theta,t}) \sum_{k=1}^K \left[ \prod_{j=1}^{k-1} (1 - w(\tilde{i}_j^t)) \right] \cdot [g_t(\tilde{i}_k^t) + h_t(\tilde{i}_k^t)] \quad (\text{C.21})$$

$$\leq \sum_{k=1}^K \left[ \prod_{j=1}^{k-1} (1 - w(\tilde{i}_j^t)) \right] \cdot [g_t(\tilde{i}_k^t) + h_t(\tilde{i}_k^t)]. \quad (\text{C.22})$$

Equality (C.20) is by applying the first equality in Lemma C.4, with  $S = S' = \tilde{S}_t$ , and  $\mathbf{w} \leftarrow \mathbf{w}$ ,  $\mathbf{w}' \leftarrow \theta_t$ . Inequality (C.21) follows the same logic as inequality (C.18).

Altogether, collating the bounds (C.15, C.16, C.19, C.22) for  $(\diamond, \clubsuit, \heartsuit, \spadesuit)$  respectively, we bound (C.14) from above (conditioned on  $H_t$ ) as follows:

$$\begin{aligned} r(S^*|\mathbf{w}) - r(S_t|\mathbf{w}) &\leq 2 \sum_{k=1}^K \left[ \prod_{j=1}^{k-1} (1 - w(\tilde{i}_j^t)) \right] (g_t(\tilde{i}_j^t) + h_t(\tilde{i}_j^t)) \\ &\quad + \sum_{k=1}^K \left[ \prod_{j=1}^{k-1} (1 - w(i_j^t)) \right] \cdot [g_t(i_k^t) + h_t(i_k^t)] + 1(\neg\mathcal{E}_{\theta,t}). \end{aligned} \quad (\text{C.23})$$

Now, observe that

$$\begin{aligned} &\mathbb{E}_{\theta_t} \left[ \sum_{k=1}^K \left[ \prod_{j=1}^{k-1} (1 - w(i_j^t)) \right] (g_t(i_j^t) + h_t(i_j^t)) \mid H_t \right] \\ &\geq \mathbb{E}_{\theta_t} \left[ \sum_{k=1}^K \left[ \prod_{j=1}^{k-1} (1 - w(i_j^t)) \right] (g_t(i_j^t) + h_t(i_j^t)) \mid H_t, S_t \in \mathcal{S}_t \right] \Pr_{\theta_t} \left[ S_t \in \mathcal{S}_t \mid H_t \right] \\ &\geq \left\{ \sum_{k=1}^K \left[ \prod_{j=1}^{k-1} (1 - w(\tilde{i}_j)) \right] (g_t(\tilde{i}_j) + h_t(\tilde{i}_j)) \right\} \cdot \left( c - \frac{1}{(t+1)^3} \right). \end{aligned} \quad (\text{C.24})$$

Thus, taking conditional expectation  $\mathbb{E}_{\theta_t}[\cdot|H_t]$  on both sides in inequality (C.23) gives

$$\begin{aligned} &\mathbb{E}_{\theta_t}[R(S^*|\mathbf{w}) - R(S_t|\mathbf{w})|H_t] \\ &\leq \left( 1 + \frac{2}{c - \frac{1}{(t+1)^3}} \right) \mathbb{E}_{\theta_t} \left[ \sum_{k=1}^K \left[ \prod_{j=1}^{k-1} (1 - w(i_j^t)) \right] \cdot [g_t(i_k^t) + h_t(i_k^t)] \mid H_t \right] + \mathbb{E}_{\theta_t}[1(\neg\mathcal{E}_{\theta,t})|H_t]. \end{aligned}$$

Finally, the lemma is proved by the assumption that  $c > 2/(t+1)^3$ , and noting from Lemma 4.2 that  $\mathbb{E}_{\theta_t}[1(\neg\mathcal{E}_{\theta,t})|H_t] \leq 1/(t+1)^3$ .  $\blacksquare$

### C.8 Proof of Lemma 5.3

**Proof Bounding probability of event  $\mathcal{E}_{\hat{\psi},t}$ :** We use Theorem B.3 with  $\eta_t(i) = W_t(i) - x(i)^\top w$ ,  $\mathcal{F}'_t = \{S_q\}_{q=1}^{t-1} \cup \{(i_k^q, W_q(i_k^q))_{k=1}^{\min\{k_t, \infty\}}\}_{q=1}^{t-1} \cup S_t \cup (x(i))_{i=1}^K$ . (Note that effectively,  $\mathcal{F}'_t$  has all the information, including the items observed, until time  $t+1$ , except for the clickness of the selected items at time  $t+1$ ). By the definition of  $\mathcal{F}'_t$ ,  $x(i)$  is  $(\mathcal{F}'_{t-1})$ -measurable, and  $\eta_t$  is  $\mathcal{F}'_t$ -measurable. Also,  $\eta_t$  is conditionally  $R$ -sub-Gaussian due to the the problem settings (refer to Remark 5.2), and is a martingale difference process:

$$\mathbb{E}[\eta_t(i)|\mathcal{F}'_{t-1}] = \mathbb{E}[W_t(i)|x(i)] - x(i)^\top w = 0.$$

Also, let

$$Q_t = I_d + \sum_{q=1}^t \sum_{k=1}^{k_q} x(i_k^q) x(i_k^q)^\top,$$

$$\zeta_t = \sum_{q=1}^t \sum_{k=1}^{k_q} x(i_k^q) \eta_q(i_k^q) = \sum_{q=1}^t \sum_{k=1}^{k_q} x(i_k^q) [W_q(i_k^q) - x(i_k^q)^\top w],$$

