
Analysis and Design of Thompson Sampling for Stochastic Partial Monitoring

Taira Tsuchiya
The University of Tokyo
RIKEN AIP
tsuchiya@ms.k.u-tokyo.ac.jp

Junya Honda
The University of Tokyo
RIKEN AIP
honda@edu.k.u-tokyo.ac.jp

Masashi Sugiyama
RIKEN AIP
The University of Tokyo
sugi@k.u-tokyo.ac.jp

Abstract

We investigate finite stochastic partial monitoring, which is a general model for sequential learning with limited feedback. While Thompson sampling is one of the most promising algorithms on a variety of online decision-making problems, its properties for stochastic partial monitoring have not been theoretically investigated, and the existing algorithm relies on a heuristic approximation of the posterior distribution. To mitigate these problems, we present a novel Thompson-sampling-based algorithm, which enables us to exactly sample the target parameter from the posterior distribution. Besides, we prove that the new algorithm achieves the logarithmic *problem-dependent expected pseudo-regret* $O(\log T)$ for a linearized variant of the problem with local observability. This result is the first regret bound of Thompson sampling for partial monitoring, which also becomes the first logarithmic regret bound of Thompson sampling for linear bandits.

1 Introduction

Partial monitoring (PM) is a general sequential decision-making problem with limited feedback (Rus-tichini, 1999; Piccolboni and Schindelhauer, 2001). PM is attracting broad interest because it includes a wide range of problems such as the multi-armed bandit problem (Lai and Robbins, 1985), a linear optimization problem with full or bandit feedback (Zinkevich, 2003; Dani et al., 2008), dynamic pricing (Kleinberg and Leighton, 2003), and label efficient prediction (Cesa-Bianchi et al., 2005).

A PM game can be seen as a sequential game that is played by two players: a learner and an opponent. At every round, the learner chooses an action, while the opponent chooses an outcome. Then, the learner suffers an unobserved loss and receives a feedback symbol, both of which are determined from the selected action and outcome. The main characteristic of this game is that the learner cannot directly observe the outcome and loss. The goal of the learner is to minimize his/her cumulative loss over all rounds. The performance of the learner is evaluated by the regret, which is defined as the difference between the cumulative losses of the learner and the optimal action (*i.e.*, the action whose expected loss is the smallest).

There are mainly two types of PM games, which are the *stochastic* and *adversarial* settings (Piccolboni and Schindelhauer, 2001; Bartók et al., 2011). In the stochastic setting, the outcome at each round is determined from the *opponent's strategy*, which is a probability vector over the opponent's possible choices. On the other hand, in the adversarial setting, the outcomes are arbitrarily decided by the

opponent. We refer to the PM game with finite actions and finite outcomes as a *finite* PM game. In this paper, we focus on the stochastic finite game.

One of the first algorithms for PM was considered by [Piccolboni and Schindelhauer \(2001\)](#). They proposed the FeedExp3 algorithm, the key idea of which is to use an unbiased estimator of the losses. They showed that the FeedExp3 algorithm attains $\tilde{O}(T^{3/4})$ minimax regret for a certain class of PM games, and showed that any algorithm suffers linear minimax regret $\Omega(T)$ for the other class. Here T is the time horizon and the notation $\tilde{O}(\cdot)$ hides polylogarithmic factors. The upper bound $\tilde{O}(T^{3/4})$ is later improved by [Cesa-Bianchi et al. \(2006\)](#) to $O(T^{2/3})$, and they also provided a game with a matching lower bound.

