
Thompson Sampling for Linearly Constrained Bandits

Vidit Saxena
KTH Royal Institute of Technology
Stockholm, Sweden
vidits@kth.se

Joseph E. Gonzalez
University of California
Berkeley
jegonzal@berkeley.edu

Joakim Jaldén
KTH Royal Institute of Technology
Stockholm, Sweden
jalden@kth.se

Abstract

We address multi-armed bandits (MAB) where the objective is to maximize the cumulative reward under a probabilistic linear constraint. For a few real-world instances of this problem, constrained extensions of the well-known Thompson Sampling (TS) heuristic have recently been proposed. However, finite-time analysis of constrained TS is challenging; as a result, only $O(\sqrt{T})$ bounds on the cumulative reward loss (i.e., the *regret*) are available. In this paper, we describe LinConTS, a TS-based algorithm for bandits that place a linear constraint on the probability of earning a reward in every round. We show that for LinConTS, the regret as well as the cumulative constraint violations are upper bounded by $O(\log T)$ for the suboptimal arms. We develop a proof technique that relies on careful analysis of the dual problem and combine it with recent theoretical work on unconstrained TS. Through numerical experiments on two real-world datasets, we demonstrate that LinConTS outperforms an asymptotically optimal upper confidence bound (UCB) scheme in terms of simultaneously minimizing the regret and the violation.

1 Introduction

Multi-armed bandits (MAB) are a systematic way of modeling sequential decision problems. In MAB problems, an agent plays a sequence of arms aimed at optimizing some cumulative objective over T discrete time intervals (“rounds”). The agent strives to achieve

this goal by sequentially *exploring* the available arms and *exploiting* historical rewards from previously selected arms. MABs have been successfully applied to problems in dynamic pricing, online procurement, and digital advertising, where the goal is to minimize the *regret* (i.e., the cumulative reward loss) over a finite time horizon [Russo et al., 2018].

Despite the popularity of MABs, the unconstrained reward maximization goal does not extend to several commonly encountered sequential decision problems. In this paper, we consider a constrained MAB problem where the objective function is subject to a specific type of linear constraint, namely, that the probability of receiving a reward in any round exceeds a fixed, predefined threshold. A few recently studied applications of this type of constraint are listed below:

Weblink selection: Several online content providers earn revenue by advertising affiliate weblinks on their page. In this context, the goal is to select a subset of links from a pool of available links, which collectively maximize the provider’s ad revenue. At the same time, the provider typically wants to avoid displaying some high-revenue links that have a low *click-through rate* (CTR), since such links may be perceived as clutter and drive down user satisfaction. In [Chen et al., 2018], this problem is formulated as a constrained MAB where the probability of clicking on one of the displayed weblinks must exceed a minimum CTR threshold.

Wireless Rate selection: Wireless communication networks strive to optimize the data transmission rate for packetized information bits. Aggressive data transmission rates carry more per-packet information but also suffer from more frequent packet failures compared to conservative rates. For latency-sensitive applications such as video streaming and online gaming, the goal of wireless rate selection is to maximize the average data throughput while simultaneously maintaining a minimum packet success frequency. Recently, [Saxena et al., 2019] proposed a constrained MAB approach for this rate selection problem that was shown to outperform state-of-the-art techniques.

Thompson sampling (TS) is an efficient Bayesian heuristic for MAB optimization [Thompson, 1933, Russo et al., 2018]. TS operates by conditioning the probability of selecting an optimal arm on the historical rewards observed in previous rounds. The regret for TS has been shown to scale as $O(\log T)$ in [Kaufmann et al., 2012, Agrawal and Goyal, 2013], where the constant factors depend on the distribution of the underlying problem parameters. These bounds asymptotically achieve the optimal regret established in [Lai and Robbins, 1985]. Further, [Ferreira et al., 2018] proposed TS-based algorithm for a stochastic packing problem commonly referred to as *Bandits with Knapsacks (BwK)*, for which the distribution-independent regret upper bounds were shown to scale as $O(\sqrt{T})$. For BwK, $O(\log T)$ distribution-dependent lower bounds were developed in [Flajolet and Jaillet, 2015].

In this paper, we propose *LinConTS*, a TS-based algorithm for the specific type of linear constraint described above. Our algorithm encodes a probabilistic arm selection policy, where the arm selection weights are obtained by solving a linear program (LP) subroutine in every round. We denote the arms supported by the stationary optimal policy as *optimal arms* and the rest as *suboptimal arms*. We show that for the rounds where the suboptimal arms are played, the LinConTS regret is upper bounded by $O(\log T)$. Further, we show that for these arms, the *violation* metric, i.e., the cumulative number of constraint violations, is also upper bounded by $O(\log T)$. For the rounds where the optimal arms are played, the regret and violation metrics scale only as fast as $O(\sqrt{T})$.

Proof technique: The recipe for MABs with linearly constrained objectives is to invoke a LP subroutine in each round. Subsequently, the central problem in analyzing constrained MABs arises from the fact that the arm parameters cannot be compared directly; instead, they interact with each other through the LP. To handle this issue, our proof technique relies on careful examination of the dual LP problem. In particular, to bound the probability of selecting a particular arm, we analyze stochastic perturbations of the convex polytope enclosing the LP feasible region. We use these perturbation results in conjunction with theoretical analysis of TS developed in [Agrawal and Goyal, 2013].

2 Problem Formulation

Consider a MAB with N arms, where at each time step $t = 1, 2, \dots, T$, one of the arms must be played. Playing the arm $i(t) \in \{1, \dots, N\}$ results in a Bernoulli-distributed *reward event* with mean $\mu_{i(t)}$. Reward events are independent across the arms and across

successive plays of an arm. In case of a successful reward event, a reward value $r_{i(t)} \in (0, 1]$ is immediately collected by the player. We assume that the μ_i s are *a priori* unknown while the r_i s are deterministic and known in advance. In the context of weblink selection, the arms (μ_i, r_i) correspond to the pool of available weblinks. For a selected weblink that is displayed on a target webpage, the reward event denotes a Bernoulli-distributed user click with an unknown mean CTR μ_i . Subsequently, the reward value r_i captures the revenue generated when a user clicks on the displayed link. If a user clicks on the displayed link, the reward event is deemed to be true and the corresponding reward value is collected. On the other hand if a user fails to click on the link, the reward event is false and no reward is collected in that round.