Note that  $M_t = Q_{t-1}$ , recall  $b_{t+1} = \sum_{q=1}^t \sum_{k=1}^{k_q} x(i_k^q) W(i_k^q)$  and  $\hat{\psi}_t = M_t^{-1} b_t$ ,

$$\begin{aligned} Q_{t-1}(\hat{\psi}_t - w) &= Q_{t-1}(M_t^{-1} b_t - w) = b_t - Q_{t-1} w \\ &= \sum_{q=1}^{t-1} \sum_{k=1}^{k_q} x(i_k^q) W(i_k^q) - \left[ I_d + \sum_{q=1}^{t-1} \sum_{k=1}^{k_q} x(i_k^q) x(i_k^q)^\top \right] w \\ &= \sum_{q=1}^{t-1} \sum_{k=1}^{k_q} x(i_k^q) [W_q(i_k^q) - x(i_k^q)^\top w] - w = \zeta_{t-1} - w \end{aligned}$$

yields that  $\hat{\psi}_t - w = Q_{t-1}^{-1}(\zeta_{t-1} - w)$ . Let for any vector  $z \in \mathbb{R}^{d \times 1}$  and matrix  $A \in \mathbb{R}^{d \times d}$ ,  $\|z\|_A := \sqrt{y^\top A y}$ . Then, for all  $i$ ,

$$\begin{aligned} \left| x(i)^\top \hat{\psi}_t - x(i)^\top w \right| &= \left| x(i)^\top Q_{t-1}^{-1}(\zeta_{t-1} - w) \right| \\ &\leq \|x(i)\|_{Q_{t-1}^{-1}} \|\zeta_{t-1} - w\|_{Q_{t-1}^{-1}} = \|x(i)\|_{M_t^{-1}} \|\zeta_{t-1} - w\|_{Q_{t-1}^{-1}}, \end{aligned}$$

where the inequality holds because  $Q_{t-1}^{-1}$  is a positive definite matrix. Using Theorem B.3, for any  $\delta' > 0$ ,  $t \geq 1$ , with probability at least  $1 - \delta'$ ,

$$\|\zeta_{t-1}\|_{Q_{t-1}^{-1}} \leq R \sqrt{d \log \left( \frac{t}{\delta'} \right)}.$$

Therefore,  $\|\zeta_{t-1} - w\|_{Q_{t-1}^{-1}} \leq \|\zeta_{t-1}\|_{Q_{t-1}^{-1}} + \|w\|_{Q_{t-1}^{-1}} \leq R \sqrt{d \log \left( \frac{t}{\delta'} \right)} + \|w\| \leq R \sqrt{d \log \left( \frac{t}{\delta'} \right)} + 1$ . Substituting  $\delta' = 1/t^2$ , we get that with probability  $1 - 1/t^2$ , for all  $i$ ,

$$\begin{aligned} \left| x(i)^\top \hat{\psi}_t - x(i)^\top w \right| &\leq \|x(i)\|_{M_t^{-1}} \cdot \left( R \sqrt{d \log \left( \frac{t}{\delta'} \right)} + 1 \right) \\ &= \|x(i)\|_{M_t^{-1}} \cdot \left( R \sqrt{3d \log t} + 1 \right) = l_t u_t(i) = g_t(i). \end{aligned}$$

This proves the bound on the probability of  $\mathcal{E}_{\hat{\psi}, t}$ .

**Bounding probability of event  $\mathcal{E}_{\rho, t}$ :** Given any filtration  $H_t$ ,  $x(i)$ ,  $M_t$  are fixed. since  $\xi_t = (\xi_t^1, \xi_t^2, \dots, \xi_t^d)^\top$  is a standard multivariate normal vector (mean  $0_{d \times 1}$  and covariance  $I_d$ ), each  $\xi_t^l$  ( $1 \leq l \leq d$ ) is a standard univariate normal random variable (mean 0 and



variance 1). Theorem B.2 implies

$$\Pr(|\xi_t^l| > \sqrt{4 \log t}) \leq \exp(-2 \log t) = \frac{1}{t^2} \quad \forall 1 \leq l \leq d,$$

$$\text{and hence } \Pr(\|\xi_t\| > \sqrt{4d \log t}) \leq \sum_{l=1}^d \Pr(|\xi_t^l| > \sqrt{4 \log t}) \leq \frac{d}{t^2}.$$

Further,

$$\begin{aligned} \left| x(i)^\top \rho_t - x(i)^\top \hat{\psi}_t \right| &= |x(i)^\top [\rho_t - \hat{\psi}_t]| = |x(i)^\top M_t^{-1/2} M_t^{-1/2} [\rho_t - \hat{\psi}_t]| \\ &\leq \lambda v_t \sqrt{K} \sqrt{x(i)^\top M_t^{-1} x(i)} \cdot \left\| \left( \frac{1}{\lambda v_t \sqrt{K}} M_t^{1/2} [\rho_t - \hat{\psi}_t] \right) \right\| \\ &= \lambda v_t \sqrt{K} u_t(i) \|\xi_t\| \leq \lambda v_t \sqrt{K} u_t(i) \sqrt{4d \log t} \end{aligned}$$

with probability at least  $1 - d/(t^2)$ .

Alternatively, we can bound  $|x(i)^\top \rho_t - x(i)^\top \hat{\psi}_t|$  for every  $i$  by considering that  $x(i)^\top \rho_t$  is a Gaussian random variable with mean  $x(i)^\top \hat{\psi}(t)$  and variance  $\lambda^2 v_t^2 K u_t(i)^2$ . Therefore, again using Theorem B.2, for every  $i$ ,

$$\Pr \left( \left| x(i)^\top \rho_t - x(i)^\top \hat{\psi}_t \right| > \lambda v_t \sqrt{K} u_t(i) \sqrt{4 \log(tL)} \right) \leq \exp \left[ -\frac{4 \log(tL)}{2} \right] = \frac{1}{t^2 L^2}.$$

Taking union bound over  $i = 1, \dots, L$ , we obtain that

$$\left| x(i)^\top \rho_t - x(i)^\top \hat{\psi}_t \right| > \lambda v_t \sqrt{K} u_t(i) \sqrt{\log(tL)}$$

holds for all  $i \in [L]$  with probability at least  $1 - 1/t^2$ .

Combined, the two bounds give that  $\mathcal{E}_{\rho,t}$  holds with probability at least  $1 - (d+1)/(t^2)$ .