In the seminal paper by [Bartók et al. \(2011\)](#), they classified PM games into four classes based on their minimax regrets. To be more specific, they classified games into trivial, easy, hard, and hopeless games, where their minimax regrets are 0 , $\Theta(\sqrt{T})$, $\Theta(T^{2/3})$, and $\Theta(T)$, respectively. Note that the easy game is also called a *locally observable* game. After their work, several algorithms have been proposed for the finite PM problem ([Bartók et al., 2012](#); [Vanchinathan et al., 2014](#); [Komiyama et al., 2015](#)). For the problem-dependent regret analysis, [Komiyama et al. \(2015\)](#) proposed an algorithm that achieves $O(\log T)$ regret with the optimal constant factor. However, it requires to solve a time-consuming optimization problem with infinitely many constraints at each round. In addition, this algorithm relies on the forced exploration to achieve the optimality, which makes the empirical performance near-optimal only after prohibitively many rounds, say, 10^5 or 10^6 .

Thompson sampling (TS, [Thompson, 1933](#)) is one of the most promising algorithms on a variety of online decision-making problems such as the multi-armed bandit ([Lai and Robbins, 1985](#)) and the linear bandit ([Agrawal and Goyal, 2013b](#)), and the effectiveness of TS has been investigated both empirically ([Chapelle and Li, 2011](#)) and theoretically ([Kaufmann et al., 2012](#); [Agrawal and Goyal, 2013a](#); [Honda and Takemura, 2014](#)). In the literature on PM, [Vanchinathan et al. \(2014\)](#) proposed a TS-based algorithm called BPM-TS (Bayes-update for PM based on TS) for stochastic PM, which empirically achieved state-of-the-art performance. Their algorithm uses Gaussian approximation to handle the complicated posterior distribution of the opponent’s strategy. However, this approximation is somewhat heuristic and can degrade the empirical performance due to the discrepancy from the exact posterior distribution. Furthermore, no theoretical guarantee is provided for BPM-TS.

Our goals are to establish a new TS-based algorithm for stochastic PM, which allows us to sample the opponent’s strategy parameter from the exact posterior distribution, and investigate whether the TS-based algorithm can achieve sub-linear regret in stochastic PM. Using the accept-reject sampling, we propose a new TS-based algorithm for PM (TSPM), which is equipped with a numerical scheme to obtain a posterior sample from the complicated posterior distribution. We derive a logarithmic regret upper bound $O(\log T)$ for the proposed algorithm on the locally observable game under a linearized variant of the problem. This is the first regret bound for TS on the locally observable game. Moreover, our setting includes the linear bandit problem and our result is also the first logarithmic expected regret bound of TS for the linear bandit, whereas a high-probability bound was provided, for example, in [Agrawal and Goyal \(2013b\)](#). Finally, we compare the performance of TSPM with existing algorithms in numerical experiments, and show that TSPM outperforms existing algorithms.

2 Preliminaries

This paper studies finite stochastic PM games ([Bartók et al., 2011](#)). A PM game with N actions and M outcomes is defined by a pair of a loss matrix $\mathbf{L} = (\ell_{i,j}) \in \mathbb{R}^{N \times M}$ and feedback matrix $\mathbf{H} = (h_{i,j}) \in [A]^{N \times M}$, where A is the number of feedback symbols and $[A] = \{1, 2, \dots, A\}$.

A PM game can be seen as a sequential game that is played by two players: the learner and the opponent. At each round $t = 1, 2, \dots, T$, the learner selects action $i(t) \in [N]$, and at the same time the opponent selects an outcome based on the opponent’s strategy $p^* \in \mathcal{P}_M$, where $\mathcal{P}_n = \{p \in \mathbb{R}^n : p_k \geq 0, \sum_{k=1}^n p_k = 1\}$ is the $(n - 1)$ -dimensional probability simplex. The outcome $j(t)$ of each round is an independent and identically distributed sample from p^* , and then, the learner suffers loss $\ell_{i(t),j(t)}$ at time t . The learner cannot directly observe the value of this loss, but instead observes the *feedback symbol* $y(t) = h_{i(t),j(t)} \in [A]$. The setting explained above has been widely studied in the literature of stochastic PM ([Bartók et al., 2011](#); [Komiyama et al., 2015](#)), and we call this the *discrete* setting. In Section 4, we also introduce a *linear* setting for theoretical analysis, which is a slightly different setting from the discrete one.