The problem objective is to maximize the cumulative expected reward value, $\sum_{t=1}^T \mu_{i(t)} r_{i(t)}$ subject to the constraint $(1/T) \sum_{t=1}^T \mu_{i(t)} \geq \eta$ where $\eta \in [0, 1]$. The constraint, which ensures that at least a fraction η of rounds result in a reward event, must be satisfied with a high probability. This problem can be formulated using the LP relaxation

$$\begin{aligned} LP(\boldsymbol{\mu}) : \text{maximize } & \sum_i x_i \mu_i r_i \\ \text{subject to } & \sum_i x_i \mu_i \geq \eta, \\ & \sum_i x_i = 1, \\ & x_i \geq 0, i = 1, \dots, N. \end{aligned} \quad (1)$$

where x_i denotes probability of selecting arm i and therefore $\mathbf{x} = [x_1, \dots, x_N]$ is the *probabilistic arm selection vector*. The arms assigned probability mass $x_i > 0$ by the LP solution constitute the set of *optimal arms* that are played with a non-zero probability.

Stationary optimal policy: If the μ_i s were known in advance, the stationary optimal policy is found by solving $LP(\boldsymbol{\mu})$. We denote the solution of this LP with $\mathbf{x}^* = [x_1^*, \dots, x_N^*]$, which is the optimal stationary probabilistic arm selection vector. The expected reward of the stationary optimal policy is therefore $r^* = \sum_i x_i^* \mu_i r_i$. The arms assigned a probability mass $x_i^* > 0$ by the optimal policy constitute the set of stationary *optimal arms* and the rest of the arms are stationary *suboptimal arms*.

Assumptions: We assume that the arm with the highest expected reward value, $i_{\max} = \operatorname{argmax}_i \mu_i r_i$, does not also satisfy the constraint, i.e., $\mu_{i_{\max}} < \eta$ (otherwise the problem reduces to an unconstrained MAB problem). Consequently, to simultaneously satisfy the reward maximization and the constraint satisfaction objectives, a mixture of high-reward-value and high-reward-event-

probability arms is required. Further, we assume that a strictly feasible solution exists for $LP(\boldsymbol{\mu})$ and that there are no degenerate cases (i.e., the stationary optimal policy is a unique mixture of the arms).

Regret and violation: The performance of any MAB algorithm is typically measured in terms of its regret, $\mathcal{R}(T)$, which is the cumulative reward loss compared to selecting the reward-maximizing arm in hindsight. However, in a constrained MAB setting, the optimal arms balance the need for reward maximization with constraint fulfilment. As a consequence, it is possible for a sub-optimal policy to exceed the cumulative reward achieved by the optimal policy, e.g., by frequently selecting a high-reward arm that instead violates the constraint. Therefore, we additionally define a *violation* metric $\mathcal{V}(T)$ to measure the cumulative constraint violations until time T . For the estimated reward event probabilities in the t^{th} round, $\tilde{\boldsymbol{\mu}}_t = [\tilde{\mu}_{1,t}, \dots, \tilde{\mu}_{N,t}]$, let the instantaneous, probabilistic arm selection vector is $x_t = [x_{1,t}, \dots, x_{N,t}]$. Then the expected violation and regret until time T are given by

$$\mathbb{E}[\mathcal{R}(T)] = \mathbb{E}\left[\left[T r^* - \sum_{t=1}^T \sum_i x_{i,t} \mu_i r_i\right]_+\right] \quad (2)$$

$$= \left[\sum_i \Delta_i \mathbb{E}[k_i(T+1)]\right]_+, \quad (3)$$

$$\mathbb{E}[\mathcal{V}(T)] = \mathbb{E}\left[\left[T \eta - \sum_{t=1}^T \sum_i x_{i,t} \mu_i\right]_+\right] \quad (4)$$

$$= \left[\sum_i (\eta - \mu_i) \mathbb{E}[k_i(T+1)]\right]_+, \quad (5)$$

where $k_i(t)$ denotes the number of times that the arm i is played until time $t - 1$, $\Delta_i = r^* - \mu_i r_i$ is the expected loss in reward for arm i , and $[x]_+ = \max\{x, 0\}$. Intuitively, the reward metric measures the amount of reward lost due to selecting suboptimal arms, or due to selecting optimal arms with a frequency different from the stationary optimal policy. In order to minimize regret, a policy could simply choose to search for the reward-maximizing arm without optimizing for the constraint. To account for this behaviour, the violation metric keeps track of the constraint violations of the policy. Both regret and violation metrics must be simultaneously optimized by any useful policy. In the next sections, we develop finite-time upper bounds on regret and violation metrics for LinConTS.

3 Algorithm and Finite-Time Analysis

3.1 LinConTS

TS assigns prior distributions over the unknown MAB parameters. Subsequently for the played arm in every

Algorithm 1 LinConTS

- 1: **Input:** Reward Values $r_{\{1, \dots, N\}}$, Constraint η
- 2: **Initialize:** $\alpha_{\{1, \dots, N\}, 0} = 1, \beta_{\{1, \dots, N\}, 0} = 1$.
- 3: **for** Time index $t = 1$ **to** T **do**
- 4: **if** $t < N$ **then**
- 5: $i(t) = t$
- 6: **else**
- 7: Sample $\tilde{\mu}_{i,t} \sim \text{Beta}(\alpha_{i,t-1}, \beta_{i,t-1})$ for each arm $i = 1, \dots, N$.
- 8: Solve, if feasible, the linear program:

$$LP(\tilde{\boldsymbol{\mu}}_t) : \text{maximize } \sum_i x_{i,t} \tilde{\mu}_{i,t} r_i$$

$$\text{subject to } \begin{cases} \sum_i x_{i,t} \tilde{\mu}_{i,t} \geq \eta \\ \sum_i x_{i,t} = 1 \\ x_{i,t} \geq 0 \forall i \in \{1, \dots, N\} \end{cases}, \quad (6)$$