Altogether, the lemma is proved.  $\blacksquare$

### C.9 Proof of Lemma 5.5

**Proof** Recall  $\alpha(1) := 1$ , and  $\alpha(k) = \prod_{j=1}^{k-1} (1 - w(j)) = \prod_{j=1}^{k-1} (1 - x(j)^\top \beta)$  for  $2 \leq k \leq K$ . By the second part of Lemma 5.3, we know that  $\Pr[\mathcal{E}_{\rho,t} | H_t] \geq 1 - (d+1)/(t^2)$ , so to complete this proof, it suffices to show that  $\Pr[(5.3) \text{ holds} | H_t] \geq c$ . For this purpose, consider

$$\begin{aligned} &\Pr_{\rho_t} \left[ \sum_{k=1}^K \alpha(k) x(k)^\top \rho_t \geq \sum_{k=1}^K \alpha(k) x(k)^\top \beta \mid H_t \right] \\ &= \Pr_{\xi_t} \left[ \sum_{k=1}^K \alpha(k) x(k)^\top (\hat{\psi}_t + \lambda v_t \sqrt{K} M_t^{-1/2} \xi_t) \geq \sum_{k=1}^K \alpha(k) x(k)^\top \beta \mid H_t \right] \end{aligned} \quad (\text{C.25})$$

$$\geq \Pr_{\xi_t} \left[ \sum_{k=1}^K \alpha(k) \left[ x(k)^\top \beta - g_t(k) \right] + \lambda v_t \sqrt{K} \sum_{k=1}^K \alpha(k) x(k)^\top M_t^{-1/2} \xi_t \geq \sum_{k=1}^K \alpha(k) x(k)^\top \beta \mid H_t \right] \quad (\text{C.26})$$

$$\begin{aligned}
 &= \Pr_{\xi_t} \left[ \lambda v_t \sum_{k=1}^K \alpha(k) x(k)^\top M_t^{-1/2} \xi_t \geq \sum_{k=1}^K \alpha(k) g_t(k) \mid H_t \right] \\
 &\geq \frac{1}{4\sqrt{\pi}} \exp \left\{ -\frac{7}{2} \left[ \sum_{k=1}^K \alpha(k) g_t(k) / \hat{\sigma}_t \right]^2 \right\} \tag{C.27}
 \end{aligned}$$

$$\begin{aligned}
 &= \frac{1}{4\sqrt{\pi}} \exp \left\{ -\frac{7}{2} \left[ \sum_{k=1}^K \alpha(k) g_t(k) \right]^2 / \lambda^2 v_t^2 \sum_{k=1}^K \alpha(k)^2 u_t(k)^2 \right\} \\
 &= \frac{1}{4\sqrt{\pi}} \exp \left\{ -\frac{7l_t^2}{2\lambda^2 v_t^2 K} \left[ \sum_{k=1}^K \alpha(k) u_t(k) \right]^2 / \sum_{k=1}^K \alpha(k)^2 u_t(k)^2 \right\} \tag{C.28}
 \end{aligned}$$

$$\geq \frac{1}{4\sqrt{\pi}} \exp \left( -\frac{7}{2\lambda^2} \right) := c(\lambda). \tag{C.29}$$

Step (C.25) is by the definition of  $\rho_t$  in Line 5 in Algorithm 3. Step (C.26) is by the Lemma assumption that  $H_t \in \mathcal{H}_{\hat{\psi}, t}$ , which indicates that  $x(k)^\top \hat{\psi}_t \geq x(k)^\top \beta - g_t(k)$  for all  $k \in [K]$ . Noted that  $\lambda v_t \sum_{k=1}^K \alpha(k) x(k)^\top M_t^{-1} \xi_t$  is a Normal random variable with mean 0 and variance

$$\hat{\sigma}_t^2 := \lambda^2 v_t^2 K \sum_{k=1}^K \alpha(k)^2 x(k)^\top M_t^{-1} x(k) = \lambda^2 v_t^2 K \sum_{k=1}^K \alpha(k)^2 u_t(k)^2,$$

step (C.27) is an application of the anti-concentration inequality of a normal random variable in Theorem B.2. Step (C.28) is by applying the definition of  $g_t(i)$ . Step (C.29) follows from  $t \geq 4$  which implying  $l_t \leq v_t$ , and Cauchy-Schwartz inequality.  $\blacksquare$

### C.10 Derivation of Inequality (5.7) in Section 5

To derive the inequality (5.7)

$$\sum_{t=1}^T \sqrt{y_t^\top M_t^{-1} y_t} \leq 5\sqrt{dT \log T},$$

we use the following result:

**Lemma C.5** [Implied by Auer et al. (2002), Lemma 11; Agrawal and Goyal (2017a), Lemma 9] *Let  $A' = A + yy^\top$ , where  $y \in \mathbb{R}^d$ ,  $A, A' \in \mathbb{R}^{d \times d}$ , and all the eigenvalues  $\lambda_j, j = 1, \dots, d$  of  $A$  are greater than or equal to 1. Then, the eigenvalues  $\lambda'_j, j = 1, \dots, d$  of  $A'$  can be arranged so that  $\lambda_j \leq \lambda'_j$  for all  $j$ , and*

$$y^\top A^{-1} y \leq 10 \sum_{j=1}^d \frac{\lambda'_j - \lambda_j}{\lambda_j}.$$

Let  $\lambda_{j,t}$  denote the eigenvalues of  $M_t$ . Note that  $M_{t+1} = M_t + y_t y_t^\top$ , and  $\lambda_{j,t} \geq 1, \forall j$ . Above implies that

$$y_t^\top M_t^{-1} y_t \leq 10 \sum_{j=1}^d \frac{\lambda_{j,t+1} - \lambda_{j,t}}{\lambda_{j,t}}.$$

This allows us to derive the inequality after some algebraic computations following along the lines of Lemma 3 in Chu et al. (2011).

### C.11 Proof of Lemma C.4

**Proof** Observe that

$$\begin{aligned} & \sum_{k=1}^K \left[ \prod_{j=1}^{k-1} (1 - w(i_j)) \right] \cdot (w(i_k) - w'(i'_k)) \cdot \left[ \prod_{j=k+1}^K (1 - w'(i'_j)) \right] \\ &= \sum_{k=1}^K \left\{ \left[ \prod_{j=1}^{k-1} (1 - w(i_j)) \right] \cdot \left[ \prod_{j=k}^K (1 - w'(i'_j)) \right] - \left[ \prod_{j=1}^k (1 - w(i_j)) \right] \cdot \left[ \prod_{j=k+1}^K (1 - w'(i'_j)) \right] \right\} \\ &= \prod_{k=1}^K (1 - w'(i'_k)) - \prod_{k=1}^K (1 - w(i_k)) = R(S|\mathbf{w}) - R(S'|\mathbf{w}'), \end{aligned}$$