The learner aims to minimize the cumulative loss over T rounds. The expected loss of action i is given by $L_i^\top p^*$, where L_i is the i -th column of \mathbf{L}^\top . We say action i is *optimal* under strategy p^* if $(L_i - L_j)^\top p^* \leq 0$ for any $j \neq i$. We assume that the optimal action is unique, and without loss of generality that the optimal action is action 1. Let $\Delta_i = (L_i - L_1)^\top p^* \geq 0$ for $i \in [N]$ and $N_i(t)$ be the number of times action i is selected before the t -th round. When the time step is clear from the context, we use n_i instead of $N_i(t)$. We adopt the pseudo-regret to measure the performance: $\text{Reg}(T) = \sum_{t=1}^T \Delta_{i(t)} = \sum_{i \in [N]} \Delta_i N_i(T+1)$. This is the relative performance of the algorithm against the *oracle*, which knows the optimal action 1 before the game starts.

We introduce the following definitions to clarify the class of PM games, for which we develop an algorithm and derive a regret upper bound. The following cell decomposition is the concept to divide the simplex \mathcal{P}_M based on the loss matrix to identify the optimal action, which depends on the opponent's strategy p^* .

Definition 1 (Cell decomposition and Pareto-optimality (Bartók et al., 2011)). For every action $i \in [N]$, cell $\mathcal{C}_i := \{p \in \mathcal{P}_M : (L_i - L_j)^\top p \leq 0, \forall j \neq i\}$ is the set of opponent's strategies for which action i is optimal. Action i is *Pareto-optimal* if there exists an opponent's strategy p^* under which action i is optimal.

Each cell is a convex closed polytope. Next, we define *neighbors* between two Pareto-optimal actions, which intuitively means that the two actions "touch" each other in their surfaces.

Definition 2 (Neighbors and neighborhood action (Bartók et al., 2011)). Two Pareto-optimal actions i and j are *neighbors* if $\mathcal{C}_i \cap \mathcal{C}_j$ is an $(M-2)$ -dimensional polytope. For two neighboring actions $i, j \in [N]$, the *neighborhood action set* is defined as $N_{i,j}^+ = \{k \in [N] : \mathcal{C}_i \cap \mathcal{C}_j \subseteq \mathcal{C}_k\}$.

Note that the neighborhood action set $N_{i,j}^+$ includes actions i and j from its definition. Next, we define the *signal matrix*, which encodes the information of the feedback matrix \mathbf{H} so that we can utilize the feedback information.

Definition 3 (Signal matrix (Komiyama et al., 2015)). The signal matrix $S_i \in \{0, 1\}^{A \times M}$ of action i is defined as $(S_i)_{y,j} = \mathbb{1}[h_{i,j} = y]$, where $\mathbb{1}[X] = 1$ if the event X is true and 0 otherwise.

Note that if we define the signal matrix as above, $S_i p^* \in \mathbb{R}^A$ is a probability vector over feedback symbols of action i . The following *local observability* condition separates easy and hard games, this condition intuitively means that the information obtained by taking actions in the neighborhood action set $N_{i,j}^+$ is sufficient to distinguish the loss difference between actions i and j .

Definition 4 (Local observability (Bartók et al., 2011)). A partial monitoring game is said to be *locally observable* if for all pairs i, j of neighboring actions, $L_i - L_j \in \bigoplus_{k \in N_{i,j}^+} \text{Im } S_k^\top$, where $\text{Im } V$ is the image of the linear map V , and $V \oplus W$ is the direct sum between the vector spaces V and W .

We also consider the concept of the *strong local observability* condition, which implies the above local observability condition.

Definition 5 (Strong local observability). A partial monitoring game is said to be *strongly locally observable* if for all pairs $i, j \in [N]$, $L_i - L_j \in \text{Im } S_i^\top \oplus \text{Im } S_j^\top$.