- 9: **if** a (feasible) optimal solution exists **then**
 - 10: Sample $i(t) \sim [x_{1,t}, \dots, x_{N,t}]$
 - 11: **else**
 - 12: Sample $i(t)$ uniformly from $\{1, \dots, N\}$.
 - 13: **end if**
 - 14: **end if**
 - 15: **Observe:** Reward event $c_{i(t),t} \in \{0, 1\}$.
 - 16: **Update:**

$$\begin{aligned} \alpha_{i(t),t} &= \alpha_{i(t),t-1} + c_{i(t),t} \\ \beta_{i(t),t} &= \beta_{i(t),t-1} + (1 - c_{i(t),t}). \end{aligned}$$
 - 17: **end for**
-

round, the arm posterior is updated using the observed reward. For the problem studied in this paper, the unknown parameters are the reward event means μ_1, \dots, μ_N . Since these parameters are Bernoulli-distributed, a suitable choice of prior is the Beta distribution [Russo et al., 2018]. The LinConTS algorithm is described in Alg. 1. At the beginning of the experiment, the reward values $r_i, i \in \{1, \dots, N\}$ and the constraint η are provided. For each arm, LinConTS assigns independent uniform priors distributed as $\text{Beta}(\alpha_{i,0} = 1, \beta_{i,0} = 1)$. At every time step $t = 1, \dots, T$, LinConTS obtains *Thompson samples* $\tilde{\mu}_{i,t} \sim \text{Beta}(\alpha_{i,t-1}, \beta_{i,t-1})$ and subsequently solves a LP parameterized by these samples to obtain the instantaneous, probabilistic arm selection vector $x_t = [x_{1,t}, \dots, x_{N,t}]$. This vector is sampled to obtain the playing arm, $i(t) \sim x_t$, if the LP is feasible. Subsequently, the reward event $c_{i(t),t} \sim \text{Bern}(\mu_{i(t)})$ is observed and the reward $c_{i(t),t} r_{i(t)}$ is collected. The parameter distribution for the played arm is updated according to the rule $\text{Beta}(\alpha_{i(t),t-1}, \beta_{i(t),t-1}) \leftarrow (c_{i(t),t}, 1 - c_{i(t),t})$.

3.2 Regret and Violation Upper Bounds

The theoretical analysis of TS is challenging, and as such only an empirical evaluation of its performance was available until recently [Chapelle and Li, 2011]. The chief reason for TS' theoretical intractability is that, owing to its randomized nature, novel techniques are required to bound the number of draws for suboptimal arms, which were first introduced in [Agrawal and Goyal, 2012] and extended in [Kaufmann et al., 2012, Agrawal and Goyal, 2013]. At a high level, these techniques compare the Thompson samples for each arm with carefully selected thresholds related to empirical estimates of the arm parameters and their true values. Subsequently, by using known bounds on the sums of Bernoulli random variables, these techniques bound the expected number of times that any suboptimal arm is played.

In case of LinConTS, the selected arm in round t depends on the solution of LP($\tilde{\boldsymbol{\mu}}_t$). Therefore, to bound the probability of selecting a suboptimal arm, we need to analyze the LP in (1) with perturbed reward event means. We address this challenge by formulating the Lagrangian dual of the LP that can be solved to obtain specific thresholds for each suboptimal arm, which, when breached under certain well-specified conditions, lead to the suboptimal arm being selected with a nonzero probability. We use these thresholds in combination with the proof technique of [Agrawal and Goyal, 2013] to show that for the suboptimal arms, the expected number of plays is bounded by $O(\log T)$. Under the assumption of no degenerate cases and that the reward-maximizing arm does not simultaneously satisfy the constraint, exactly two arms support (1). This well-known result for LPs satisfies the following intuition: the optimal solution to (1) frequently selects a high-reward-value arm, which comes at the cost of frequent failures (i.e., no-reward events). Hence, to satisfy the constraint, the optimal solution sometimes picks another, low-reward-value arm that has a high success probability. Without loss of generality, we denote the optimal arm indices with 1 and 2 respectively with parameters that satisfy $0 < \mu_1 < \eta < \mu_2 \leq 1$ and $r_1\mu_1 > r_2\mu_2$. Subsequently, we obtain the following regret and violation bounds:

Theorem 1. *The expected regret for LinConTS,*

$$\begin{aligned} \mathbb{E}[\mathcal{R}(T)] &\leq \left[\sum_{i \neq \{1,2\}} \frac{(1+\gamma)^2}{d(\mu_i, \xi_i)} \Delta_i^+ \right] \log T + \Delta_{\max}^+ \cdot 18\sqrt{2T \log 2} + O\left(\frac{N}{\gamma^2}\right), \end{aligned}$$

$$\begin{aligned} \text{where } \xi_i &= \frac{(r_1-r_2)\mu_1\mu_2}{(r_i-r_2)\mu_2-(r_i-r_1)\mu_1} > \mu_i, \quad \gamma \in (0, 1] \\ d(\mu_i, \xi_i) &= \mu_i \log \frac{\mu_i}{\xi_i} + (1-\mu_i) \log \frac{1-\mu_i}{1-\xi_i}, \quad \Delta_i^+ = \\ &\max\{0, r^* - \mu_i r_i\}, \text{ and } \Delta_{\max}^+ = \max_{i \in [N]} \Delta_i^+. \end{aligned}$$

Theorem 2. *The expected violation for LinConTS,*

$$\begin{aligned} \mathbb{E}[\mathcal{V}(T)] &\leq \left[\sum_{i \neq \{1,2\}} \frac{(1+\gamma)^2}{d(\mu_i, \xi_i)} \delta_i^+ \right] \log T + \delta_{\max}^+ \cdot 18\sqrt{2T \log 2} + O\left(\frac{N}{\gamma^2}\right), \end{aligned}$$

where $\delta_i^+ = \max\{0, \eta - \mu_i\}$ and $\delta_{\max}^+ = \max_{i \in [N]} \delta_i^+$