and also that (actually we can also see this by a symmetry argument)

$$\begin{aligned} & \sum_{k=1}^K \left[ \prod_{j=1}^{k-1} (1 - w'(i'_j)) \right] \cdot (w(i_k) - w'(i'_k)) \cdot \left[ \prod_{j=k+1}^K (1 - w(i_j)) \right] \\ &= \sum_{k=1}^K \left\{ \left[ \prod_{j=1}^k (1 - w'(i'_j)) \right] \cdot \left[ \prod_{j=k+1}^K (1 - w(i_j)) \right] - \left[ \prod_{j=1}^{k-1} (1 - w'(i'_j)) \right] \cdot \left[ \prod_{j=k}^K (1 - w(i_j)) \right] \right\} \\ &= \prod_{k=1}^K (1 - w'(i'_k)) - \prod_{k=1}^K (1 - w(i_k)) = R(S|\mathbf{w}) - R(S'|\mathbf{w}'). \end{aligned}$$

This completes the proof. ■

### C.12 Proof of Lemma 6.2

**Proof** Let  $Q = |S \setminus S^{*,\ell}|$ . It is clear that  $r(S|w^{(\ell)}) = 1 - [1 - \frac{1-\epsilon}{K}]^Q [1 - \frac{1+\epsilon}{K}]^{K-Q}$ , and then

$$\begin{aligned} r(S^{*,\ell}|w^{(\ell)}) - r(S|w^{(\ell)}) &= \left[ 1 - \frac{1+\epsilon}{K} \right]^{K-Q} \left[ \left( 1 - \frac{1-\epsilon}{K} \right)^Q - \left( 1 - \frac{1+\epsilon}{K} \right)^Q \right] \\ &= \left[ 1 - \frac{1+\epsilon}{K} \right]^{K-Q} \frac{Q}{K} \left( 1 - \frac{1-\bar{\epsilon}}{K} \right)^{Q-1} \cdot 2\epsilon \end{aligned} \quad (\text{C.30})$$

$$\begin{aligned}
 &\geq \frac{2Q}{K} \left(1 - \frac{1+\epsilon}{K}\right)^{K-1} \epsilon > \frac{2Q}{K} \left(1 - \frac{1+\epsilon}{K}\right)^K \epsilon \\
 &\geq \frac{2Q\epsilon}{K} \exp[-2(1+\epsilon)] \geq \frac{2Q\epsilon}{e^4 K} = \frac{2|S \setminus S^{*,\ell}| \epsilon}{e^4 K}. \tag{C.31}
 \end{aligned}$$

Step (C.30) is by the application of the Mean Value Theorem on function  $f : [-\epsilon, \epsilon] \rightarrow \mathbb{R}$  defined as  $f(\epsilon) = \left(1 - \frac{1-\epsilon}{K}\right)^Q$ . In Step (C.30),  $\bar{\epsilon}$  is some number lying in  $[-\epsilon, \epsilon]$ . Step (C.31) is by the fact that  $1 - \delta > \exp(-2\delta)$  for all  $\delta \in [0, 1/4]$ , and applying  $\delta = (1+\epsilon)/K \in [0, 1/4]$  to the fact.  $\blacksquare$

### C.13 Proof of Lemma 6.3

**Proof** To proceed, we define a uni-variate random variable  $J^\pi(S)$ , for  $S$  being a  $K$ -subset of  $\{1, \dots, L\}$ :

$$J^\pi(S) = \sum_{t=1}^T \sum_{j \in [L]} \frac{1(j \in S_t^\pi) \cdot 1(j \in S)}{KT}.$$

The random variable  $J^\pi(S) \in [0, 1]$  always, and note that for any fixed  $S$ , the random variable  $J^\pi(S)$  is a deterministic function of the stochastic process  $\{S_t^\pi, O_t^\pi\}_{t=1}^T$ . We denote the probability mass functions of  $J^\pi(S)$ ,  $\{S_t^\pi, O_t^\pi\}_{t=1}^T$  under instance  $\ell$  as  $P_{J^\pi(S)}^{(\ell)}$ ,  $P_{\{S_t^\pi, O_t^\pi\}_{t=1}^T}^{(\ell)}$  respectively, for each  $\ell \in \{0, 1, \dots, L\}$ . Now, we have

$$\begin{aligned}
 &\frac{1}{L} \sum_{\ell=1}^L \sum_{j \in [L] \setminus S^{*,\ell}} \mathbb{E}^{(\ell)} \left[ \sum_{t=1}^T 1(j \in S_t^\pi) \right] \\
 &= \frac{1}{L} \sum_{\ell=1}^L \sum_{j \in [L]} \mathbb{E}^{(\ell)} \left[ \sum_{t=1}^T 1(j \in S_t^\pi) \right] - \frac{1}{L} \sum_{\ell=1}^L \sum_{j \in S^{*,\ell}} \mathbb{E}^{(\ell)} \left[ \sum_{t=1}^T 1(j \in S_t^\pi) \right] \\
 &= KT - \frac{1}{L} \sum_{\ell=1}^L \mathbb{E}^{(\ell)} \left[ \sum_{t=1}^T \sum_{j \in [L]} 1(j \in S_t^\pi) 1(j \in S^{*,\ell}) \right] \\
 &= KT \left\{ 1 - \frac{1}{L} \sum_{\ell=1}^L \mathbb{E}^{(\ell)} [J^\pi(S^{*,\ell})] \right\} \\
 &\geq KT \left\{ 1 - \frac{1}{L} \sum_{\ell=1}^L \mathbb{E}^{(0)} [J^\pi(S^{*,\ell})] - \frac{1}{L} \sum_{\ell=1}^L \sqrt{\frac{1}{2} \text{KL} \left( P_{J^\pi(S^{*,\ell})}^{(0)}, P_{J^\pi(S^{*,\ell})}^{(\ell)} \right)} \right\}. \tag{C.32}
 \end{aligned}$$

$$= KT \left\{ 1 - \frac{K}{L} - \frac{1}{L} \sum_{\ell=1}^L \sqrt{\frac{1}{2} \text{KL} \left( P_{J^\pi(S^{*,\ell})}^{(0)}, P_{J^\pi(S^{*,\ell})}^{(\ell)} \right)} \right\} \tag{C.33}$$

$$\geq KT \left\{ 1 - \frac{K}{L} - \frac{1}{L} \sum_{\ell=1}^L \sqrt{\frac{1}{2} \text{KL} \left( P_{\{S_t^\pi, O_t^\pi\}_{t=1}^T}^{(0)}, P_{\{S_t^\pi, O_t^\pi\}_{t=1}^T}^{(\ell)} \right)} \right\} \tag{C.34}$$

$$\geq KT \left\{ 1 - \frac{K}{L} - \sqrt{\frac{1}{2L} \sum_{\ell=1}^L \text{KL} \left( P_{\{S_t^\pi, O_t^\pi\}_{t=1}^T}^{(0)}, P_{\{S_t^\pi, O_t^\pi\}_{t=1}^T}^{(\ell)} \right)} \right\}. \quad (\text{C.35})$$