This condition was assumed in the theoretical analysis in Vanchinathan et al. (2014), and we also assume this condition in theoretical analysis in Section 4. Note that the strong local observability means that, for any $j \neq k$, there exists $z_{j,k} \neq 0 \in \mathbb{R}^{2A}$ such that $L_j - L_k = (S_j^\top, S_k^\top) z_{j,k}$.

Notation. Let $\|\cdot\|$ and $\|\cdot\|_p$ be the Euclidian norm and p -norm, and let $\|x\|_A = \sqrt{x^\top A x}$ be the norm induced by the positive semidefinite matrix $A \succeq 0$. Let $\mathcal{D}_{\text{KL}}(p||q) = \sum_{a=1}^A p_a \log(p_a/q_a)$ be the Kullback-Leibler divergence of p from q . The vector $e_y \in \mathbb{R}^M$ is the y -th orthonormal basis of \mathbb{R}^M , and $\mathbf{1}_n = [1, \dots, 1]^\top$ is the n -dimensional all-one vector. Let $q_i^{(t)}$ be the empirical feedback distribution of action i at time t , i.e., $q_i^{(t)} = [n_{i1}/n_i, \dots, n_{iA}/n_i]^\top \in \mathcal{P}_A$, where $n_{iy} = \sum_{s=1}^t \mathbb{1}[i(s) = i, y(s) = y]$ and $n_i = \sum_{y=1}^A n_{iy}$. The notation is summarized in Appendix A.

Methods for Sampling from Posterior Distribution. We briefly review the methods to draw a sample from the posterior distribution. While TS is one of the most promising algorithms, the posterior distribution can be in a quite complicated form, which makes obtaining a sample from it

Algorithm 1: TSPM Algorithm

Input: prior parameter $\lambda > 0$

- 1 Set $B_0 \leftarrow \lambda I_M, b_0 \leftarrow 0$.
 - 2 Take each action for $n \geq 1$ times.
 - 3 **for** $t = 1, 2, \dots, T$ **do**
 - 4 Sample $\tilde{p}_t \sim \pi(p \mid \{i(s), y(s)\}_{s=1}^t)$ based on the accept-reject sampling (Algorithm 2).
 - 5 Take action $i(t) = \arg \max_{i \in [N]} L_i^\top \tilde{p}_t$ and observe feedback $y(t)$.
 - 6 Update $B_t \leftarrow B_{t-1} + S_{i(t)}^\top S_{i(t)}, b_t \leftarrow b_{t-1} + S_{i(t)}^\top e_{y(t)}$.
-

Algorithm 2: Accept-Reject Sampling

Input: constant $R \in [0, 1]$

- 1 **while** true **do**
 - 2 Sample $\tilde{p}_t \sim g_t(p)$ (Algorithm 3).
 - 3 Sample $\tilde{u} \sim \mathcal{U}([0, 1])$.
 - 4 **if** $R\tilde{u} < F_t(\tilde{p}_t)/G_t(\tilde{p}_t)$ **then**
 - 5 **return** \tilde{p}_t .
-

Algorithm 3: Sampling from $g_t(p)$

- 1 Compute \tilde{B}_t, \tilde{b}_t from B_t, b_t .
 - 2 **repeat**
 - 3 | Sample $p^{(\alpha)} \sim \mathcal{N}(\tilde{B}_t^{-1} \tilde{b}_t, \tilde{B}_t^{-1})$.
 - 4 **until** $p^{(\alpha)} \in \mathcal{P}_{M-1}$;
 - 5 **return** $\tilde{p} = [p^{(\alpha)\top}, 1 - \sum_{i=1}^{M-1} (p^{(\alpha)})_i]^\top$.
-

computationally hard. To overcome this issue, a variety of approximate posterior sampling methods have been considered, such as Gibbs sampling, Langevin Monte Carlo, Laplace approximation, and the bootstrap (Russo et al., 2018, Section 5). Recent work (Lu and Van Roy, 2017) proposed a flexible approximation method, which can even efficiently be applied to quite complex models such as neural networks. However, more recent work revealed that algorithms based on such an approximation procedure *can* suffer a linear regret (Phan et al., 2019), even if the approximation error in terms of the α -divergence is small enough.