3.3 LP Perturbation Analysis

The Lagrangian dual $\mathcal{L}(\mathbf{x}, \lambda, \nu, \boldsymbol{\psi})$ of the primal LP in (1) can be formulated as

$$\begin{aligned} \mathcal{L}(\mathbf{x}, \lambda, \nu, \boldsymbol{\psi}) &= - \sum_i x_i \mu_i r_i \\ &\quad + \lambda \left(\eta - \sum_i x_i \mu_i \right) \\ &\quad + \nu \left(\sum_i x_i - 1 \right) \\ &\quad - \sum_i \psi_i x_i, \\ &= \lambda \eta - \nu - \sum_i (r_i \mu_i + \lambda \mu_i - \nu + \psi_i) x_i, \quad (7) \end{aligned}$$

where $\mathbf{x} \in \mathbf{R}^N$ is the optimization variable, and $\lambda \in \mathbf{R}$, $\nu \in \mathbf{R}$, and $\boldsymbol{\psi} \in \mathbf{R}^N$ are the Lagrange dual variables. The corresponding dual function

$$\begin{aligned} g(\lambda, \nu, \boldsymbol{\psi}) &= \inf_{\mathbf{x}} \mathcal{L}(\mathbf{x}, \lambda, \nu, \boldsymbol{\psi}) \quad (8) \\ &= \begin{cases} \lambda \eta - \nu & r_i \mu_i + \lambda \mu_i - \nu + \psi_i = 0, \\ & \lambda \geq 0, \\ & \psi_i \geq 0. \\ -\infty & \text{otherwise.} \end{cases} \quad (9) \end{aligned}$$

We assume that a strictly feasible solution exists to the LP($\boldsymbol{\mu}$), otherwise there can be no policy that achieves the objective. Then, strong duality holds from Slater's condition and the primal optimal is equal to the dual optimal, i.e., there no duality gap. Hence, from the KKT optimality conditions,

$$\lambda^* (\eta - \sum_i x_i^* \mu_i) = 0 \quad (10)$$

$$\nu^* (\sum_i x_i^* - 1) = 0, \text{ and} \quad (11)$$

$$\sum_i \psi_i^* x_i^* = 0, \quad (12)$$

where $\lambda^*, \nu^*, \boldsymbol{\psi}^*$ are the optimal duals. Since $\psi_i^* \geq 0, x_i^* \geq 0$, each term of $\sum_i \psi_i^* x_i^*$ is non-negative. Consequently for $\sum_i \psi_i^* x_i^* = 0$ to hold, $x_i^* > 0$ implies that $\psi_i^* = 0$ and $x_i^* > 0$ implies that $\psi_i^* = 0$. Under the assumption of no degenerate cases, strict

complementarity holds, i.e., either $x_i^* > 0$ or $\psi_i^* > 0$ [Boyd and Vandenberghe, 2004, Chap. 5].

Optimal arms: The arms for which $\psi_i^* = 0$ are assigned positive probabilities $x_i^* > 0$ by the solution to $LP(\boldsymbol{\mu})$. We can then obtain the optimal dual variables by solving the system of linear equations

$$r_i \mu_i + \lambda^* \mu_i - \nu^* = 0, \quad i \in \{j : \psi_j^* = 0, j \in \{1, \dots, N\}\},$$

which gives $\lambda^* = \frac{r_1 \mu_1 - r_2 \mu_2}{\mu_2 - \mu_1} \geq 0$ and $\nu^* = \frac{(r_1 - r_2) \mu_1 \mu_2}{\mu_2 - \mu_1}$. Note that λ^* is the slope of the hyperplane that optimizes $LP(\boldsymbol{\mu})$. Further, by solving $\sum_{i=1,2} p_i^* \mu_i = \eta$, we get the optimal selection probabilities $x_1^* = (\mu_2 - \eta) / (\mu_2 - \mu_1)$ and $x_2^* = (\eta - \mu_1) / (\mu_2 - \mu_1)$.

Suboptimal arms and complementary slackness: In contrast to the optimal arms, the arms where $x_i^* = 0$, $\psi_i^* > 0$ constitute suboptimal arms that are assigned zero probability mass by the optimal solution. For these arms, the constraints are *slack* so that sufficiently small perturbations of their parameters does not alter the optimal solution to $LP(\boldsymbol{\mu})$. We now quantify the slack for each suboptimal arm $i \notin \{1, 2\}$. For this, we observe that the arms stay suboptimal as long as $\psi_i^* > 0$. Rearranging (9) to obtain $r_i \mu_i + \lambda^* \mu_i - \nu^* = -\psi_i^* < 0$, we get $\mu_i < \frac{\nu^*}{r_i + \lambda^*}$. Hence, we define

$$\xi_i := \frac{\nu^*}{r_i + \lambda^*} = \frac{(r_1 - r_2) \mu_1 \mu_2}{(r_i - r_2) \mu_2 - (r_i - r_1) \mu_1} > \mu_i. \quad (13)$$

The value $\xi_i - \mu_i > 0$ is the *complementary slackness* for arm i . As long as any (perturbed) mean value, μ'_i , for arm i satisfies $\mu'_i < \xi_i$ the arm i stays suboptimal, i.e., the arm is assigned zero selection probability mass.

3.4 Arm Selection Bounds

To bound the regret and violation for LinConTS, we first bound $\mathbb{E}[k_i(T+1)]$, the expected number of times any suboptimal arm $i \notin \{1, 2\}$ is played until time $T+1$. Our proof technique is inspired by [Agrawal and Goyal, 2013]: First, we show that the probability of playing a sub-optimal arm is a linear function of the probability of playing one of the optimal arms. Next, we show that the coefficient of this linear function decreases exponentially with successive plays of the arm. For this, we use the concentration results from [Agrawal and Goyal, 2013] to upper bound the number of plays for each suboptimal arm. We begin by defining the following:

Definition 1: The number of successes observed for arm $i = 1, \dots, N$ until time step $t-1$ is denoted by $S_i(t)$. The empirical mean at time t is $\hat{\mu}_i(t) = \frac{S_i(t)}{k_i(t)}$, with $\hat{\mu}_i(t) = 1$ when $k_i(t) = 0$.

Definition 2: The history of plays until time $t-1$ are

denoted by the *filtration* \mathcal{F}_{t-1} , i.e.,

$$\mathcal{F}_{t-1} = \{i(w), c_{i(w),w}, w = 1, \dots, t-1\}.$$

Definition 3: For each suboptimal arm $i \notin \{1, 2\}$, we choose two thresholds y_i and z_i such that $\mu_i < y_i < z_i < \xi_i$, which are set at appropriate points of the proof. Further, we define $L_i(T) = \frac{\log T}{d(y_i, z_i)}$.