Step (C.32) is by applying the following version of the Pinsker's inequality:

**Theorem C.6 (Pinsker's Inequality)** *Consider two probability mass functions  $P_X, P_Y$  defined on the same discrete probability space  $\mathcal{A} \subset [0, 1]$ . The following inequalities hold:*

$$|\mathbb{E}_{X \sim P_X}[X] - \mathbb{E}_{Y \sim P_Y}[Y]| \leq \sup_{A \subset \mathcal{A}} \left\{ \sum_{a \in A} P_X(a) - \sum_{a \in A} P_Y(a) \right\} \leq \sqrt{\frac{1}{2} \text{KL}(P_X, P_Y)}.$$

Step (C.33) is by the uniform cover property of the size- $K$  sets  $\{S^{*,\ell}\}_{\ell=1}^L$ :

$$\begin{aligned} \frac{1}{L} \sum_{\ell=1}^L J_t^\pi(S^{*,\ell}) &= \frac{1}{LKT} \sum_{t=1}^T \sum_{\ell=1}^L \sum_{j \in [L]} \mathbf{1}(j \in S_t^\pi) \mathbf{1}(j \in S^{*,\ell}) \\ &= \frac{1}{LKT} \sum_{t=1}^T \sum_{j \in [L]} \mathbf{1}(j \in S_t^\pi) \cdot \left[ \sum_{\ell=1}^L \mathbf{1}(j \in S^{*,\ell}) \right] \\ &= \frac{1}{LT} \sum_{t=1}^T \sum_{j \in [L]} \mathbf{1}(j \in S_t^\pi) \\ &= \frac{K}{L}. \end{aligned} \quad (\text{C.36})$$

Step (C.36) is because each  $j \in [L]$  lies in exactly  $K$  sets in the collection  $\{S^{*,\ell}\}_{\ell=1}^L$ . Step (C.34) is by the fact that  $J_t^\pi(S^{*,\ell})$  is  $\sigma(\{S_t^\pi, O_t^\pi\}_{t=1}^T)$ -measurable for any  $\ell$ , so that

$$\text{KL} \left( P_{J_t^\pi(S^{*,\ell})}^{(0)}, P_{J_t^\pi(S^{*,\ell})}^{(\ell)} \right) \leq \text{KL} \left( P_{\{S_t^\pi, O_t^\pi\}_{t=1}^T}^{(0)}, P_{\{S_t^\pi, O_t^\pi\}_{t=1}^T}^{(\ell)} \right).$$

Finally, step (C.35) follows from Jensen's inequality. Altogether, the lemma is proved.  $\blacksquare$

#### C.14 Proof of Lemma 6.4

**Proof** In the following calculation, the notation  $\sum_{\{s_\tau, o_\tau\}_{\tau=1}^t}$  means  $\sum_{\{s_\tau, o_\tau\}_{\tau=1}^t \in ([L]^{(K)} \times \mathcal{O})^t}$ . In addition, to ease the notation, we drop the subscript  $\{S_t^\pi, O_t^\pi\}_{t=1}^T$  in the notation  $P_{\{S_t^\pi, O_t^\pi\}_{t=1}^T}^{(\ell)}$ , since in this section we are only works with the stochastic process  $\{S_t^\pi, O_t^\pi\}$ , but not  $J_t^\pi(S)$ . Lastly, we adopt the convention that  $0 \log(0/0) = 0$ . Now,

$$\begin{aligned} &\text{KL} \left( P_{\{S_t^\pi, O_t^\pi\}_{t=1}^T}^{(0)}, P_{\{S_t^\pi, O_t^\pi\}_{t=1}^T}^{(\ell)} \right) \\ &= \sum_{\{s_t, o_t\}_{t=1}^T} P^{(0)} [\{S_t^\pi, O_t^\pi\}_{t=1}^T = \{s_t, o_t\}_{t=1}^T] \log \frac{P^{(0)} [\{S_t^\pi, O_t^\pi\}_{t=1}^T = \{s_t, o_t\}_{t=1}^T]}{P^{(\ell)} [\{S_t^\pi, O_t^\pi\}_{t=1}^T = \{s_t, o_t\}_{t=1}^T]} \end{aligned}$$

$$\begin{aligned}
 &= \sum_{t=1}^T \sum_{\{s_\tau, o_\tau\}_{\tau=1}^t} P^{(0)} [\{S_\tau^{\pi,0}, O_\tau^{\pi,0}\}_{\tau=1}^t = \{s_\tau, o_\tau\}_{\tau=1}^t] \\
 &\quad \cdot \log \frac{P^{(0)} [(S_t^\pi, O_t^\pi) = (s_t, o_t) \mid \{S_\tau^\pi, O_\tau^\pi\}_{\tau=1}^{t-1} = \{s_\tau, o_\tau\}_{\tau=1}^{t-1}]}{P^{(\ell)} [(S_t^\pi, O_t^\pi) = (s_t, o_t) \mid \{S_\tau^\pi, O_\tau^\pi\}_{\tau=1}^{t-1} = \{s_\tau, o_\tau\}_{\tau=1}^{t-1}]} \tag{C.37}
 \end{aligned}$$