Although BPM-TS is one of the best methods for stochastic PM, it approximates the posterior by a Gaussian distribution in a heuristic way, which can degrade the empirical performance due to the distributional discrepancy from the exact posterior distribution. Furthermore, no theoretical guarantee is provided for BPM-TS. In this paper, we mitigate these problems by providing a new algorithm for stochastic PM, which allows us to exactly draw samples from the posterior distribution. We also give theoretical analysis for the proposed algorithm.

3 Thompson-sampling-based Algorithm for Partial Monitoring

In this section, we present a new algorithm for stochastic PM games, where we name the algorithm TSPM (TS-based algorithm for PM). The algorithm is given in Algorithm 1, and we will explain the subroutines in the following.

3.1 Accept-Reject Sampling

We adopt the accept-reject sampling (Casella et al., 2004) to *exactly* draw samples from the posterior distribution. The accept-reject sampling is a technique to draw samples from a specific distribution f , and a key feature is to use a *proposal distribution* g , from which we can easily draw a sample and whose ratio to f , that is f/g , is bounded by a constant value R . To obtain samples from f , (i) we generate samples $X \sim g$; (ii) accept X with probability $f(X)/Rg(X)$. Note that f and g do not have to be normalized when the acceptance probability is calculated.

Let $\pi(p)$ be a prior distribution for p . Then an unnormalized density of the posterior distribution for p can be expressed as

$$F_t(p) = \pi(p) \prod_{i=1}^N \exp\left\{-n_i \mathcal{D}_{\text{KL}}\left(q_i^{(t)} \parallel S_i p\right)\right\}, \quad (1)$$

the detailed derivation of which is given in Appendix B. We use the proposal distribution with unnormalized density

$$G_t(p) = \pi(p) \prod_{i=1}^N \exp\left\{-\frac{1}{2}n_i\|q_i^{(t)} - S_i p\|^2\right\}. \quad (2)$$

Based on these distributions, we use Algorithm 2 for exact sampling from the posterior distribution, where $\mathcal{U}([0, 1])$ is the uniform distribution over $[0, 1]$ and $g_t(p)$ is the distribution corresponding to the unnormalized density $G_t(p)$ in (2). The following proposition shows that setting $R = 1$ realizes the exact sampling.

Proposition 1. *Let $f_t(p)$ be the distribution corresponding to the unnormalized density $F_t(p)$ in (1). Then, the output of Algorithm 2 with $R = 1$ follows $f_t(p)$.*

This proposition can easily be proved by Pinsker’s inequality, which is detailed in Appendix B.

In practice, $R \in [0, 1]$ is a parameter to balance the amount of over-exploration and the computational efficiency. As R decreases from 1, the algorithm tends to accept a point p far from the mode. The case $R = 0$ corresponds the TSPM algorithm where the proposal distribution is used without the accept-reject sampling, which we call *TSPM-Gaussian*. As we will see in Section 4, TSPM-Gaussian corresponds to exact sampling of the posterior distribution when the feedback follows a Gaussian distribution rather than a multinomial distribution.

TSPM-Gaussian can be related to BPM-TS (Vanchinathan et al., 2014) in the sense that both of them use samples from Gaussian distributions. Nevertheless, they use different Gaussians and TSPM-Gaussian performs much better than BPM-TS as we will see in the experiments. Details on the relation between TSPM-Gaussian and BPM-TS are described in Appendix D.