Definition 4: We define the variable, $\kappa_1, \dots, \kappa_N$ such that

$$\frac{\kappa_j r_j - z_i r_i}{\kappa_j - z_i} = \lambda^*,$$

where $\kappa_i := z_i$. The hyperplane supported by any two points $(z_i, z_i r_i)$ and $(\kappa_j, \kappa_j r_j)$, $j \neq i$ runs parallel to the optimal hyperplane for $LP(\boldsymbol{\mu})$.

Definition 5: We define the variable

$$\epsilon_{1,i} = \frac{\kappa_2 - \eta}{\kappa_2 - \kappa_1} > x_{1,t} > 0,$$

which lower bounds the selection probability mass assigned to arm 1, $x_{1,t}$, under certain conditions described later in the proof.

Definition 6: We define two events: $E_i^\mu(t) : \{\hat{\mu}_i(t) \leq y_i\}$ and $E_i^\theta(t) : \{\theta_{i,t} \leq z_i\}$, where $\theta_{i,t} := \tilde{\mu}_{i,t}$ denotes the Thompson sample at time t . Intuitively, $E_i^\mu(t)$ and $E_i^\theta(t)$ denote the events that the empirical and sampled means for arm i at time t do not exceed the true arm mean by a large amount.

Definition 7: We define the probability

$$p_{i,t} = \Pr(\{\theta_{1,t} > \kappa_1\} | \mathcal{F}_{t-1}) \quad (14)$$

Next, we prove the following lemma that establishes the relationship between the number of plays of any suboptimal arm $i \notin \{1, 2\}$ and the optimal arm $i = 1$.

Lemma 1. For all $t \in \{1, \dots, T\}$, and $i \notin \{1, 2\}$,

$$\begin{aligned} & \Pr(i(t) = i, E_i^\mu(t), E_i^\theta(t) | \mathcal{F}_{t-1}) \\ & \leq \frac{1}{\epsilon_{1,i}} \cdot \frac{1 - p_{i,t}}{p_{i,t}} \cdot \Pr(i(t) = 1, E_i^\mu(t), E_i^\theta(t) | \mathcal{F}_{t-1}). \end{aligned} \quad (15)$$

Proof. Similar to [Agrawal and Goyal, 2013, Lemma 1], we note that $E_i^\mu(t)$ is determined by \mathcal{F}_{t-1} . Further, when $E_i^\mu(t)$ is false, the left hand side of (15) is zero and the inequality is trivially satisfied. Hence, we assume that \mathcal{F}_{t-1} is such that $E_i^\mu(t)$ is true. Subsequently, to prove (15), we show that under a suitably chosen set of conditions $M_i(t)$ that hold with a non-zero probability, the following relationships also hold:

$$\begin{aligned} \Pr(i(t) = i | E_i^\theta(t), \mathcal{F}_{t-1}) & \leq (1 - p_{i,t}) \Pr(M_i(t) | E_i^\theta(t), \mathcal{F}_{t-1}), \\ \Pr(i(t) = 1 | E_i^\theta(t), \mathcal{F}_{t-1}) & \geq (\epsilon_{1,i} \cdot p_{i,t}) \Pr(M_i(t) | E_i^\theta(t), \mathcal{F}_{t-1}), \end{aligned}$$

which immediately gives the Lemma above. The proof details are provided in Appendix A. \square

Subsequently, to bound the expected number of plays for any suboptimal arm i , we use Lemma 2, 3, and 4 obtained in [Agrawal and Goyal, 2013],

$$\sum_{t=1}^T \Pr(i(t) = i, \overline{E_i^\mu(t)}) \leq \frac{1}{d(y_i, \mu_i)} + 1, \quad (16)$$

which relies on Chernoff-Hoeffding bounds on the concentration of the empirical mean around the true mean,

$$\sum_{t=1}^T \Pr(i(t) = i, E_i^\mu(t), \overline{E_i^\theta(t)}) \leq L_i(t) + 1, \quad (17)$$

which is based on the fact that after the arm has played a sufficient number of times, $L_i(t)$, the Thompson sample will be close to its mean value, and

$$\mathbb{E} \left[\frac{1}{p_{i, \tau_j + 1}} \right] \leq \begin{cases} 1 + \frac{3}{\Delta_i'} & j < \frac{8}{\Delta_i'}, \\ 1 + O\left(e^{-\frac{\Delta_i'^2 j}{2}} + \frac{1}{(j+1)\Delta_i'^2} e^{-D_i j} + \frac{1}{e^{\Delta_i'^2 j/4} - 1}\right) & j > \frac{8}{\Delta_i'} \end{cases}, \quad (18)$$

respectively, where τ_j denotes the time step for the j^{th} trial of arm 1, $\Delta_i' = \mu_1 - \kappa_1$ and $D_i = z_i \log \frac{z_i}{\mu_1} + (1 - z_i) \log \frac{1 - z_i}{1 - \mu_1}$. This last result uses novel algebraic analysis developed from the concentration of Binomial sums. Based on the above, we obtain the following bound on the expected number of plays of the suboptimal arm i until time T :

Lemma 2. *The expected number of plays for any stationary suboptimal arm, $i \in \{1, \dots, K\} \setminus \{1, 2\}$, is upper bounded by*

$$\begin{aligned} \mathbb{E}[k_i(T)] &\leq 2 + L_i(T) + \frac{1}{d(y_i, \xi_i)} + \frac{24}{\epsilon_{1,i} \Delta_i'^2} \\ &\quad + \frac{1}{\epsilon_{1,i}} \sum_{j=0}^{T-1} O\left(e^{-\frac{\Delta_i'^2 j}{2}} + \frac{1}{(j+1)\Delta_i'^2} e^{-D_i j} + \frac{1}{e^{\Delta_i'^2 j/4} - 1}\right). \end{aligned} \quad (19)$$

Proof. This proof essentially sums over the expressions obtained in (16), (17), and (18) until time T . The details are provided in Appendix B. \square

Lemma 3. *For the stationary optimal arms $i \in \{1, 2\}$,*

$$\begin{aligned} \left[\sum_{i=1,2} \Delta_i \mathbb{E}[k_i(T)] \right]_+ &\leq \Delta_{\max}^+ \cdot 18\sqrt{2T \log 2}, \\ \left[\sum_{i=1,2} (\eta - \mu_i) \mathbb{E}[k_i(T)] \right]_+ &\leq \delta_{\max}^+ \cdot 18\sqrt{2T \log 2}, \end{aligned} \quad (20)$$

where $\Delta_i = r^* - \mu_i r_i$, $\Delta_{\max}^+ = \max_{i \in [N]} \Delta_i^+$ and $\delta_{\max}^+ = \max_{i \in [N]} \delta_i^+$.