Step (C.37) is by the Chain Rule for the KL divergence. To proceed from step (C.37), we invoke the assumption that, policy  $\pi$  is non-anticipatory and deterministic, to simplify the conditional probability terms. For any  $t, \ell$ , we have

$$\begin{aligned}
 &P^{(\ell)} [(S_t^\pi, O_t^\pi) = (s_t, o_t) \mid \{S_\tau^\pi, O_\tau^\pi\}_{\tau=1}^{t-1} = \{s_\tau, o_\tau\}_{\tau=1}^{t-1}] \\
 &= P^{(\ell)} [(S_t^\pi, O_t^\pi) = (s_t, o_t) \mid \{S_\tau^\pi, O_\tau^\pi\}_{\tau=1}^{t-1} = \{s_\tau, o_\tau\}_{\tau=1}^{t-1}, S_t^\pi = \pi_t(\{s_\tau, o_\tau\}_{\tau=1}^{t-1})] \tag{C.38}
 \end{aligned}$$

$$\begin{aligned}
 &= P^{(\ell)} [(S_t^\pi, O_t^\pi) = (s_t, o_t) \mid S_t^\pi = \pi_t(\{s_\tau, o_\tau\}_{\tau=1}^{t-1})] \tag{C.39} \\
 &= \begin{cases} P^{(\ell)} [O_t^\pi = o_t \mid S_t^\pi = s_t] & \text{if } s_t = \pi_t(\{s_\tau, o_\tau\}_{\tau=1}^{t-1}) \\ 0 & \text{otherwise} \end{cases}.
 \end{aligned}$$

Step (C.38) is by our supposition on  $\pi$ , which implies  $S_t^\pi = \pi_t(S_1^\pi, O_1^\pi, \dots, S_{t-1}^\pi, O_{t-1}^\pi)$  for each  $\ell \in \{0, \dots, L\}$ , where  $\pi_t$  is a deterministic function. Step (C.39) is by the model assumption that, conditioned on the arm  $S_t^\pi$  pulled at time step  $t$ , the outcome  $O_t^\pi$  is independent of the history from time 1 to  $t-1$ .

By our adopted convention that  $0 \log(0/0) = 0$  (equivalently, we are removing the terms with  $0 \log(0/0)$  from the sum), we simplify the sum in (C.37) as follows:

$$\begin{aligned}
 \text{(C.37)} &= \sum_{t=1}^T \sum_{\{s_\tau, o_\tau\}_{\tau=1}^t} \left\{ P^{(0)} [O_t^\pi = o_t \mid S_t^\pi = s_t] P^{(0)} [\{S_\tau^\pi, O_\tau^\pi\}_{\tau=1}^{t-1} = \{s_\tau, o_\tau\}_{\tau=1}^{t-1}] \right. \\
 &\quad \left. \cdot \mathbf{1}(s_t = \pi_t(\{s_\tau, o_\tau\}_{\tau=1}^{t-1})) \cdot \log \frac{P_{O_t^\pi | S_t^\pi}^{(0)} [O_t^\pi = o_t \mid S_t^\pi = s_t]}{P_{O_t^\pi | S_t^\pi}^{(\ell)} [O_t^\pi = o_t \mid S_t^\pi = s_t]} \right\} \\
 &= \sum_{t=1}^T \sum_{s_t, o_t} P^{(0)} [S_t^\pi = s_t] P^{(0)} [O_t^\pi = o_t \mid S_t^\pi = s_t] \log \frac{P_{O_t^\pi | S_t^\pi}^{(0)} [O_t^\pi = o_t \mid S_t^\pi = s_t]}{P_{O_t^\pi | S_t^\pi}^{(\ell)} [O_t^\pi = o_t \mid S_t^\pi = s_t]} \\
 &= \sum_{t=1}^T \sum_{s_t \in [L]^{(K)}} P^{(0)} [S_t^\pi = s_t] \sum_{k=0}^K P_{O_t^\pi | S_t^\pi}^{(0)} [O_t^\pi = o^k \mid S_t^\pi = s_t] \log \frac{P_{O_t^\pi | S_t^\pi}^{(0)} [O_t^\pi = o^k \mid S_t^\pi = s_t]}{P_{O_t^\pi | S_t^\pi}^{(\ell)} [O_t^\pi = o^k \mid S_t^\pi = s_t]}. \tag{C.40}
 \end{aligned}$$

$$= \sum_{t=1}^T \sum_{s_t \in [L]^{(K)}} P^{(0)} [S_t^\pi = s_t] \cdot \text{KL} \left( P_{O_t^\pi | S_t^\pi}^{(0)}(\cdot \mid s_t) \parallel P_{O_t^\pi | S_t^\pi}^{(\ell)}(\cdot \mid s_t) \right). \tag{C.41}$$

Altogether, the Lemma is proved. ■

**C.15 Proof of Lemma 6.5**

**Proof** Consider a fixed  $\ell \in \{1, \dots, L\}$ , and recall our notation  $\underline{a} = a$  if  $1 \leq a \leq L$ , and  $\underline{a} = a - L$  if  $L < a \leq 2L$ . Recall the weigh parameters  $\{w^{(0)}(i)\}_{i \in [L]}$ ,  $\{w^{(\ell)}(i)\}_{i \in [L]}$  of instances  $0, \ell$ :

- For  $i \in \{\underline{\ell}, \underline{\ell} + 1, \dots, \underline{\ell} + K - 1\}$ ,  $w^{(\ell)}(i) = \frac{1+\epsilon}{K}$ , while  $w^{(0)}(i) = \frac{1-\epsilon}{K}$ .
- For  $i \in [L] \setminus \{\underline{\ell}, \underline{\ell} + 1, \dots, \underline{\ell} + K - 1\}$ ,  $w^{(\ell)}(i) = w^{(0)}(i) = \frac{1-\epsilon}{K}$ .

By an abuse of notation, we see  $S^{*,\ell}, S$  as size- $K$  subsets (instead of order sets) when we take the intersection  $S^{*,\ell} \cap S$ . More precisely, for  $S = (i_1, \dots, i_k)$ , we denote  $S^{*,\ell} \cap S = \{\iota = \underline{j} \text{ for some } j \in \{\ell, \dots, \ell + K - 1\}, \text{ and } \iota = i_j \text{ for some } j \in \{1, \dots, K\}\}$ .