Proof. From Lemma 2, the rate of playing any suboptimal is a logarithmic function of T . Hence, the set of optimal arms is played at a linear rate. Viewing the play of optimal arms as a two-armed bandit, we can use existing results for constrained TS to bound the regret and violation contribution for the rounds when only the optimal arms are assigned nonzero arm selection probabilities by $LP(\tilde{\mu}_t)$. First, we decompose the regret and violation expressions separately using the decomposition in [Ferreira et al., 2018, EC.4] that relies on a frequentist upper and lower bounds that holds with high probability for every $t \in [T]$. Subsequently, we apply [Ferreira et al., 2018, Lemma EC.2], which provides an upper bound on the expected regret and violation with respect to the frequentist upper and lower bounds. This directly gives us the bounds in (20) where we set $K = 2$ since only the rounds where the optimal arms are played are taken into account. \square

3.5 Proof of Regret and Violation Bounds

For some $\gamma \in (0, 1]$ we choose the thresholds $y_i \in (\mu_i, \xi_i)$ and $z_i \in (y_i, \xi_i)$ such that $d(y_i, \xi_i) = d(\mu_i, \xi_i)/(1 + \gamma)$ and $d(y_i, z_i) = d(y_i, \xi_i)/(1 + \gamma) = d(\mu_i, \xi_i)/(1 + \gamma)^2$. This leads to

$$L_i(T) = \frac{\log T}{d(y_i, z_i)} = (1 + \gamma)^2 \frac{\log T}{d(\mu_i, \xi_i)}.$$

Following the ideas in [Agrawal and Goyal, 2013], we obtain $y_i - \mu_i \geq \frac{\gamma}{1 + \gamma} \cdot d(\mu_i, \xi) / \log \frac{\xi_i(1 - \mu_i)}{\mu_i(1 - \xi_i)}$, which gives $\frac{1}{d(y_i, \mu_i)} \leq \frac{2}{(y_i - \mu_i)^2} = O(1/\gamma^2)$.

Proof of regret bound: From Lemma 2 and Lemma 3, we obtain

$$\begin{aligned} \mathbb{E}[\mathcal{R}(T)] &= \left[\sum_i \Delta_i \mathbb{E}[k_i(T)] \right]_+ \\ &\leq \left[\sum_{i \neq \{1,2\}} \Delta_i \mathbb{E}[k_i(T)] \right]_+ + \left[\sum_{i=1,2} \Delta_i \mathbb{E}[k_i(T)] \right]_+ \\ &\leq \sum_{i \neq \{1,2\}} \Delta_i^+ \mathbb{E}[k_i(T)] + \Delta_{\max}^+ \cdot 18\sqrt{2T \log 2} \\ &\leq \sum_{i \neq \{1,2\}} (1 + \gamma)^2 \frac{\log T}{d(\mu_i, \xi_i)} \Delta_i^+ + O\left(\frac{N}{\gamma^2}\right) \\ &\quad + \Delta_{\max}^+ \cdot 18\sqrt{2T \log 2}, \end{aligned}$$

where the first and second inequalities use $[a + b]_+ \leq [a]_+ + [b]_+$. This completes the regret bound.

Proof of violation bound: From Lemma 2 and Lemma 3, we obtain

$$\begin{aligned} \mathbb{E}[\mathcal{V}(T)] &= \left[\sum_i (\eta - \mu_i) \mathbb{E}[k_i(T)] \right]_+ \\ &\leq \left[\sum_{i \neq \{1,2\}} \delta_i \mathbb{E}[k_i(T)] \right]_+ + \left[\sum_{i=\{1,2\}} \delta_i \mathbb{E}[k_i(T)] \right]_+ \\ &\leq \sum_{i \neq \{1,2\}} \frac{(1+\gamma)^2 \log T}{d(\mu_i, \xi_i)} \delta_i^+ + O\left(\frac{N}{\gamma^2}\right) + \delta_{\max}^+ \cdot 18\sqrt{2T \log 2} \\ &= \left[\sum_{i \neq \{1,2\}} \frac{(1+\gamma)^2}{d(\mu_i, \xi_i)} \delta_i^+ \right] \log T + O\left(\frac{N}{\gamma^2}\right) \\ &\quad + \delta_{\max}^+ \cdot 18\sqrt{2T \log 2}, \end{aligned}$$

where the first and second inequalities use $[a+b]_+ \leq [a]_+ + [b]_+$. This completes the violation bound.

4 Related Work

Sequential learning under probabilistic constraints has been studied in [Meisami et al., 2018], and Thompson sampling under general problem settings was studied in [Gopalan et al., 2014]. In [Agarwal et al., 2011], distribution-free regret bounds for convex optimization with bandit feedback were shown to scale as $O(\sqrt{T})$. Bandits with concave rewards and convex knapsacks were studied within a very general framework in [Agrawal and Devanur, 2014], which subsumed the previous bandits with knapsack (BwK) formulation in [Badanidiyuru et al., 2013]. In [Xia et al., 2015], a Thompson sampling algorithm for budgeted MABs was proposed that outperforms the UCB BwK algorithm. Subsequently TS was applied to network revenue management in [Ferreira et al., 2018], where distribution-free bounds that scale as $O(\sqrt{T})$ were shown to hold. In [Chen et al., 2018], a horizon-independent UCB approach was proposed whose regret and violation performance was shown to scale with $O(\sqrt{T})$. Recently, applications of linearly constrained MABs have been proposed for advertiser portfolio optimization [Pani et al., 2017], wireless communications [Saxena et al., 2019, Saxena and Jaldén, 2020], and real-time electricity pricing [Tucker et al., 2019].

5 Numerical Experiments

We evaluate the regret and violation performance of LinConTS on two real-world datasets, *Coupon-Purchase* [Kaggle, 2016] and *edX-Course* [Chuang and Ho, 2016], respectively. The Coupon-Purchase and edX-Course datasets have been explored previously for constrained bandit problems, albeit in a multi-play setting that allowed multiple arms to be played in each time

step and the assumed that the reward values were stochastic [Chen et al., 2018, Cai et al., 2017]. The experiments are implemented in Python using Jupyter notebooks and have been made publicly available at [Saxena, 2019].

For the Coupon-Purchase dataset, which contains discount coupons applied to online purchases, we extract all the $N = 142$ coupons for products priced equal to or below 200 price units that have been purchased by at least one customer. For these coupons, we obtain the purchase rate and the discounted selling price from the dataset. The edX-Course dataset contains enrolment information for $N = 290$ Harvard and MIT courses. We process this dataset according to previous experiments [Chen et al., 2018]: the course participation rates are estimated by max-min normalization of the number of participants in each course, and the course certification rates are obtained by dividing the number of certified participants in each course by the number of course participants.

We model each coupon in the Coupon-Purchase dataset and each course in the edX-Course dataset using independent bandit arms. Further, for the Coupon-Purchase dataset, we generate independent Bernoulli-distributed reward events with mean values obtained from the purchase rates and a deterministic reward value defined as the final selling price normalized by 200. For the processed Coupon-Purchase dataset, the reward event means and the rewards values are found to lie between $[0, 0.30]$ and $(0, 1]$ respectively. Analogously for the edX-Course dataset, the reward event means are obtained from the course participation rates and the corresponding reward values are the course certification rates that are assumed to be known in advance, e.g. from historical course data. For this dataset, the reward event means and the rewards values are found to span $[0, 1]$ and $(0, 0.40]$ respectively.

We implement two linearly constrained bandit algorithms, LinConTS and LinCon-KL-UCB respectively. The LinConTS algorithm is described in Sec. 3 and the pseudocode is available in Alg. 1. The LinCon-KL-UCB, described in Appendix C, is inspired by the ConTS algorithm proposed for multi-play linearly constrained MABs in [Chen et al., 2018]. However, compared to ConTS that relies on an index-based UCB, LinCon-KL-UCB estimates the UCB for each arm using the Kullback-Liebler (KL) divergence metric. For the Bernoulli-distributed rewards considered in here, KL-based UCB has been shown to achieve optimal regret and significantly outperform index-based UCB schemes [Garivier and Cappé, 2011].

The performance of LinCon-KL-UCB and LinConTS for the Coupon-Purchase dataset and for $\eta = 0.25$

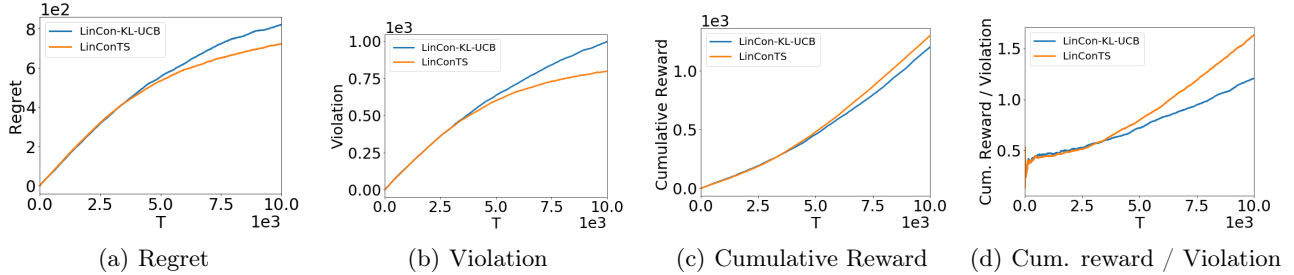


Figure 1: Experimental results for the Coupon-Purchase dataset for $N = 142$ coupons with $\eta = 0.25$. The LinConTS approach outperforms the competing LinCon-KL-UCB approach by achieving a lower regret and violation, higher cumulative rewards and a higher ratio of cumulative rewards to violation.




Figure 2: Experimental results for the edX-Course dataset for $N = 290$ courses with $\eta = 0.50$. The LinConTS approach outperforms the competing LinCon-KL-UCB approach by achieving a lower regret and violation, higher cumulative rewards and a higher ratio of cumulative rewards to violation.

is shown in Fig. 1, where the results have been averaged over 16 test runs. LinConTS achieves a lower regret compared to LinCon-KL-UCB in Fig. 1(a). This demonstrates that compared to LinCon-KL-UCB, LinConTS is closer to the cumulative rewards achieved by the stationary optimal policy. We study the cumulative violations for each approach in Fig. 1(b), where LinConTS demonstrates fewer constraint violations than LinCon-KL-UCB. Interestingly, in Fig. 1(c), LinConTS achieves higher cumulative rewards as well compared to LinCon-KL-UCB. We combine the related effects of regret and violation minimization by calculating the ratio between the cumulative rewards and the cumulative violations at each time step. This quantity, depicted in Fig. 1(d), can be interpreted as the additional reward earned for every constraint violation. We observe that LinConTS achieves a higher ratio, which demonstrated that LinConTS is more efficient in exploiting the infrequent constraint violations.

For the edX-Course dataset and $\eta = 0.50$, the experimental results for LinConTS and LinCon-KL-UCB schemes are averaged over 16 test runs and shown in Fig. 2. Similar to the results for the Coupon-Purchase dataset, here also LinConTS has lower regret and violation than LinCon-KL-UCB in Fig. 2(a) and Fig. 2(b) respectively. Also, from Fig. 2(c) and Fig. 2(d), Lin-

ConTS increases the cumulative reward, and the ratio of cumulative reward to violations, respectively.

6 Conclusions and Further Work

For constrained bandit problems, LP subroutines enable powerful sequential arm selection techniques. Combined with an underlying Thompson Sampling approach, these techniques promise efficient and robust solutions that are optimal in terms of their learning rate. We have addressed the MAB problem of maximizing the cumulative reward when the reward event probability is constrained above a fixed threshold in every round. For this problem, we described LinConTS, which incorporates a LP subroutine in every step of the Thompson Sampling heuristic. We have provided the first instance-dependent asymptotic analysis of our algorithm. Through numerical results for two real-world datasets, we have showed that LinConTS outperforms an optimal UCB-based algorithm in terms of the regret and violation metrics. The proof technique developed in this paper can inspire solutions to other constrained MAB problems, for example to problems that deal with more general type of constraints compared to the one considered here.