It is clear that, for any ordered  $K$  set  $S \in [L]^{(K)}$ ,

$$\begin{aligned} \text{KL} \left( P_{O_t^\pi | S_t^\pi}^{(0)}(\cdot | S) \parallel P_{O_t^\pi | S_t^\pi}^{(\ell)}(\cdot | S) \right) &= \sum_{i \in S^{*,\ell} \cap S} P^{(0)}(i \text{ is observed}) \cdot \text{KL} \left( \frac{1-\epsilon}{K}, \frac{1+\epsilon}{K} \right) \\ &\leq |S^{*,\ell} \cap S| \cdot \text{KL} \left( \frac{1-\epsilon}{K}, \frac{1+\epsilon}{K} \right). \end{aligned}$$

For any  $S \in [L]^{(K)}$ ,

$$\begin{aligned} \frac{1}{L} \sum_{\ell=1}^L |S^{*,\ell} \cap S| &= \frac{1}{L} \sum_{\ell=1}^L \sum_{i \in S} \mathbf{1}\{i \in S^{*,\ell}\} = \sum_{i \in S} \frac{1}{L} \sum_{\ell=1}^L \mathbf{1}\{i \in S^{*,\ell}\} \\ &= \sum_{i \in S} \frac{K}{L} = K \cdot \frac{K}{L} = \frac{K^2}{L}. \end{aligned}$$

Therefore,

$$\begin{aligned} &\frac{1}{L} \sum_{\ell=1}^L \sum_{s_t \in [L]^{(K)}} P^{(0)}[S_t^\pi = s_t] \cdot \text{KL} \left( P_{O_t^\pi | S_t^\pi}^{(0)}(\cdot | s_t) \parallel P_{O_t^\pi | S_t^\pi}^{(\ell)}(\cdot | s_t) \right) \\ &= \frac{1}{L} \sum_{\ell=1}^L \sum_{s_t \in [L]^{(K)}} P^{(0)}[S_t^{\pi,0} = s_t] \cdot |S^{*,\ell} \cap s_t| \cdot \text{KL} \left( \frac{1-\epsilon}{K}, \frac{1+\epsilon}{K} \right) \\ &= \text{KL} \left( \frac{1-\epsilon}{K}, \frac{1+\epsilon}{K} \right) \cdot \sum_{s_t \in [L]^{(K)}} P^{(0)}[S_t^{\pi,0} = s_t] \cdot \left( \frac{1}{L} \sum_{\ell=1}^L |S^{*,\ell} \cap s_t| \right) \\ &= \text{KL} \left( \frac{1-\epsilon}{K}, \frac{1+\epsilon}{K} \right) \cdot \frac{K^2}{L} = \frac{K}{L} \left[ (1-\epsilon) \log \left( \frac{1-\epsilon}{1+\epsilon} \right) + (K-1+\epsilon) \log \left( \frac{K-1+\epsilon}{K-1-\epsilon} \right) \right]. \end{aligned}$$

■

## Appendix D. Additional numerical results

In this section, we present a long table of extra numerical results in the standard cascading setting, which implies the same conclusion as in Section 7.1.

Table 4: The performances of TS-CASCADE, CTS, CASCADEUCB1 and CASCADEKL-UCB when  $L \in \{16, 32, 64, 256, 512, 1024, 2048\}$ ,  $K \in \{2, 4\}$ . The first column shows the mean and the standard deviation of  $\text{Reg}(T)$  and the second column shows the average running time in seconds.

|               | $L$ | $K$ | $\text{Reg}(T)$                         | Running time |
|---------------|-----|-----|---|--------------|
| TS-CASCADE    | 16  | 2   | $4.10 \times 10^2 \pm 3.03 \times 10^1$ | 1.68         |
| CTS           | 16  | 2   | $1.61 \times 10^2 \pm 2.07 \times 10^1$ | 4.96         |
| CASCADEUCB1   | 16  | 2   | $1.35 \times 10^3 \pm 5.77 \times 10^1$ | 1.33         |
| CASCADEKL-UCB | 16  | 2   | $1.64 \times 10^4 \pm 6.49 \times 10^3$ | 2.18         |
| TS-CASCADE    | 16  | 4   | $3.54 \times 10^2 \pm 2.70 \times 10^1$ | 3.75         |
| CTS           | 16  | 4   | $1.20 \times 10^2 \pm 1.56 \times 10^1$ | 6.29         |
| CASCADEUCB1   | 16  | 4   | $1.10 \times 10^3 \pm 4.60 \times 10^1$ | 1.23         |
| CASCADEKL-UCB | 16  | 4   | $2.02 \times 10^4 \pm 6.06 \times 10^3$ | 4.84         |
| TS-CASCADE    | 32  | 2   | $7.55 \times 10^2 \pm 2.38 \times 10^1$ | 1.85         |
| CTS           | 32  | 2   | $3.30 \times 10^2 \pm 2.84 \times 10^1$ | 5.93         |
| CASCADEUCB1   | 32  | 2   | $2.77 \times 10^3 \pm 9.72 \times 10^1$ | 1.53         |
| CASCADEKL-UCB | 32  | 2   | $1.86 \times 10^4 \pm 7.04 \times 10^3$ | 2.33         |
| TS-CASCADE    | 32  | 4   | $6.52 \times 10^2 \pm 2.95 \times 10^1$ | 1.85         |
| CTS           | 32  | 4   | $2.84 \times 10^2 \pm 8.25 \times 10^1$ | 5.88         |
| CASCADEUCB1   | 32  | 4   | $2.36 \times 10^3 \pm 5.53 \times 10^1$ | 1.52         |
| CASCADEKL-UCB | 32  | 4   | $2.21 \times 10^4 \pm 7.13 \times 10^3$ | 2.90         |
| TS-CASCADE    | 64  | 2   | $1.39 \times 10^3 \pm 3.37 \times 10^1$ | 2.23         |
| CTS           | 64  | 2   | $6.67 \times 10^2 \pm 3.43 \times 10^1$ | 7.85         |
| CASCADEUCB1   | 64  | 2   | $5.52 \times 10^3 \pm 1.11 \times 10^2$ | 2.05         |
| CASCADEKL-UCB | 64  | 2   | $2.17 \times 10^4 \pm 8.03 \times 10^3$ | 2.94         |



|               |     |   |   |                    |
|---------------|-----|---|---|--------------------|
| TS-CASCADE    | 64  | 4 | $1.22 \times 10^3 \pm 3.08 \times 10^1$ | 2.23               |
| CTS           | 64  | 4 | $5.27 \times 10^2 \pm 2.89 \times 10^1$ | 7.83               |
| CASCADEUCB1   | 64  | 4 | $4.75 \times 10^3 \pm 7.54 \times 10^1$ | 2.06               |
| CASCADEKL-UCB | 64  | 4 | $2.72 \times 10^4 \pm 8.71 \times 10^3$ | 3.32               |
| TS-CASCADE    | 128 | 2 | $2.46 \times 10^3 \pm 4.63 \times 10^1$ | 3.08               |
| CTS           | 128 | 2 | $1.36 \times 10^3 \pm 4.39 \times 10^1$ | $1.17 \times 10^1$ |
| CASCADEUCB1   | 128 | 2 | $1.07 \times 10^4 \pm 1.70 \times 10^2$ | 3.74               |
| CASCADEKL-UCB | 128 | 2 | $2.19 \times 10^4 \pm 8.11 \times 10^3$ | 4.49               |
| TS-CASCADE    | 128 | 4 | $2.25 \times 10^3 \pm 5.95 \times 10^1$ | 3.09               |
| CTS           | 128 | 4 | $1.06 \times 10^3 \pm 4.95 \times 10^1$ | $1.17 \times 10^1$ |
| CASCADEUCB1   | 128 | 4 | $9.23 \times 10^3 \pm 1.27 \times 10^2$ | 3.58               |
| CASCADEKL-UCB | 128 | 4 | $2.95 \times 10^4 \pm 1.05 \times 10^4$ | 4.71               |
| TS-CASCADE    | 256 | 2 | $4.31 \times 10^3 \pm 2.27 \times 10^2$ | 4.73               |
| CTS           | 256 | 2 | $2.76 \times 10^3 \pm 2.37 \times 10^2$ | $1.93 \times 10^1$ |
| CASCADEUCB1   | 256 | 2 | $1.92 \times 10^4 \pm 2.73 \times 10^2$ | 7.57               |
| CASCADEKL-UCB | 256 | 2 | $2.19 \times 10^4 \pm 8.11 \times 10^3$ | 7.24               |
| TS-CASCADE    | 256 | 4 | $3.99 \times 10^3 \pm 1.37 \times 10^2$ | 4.79               |
| CTS           | 256 | 4 | $2.16 \times 10^3 \pm 8.27 \times 10^1$ | $1.93 \times 10^1$ |
| CASCADEUCB1   | 256 | 4 | $1.74 \times 10^4 \pm 1.65 \times 10^2$ | 6.59               |
| CASCADEKL-UCB | 256 | 4 | $2.98 \times 10^4 \pm 1.03 \times 10^4$ | 7.48               |
| TS-CASCADE    | 512 | 2 | $6.83 \times 10^3 \pm 5.28 \times 10^2$ | 8.05               |
| CTS           | 512 | 2 | $5.39 \times 10^3 \pm 1.92 \times 10^2$ | $3.49 \times 10^1$ |
| CASCADEUCB1   | 512 | 2 | $2.51 \times 10^4 \pm 1.47 \times 10^2$ | $1.30 \times 10^1$ |
| CASCADEKL-UCB | 512 | 2 | $2.19 \times 10^4 \pm 8.10 \times 10^3$ | $1.31 \times 10^1$ |
| TS-CASCADE    | 512 | 4 | $6.77 \times 10^3 \pm 2.00 \times 10^2$ | 8.20               |
| CTS           | 512 | 4 | $4.27 \times 10^3 \pm 1.73 \times 10^2$ | $3.50 \times 10^1$ |
| CASCADEUCB1   | 512 | 4 | $3.03 \times 10^4 \pm 3.33 \times 10^2$ | $1.27 \times 10^1$ |
| CASCADEKL-UCB | 512 | 4 | $2.98 \times 10^4 \pm 1.03 \times 10^4$ | $1.34 \times 10^1$ |

|               |      |   |   |                    |
|---------------|------|---|---|--------------------|
| TS-CASCADE    | 1024 | 2 | $9.82 \times 10^3 \pm 6.82 \times 10^2$ | $1.47 \times 10^1$ |
| CTS           | 1024 | 2 | $1.06 \times 10^4 \pm 2.26 \times 10^2$ | $6.73 \times 10^1$ |
| CASCADEUCB1   | 1024 | 2 | $2.60 \times 10^4 \pm 3.32 \times 10^1$ | $1.01 \times 10^2$ |
| CASCADEKL-UCB | 1024 | 2 | $2.19 \times 10^4 \pm 8.05 \times 10^3$ | $2.57 \times 10^1$ |
| TS-CASCADE    | 1024 | 4 | $1.06 \times 10^4 \pm 6.23 \times 10^2$ | $1.51 \times 10^1$ |
| CTS           | 1024 | 4 | $8.59 \times 10^3 \pm 2.94 \times 10^2$ | $6.84 \times 10^1$ |
| CASCADEUCB1   | 1024 | 4 | $3.87 \times 10^4 \pm 1.32 \times 10^2$ | $2.56 \times 10^1$ |
| CASCADEKL-UCB | 1024 | 4 | $3.00 \times 10^4 \pm 1.01 \times 10^4$ | $4.07 \times 10^1$ |
| TS-CASCADE    | 2048 | 2 | $1.27 \times 10^4 \pm 1.72 \times 10^3$ | $4.40 \times 10^1$ |
| CTS           | 2048 | 2 | $1.99 \times 10^4 \pm 5.85 \times 10^2$ | $1.25 \times 10^2$ |
| CASCADEUCB1   | 2048 | 2 | $2.62 \times 10^4 \pm 9.19$             | $5.67 \times 10^1$ |
| CASCADEKL-UCB | 2048 | 2 | $2.19 \times 10^4 \pm 8.05 \times 10^3$ | $1.08 \times 10^2$ |
| TS-CASCADE    | 2048 | 4 | $1.50 \times 10^4 \pm 7.27 \times 10^2$ | $4.55 \times 10^1$ |
| CTS           | 2048 | 4 | $1.67 \times 10^4 \pm 3.99 \times 10^2$ | $1.25 \times 10^2$ |
| CASCADEUCB1   | 2048 | 4 | $4.01 \times 10^4 \pm 4.18 \times 10^1$ | $5.57 \times 10^1$ |
| CASCADEKL-UCB | 2048 | 4 | $3.00 \times 10^4 \pm 1.00 \times 10^4$ | $5.47 \times 10^1$ |