Acknowledgements

This work was partially supported by the Wallenberg Artificial Intelligence, Autonomous Systems and Software Program (WASP) funded by Knut and Alice Wallenberg Foundation, and by the European Research Council project AGNOSTIC (742648). We thank Simon Lindstahl for noting a flaw in the original formulation of Lemma 3, which has been updated in this version of the paper.

References

- [Agarwal et al., 2011] Agarwal, A., Foster, D. P., Hsu, D. J., Kakade, S. M., and Rakhlin, A. (2011). Stochastic convex optimization with bandit feedback. In *Advances in Neural Information Processing Systems*, pages 1035–1043.
- [Agrawal and Devanur, 2014] Agrawal, S. and Devanur, N. R. (2014). Bandits with concave rewards and convex knapsacks. In *Proceedings of the Fifteenth ACM Conference on Economics and Computation*, EC '14, pages 989–1006, New York, NY, USA. ACM.
- [Agrawal and Goyal, 2012] Agrawal, S. and Goyal, N. (2012). Analysis of thompson sampling for the multi-armed bandit problem. In *Conference on Learning Theory*, pages 39–1.
- [Agrawal and Goyal, 2013] Agrawal, S. and Goyal, N. (2013). Further optimal regret bounds for thompson sampling. In *Artificial Intelligence and Statistics*, pages 99–107.
- [Badanidiyuru et al., 2013] Badanidiyuru, A., Kleinberg, R., and Slivkins, A. (2013). Bandits with knapsacks. In *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on*, pages 207–216. IEEE.
- [Boyd and Vandenberghe, 2004] Boyd, S. and Vandenberghe, L. (2004). *Convex Optimization*. Cambridge university press.
- [Cai et al., 2017] Cai, K., Chen, K., Huang, L., and Lui, J. C. (2017). Multi-level feedback web links selection problem: Learning and optimization. In *2017 IEEE International Conference on Data Mining (ICDM)*, pages 763–768. IEEE.
- [Chapelle and Li, 2011] Chapelle, O. and Li, L. (2011). An empirical evaluation of thompson sampling. In *Advances in neural information processing systems*, pages 2249–2257.
- [Chen et al., 2018] Chen, K., Cai, K., Huang, L., and Lui, J. C. (2018). Beyond the click-through rate: Web link selection with multi-level feedback. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*, pages 3308–3314. International Joint Conferences on Artificial Intelligence Organization.
- [Chuang and Ho, 2016] Chuang, I. and Ho, A. (2016). HarvardX and MITx: Four years of open online courses—fall 2012–summer 2016. *Available at SSRN 2889436*.
- [Ferreira et al., 2018] Ferreira, K. J., Simchi-Levi, D., and Wang, H. (2018). Online network revenue management using thompson sampling. *Operations research*, 66(6):1586–1602.
- [Flajolet and Jaillet, 2015] Flajolet, A. and Jaillet, P. (2015). Logarithmic regret bounds for bandits with knapsacks. *arXiv preprint arXiv:1510.01800*.
- [Garivier and Cappé, 2011] Garivier, A. and Cappé, O. (2011). The kl-ucb algorithm for bounded stochastic bandits and beyond. In *Proceedings of the 24th annual Conference On Learning Theory*, pages 359–376.
- [Gopalan et al., 2014] Gopalan, A., Mannor, S., and Mansour, Y. (2014). Thompson sampling for complex online problems. In *International Conference on Machine Learning*, pages 100–108.
- [Kaggle, 2016] Kaggle (2016). Coupon Purchase Data. <https://www.kaggle.com/c/coupon-purchase-prediction/data>.
- [Kaufmann et al., 2012] Kaufmann, E., Korda, N., and Munos, R. (2012). Thompson sampling: An asymptotically optimal finite-time analysis. In *International Conference on Algorithmic Learning Theory*, pages 199–213. Springer.
- [Lai and Robbins, 1985] Lai, T. L. and Robbins, H. (1985). Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics*, 6(1):4–22.
- [Meisami et al., 2018] Meisami, A., Lam, H., Dong, C., and Pani, A. (2018). Sequential learning under probabilistic constraints. In Globerson, A. and Silva, R., editors, *Proceedings of the Thirty-Fourth Conference on Uncertainty in Artificial Intelligence, UAI 2018, Monterey, California, USA, August 6-10, 2018*, pages 621–631. AUAI Press.
- [Pani et al., 2017] Pani, A., Raghavan, S., and Sahin, M. (2017). Large-scale advertising portfolio optimization in online marketing. Technical report, Working Paper.
- [Russo et al., 2018] Russo, D. J., Van Roy, B., Kazerouni, A., Osband, I., and Wen, Z. (2018). A tutorial on thompson sampling. *Found. Trends Mach. Learn.*, 11(1):1–96.

- [Saxena, 2019] Saxena, V. (2019). LinConTS. <https://github.com/vidits-kth/LinConTS>.
- [Saxena and Jaldén, 2020] Saxena, V. and Jaldén, J. (2020). Bayesian link adaptation under a bler target. In *IEEE 21st International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*.
- [Saxena et al., 2019] Saxena, V., Jaldén, J., Gonzalez, J. E., Stoica, I., and Tullberg, H. (2019). Constrained thompson sampling for wireless link optimization. *arXiv preprint arXiv:1902.11102*.
- [Thompson, 1933] Thompson, W. R. (1933). On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285–294.
- [Tucker et al., 2019] Tucker, N., Moradipari, A., and Alizadeh, M. (2019). Constrained thompson sampling for real-time electricity pricing with grid reliability constraints. *arXiv preprint arXiv:1908.07964*.
- [Xia et al., 2015] Xia, Y., Li, H., Qin, T., Yu, N., and Liu, T.-Y. (2015). Thompson sampling for budgeted multi-armed bandits. In *Proceedings of the 24th International Conference on Artificial Intelligence, IJCAI'15*, page 3960–3966. AAAI Press.