In general, we can realize efficient sampling with a small number of rejections if the proposal distribution and the target distribution are close to each other. On the other hand, in our problem, the densities in (1) and (2) for each fixed point p exponentially decay with the number of samples n_i if the empirical feedback distribution $q_i^{(t)}$ converges. This means that $F_t(p)$ and $G_t(p)$ have an exponentially large relative gap in most rounds. Nevertheless, the number of rejections does not increase with t as we will see in the experiments, which suggests that the proposal distribution approximates the target distribution well with high probability.

3.2 Sampling from Proposal Distribution

When we consider Gaussian density $\mathcal{N}(0, \lambda I_M)$ truncated over \mathcal{P}_M as a prior, the proposal distribution also has the Gaussian density $\mathcal{N}(B_t^{-1}b_t, B_t^{-1})$ over \mathcal{P}_M , where

$$B_t = \lambda I_M + \sum_{i=1}^N n_i S_i^\top S_i = B_{t-1} + S_{i(t)}^\top S_{i(t)}, \quad b_t = \sum_{i=1}^N n_i S_i^\top q_i^{(t)} = b_{t-1} + S_{i(t)}^\top e_{y(t)}. \quad (3)$$

Here note that the probability simplex \mathcal{P}_M is in an $(M - 1)$ -dimensional space and a sample from $\mathcal{N}(0, \lambda I_M)$ is not contained in \mathcal{P}_M with probability one. In the literature, *e.g.*, Altmann et al. (2014), sampling methods for Gaussian distributions truncated on a simplex have been discussed. We use one of these procedures summarized in Algorithm 3, where we first sample $M - 1$ elements of p from another Gaussian distribution and determine the remaining element by the constraint $\sum_{i=1}^M p_i = 1$.

Proposition 2. *Sampling from $g_t(p)$ is equivalent to Algorithm 3 with*

$$\tilde{B}_t = C_t - 2D_t + f_t \mathbf{1}_{M-1} \mathbf{1}_{M-1}^\top, \quad \tilde{b}_t = f_t \mathbf{1}_{M-1} - d_t + b_t^{(\alpha)} - b^{(M)} \mathbf{1}_{M-1},$$

where $B_t = \begin{bmatrix} C_t & d_t \\ d_t^\top & f_t \end{bmatrix}$ for $C_t \in \mathbb{R}^{(M-1) \times (M-1)}$, $d_t \in \mathbb{R}^{M-1}$, $f_t \in \mathbb{R}$, $b_t = [b_t^{(\alpha)\top}, b_t^{(M)\top}]^\top \in \mathbb{R}^{M-1} \times \mathbb{R}$, and $D_t = \frac{1}{2}(d_t \mathbf{1}_{M-1}^\top + \mathbf{1}_{M-1} d_t^\top)$.

We give the proof of this proposition for self-containedness in Appendix C.

4 Theoretical Analysis

This section considers a regret upper bound of the TSPM algorithm.

In the theoretical analysis, we consider a *linear* setting of PM. In the linear PM, the learner suffers the expected loss $L_{i(t)}^\top p^*$ as in the discrete setting, and receives feedback vector $y(t) = S_i p^* + \epsilon_t$ for $\epsilon_t \sim \mathcal{N}(0, I_M)$ whereas the one-hot representation of $y(t)$ is distributed by the probability vector $S_i p^*$ in the discrete setting. Therefore, if ϵ_t can be regarded as a sub-Gaussian random variable as in [Kirschner et al. \(2020\)](#) then the linear PM includes the discrete PM, though our theoretical analysis requires ϵ_t to be Gaussian. The relation between discrete and linear settings can also be seen from the observation that bandit problems with Bernoulli and Gaussian rewards can be expressed as discrete and linear PM, respectively. The linear PM also includes the linear bandit problem, where the feedback vector is expressed as $L_i^\top p^* + \epsilon_t$.

In the linear PM, $G_t(p)$ in (2) becomes the exact posterior distribution rather than a proposal distribution. The definition of the cell decomposition for this setting is largely the same as that of discrete setting and detailed in Appendix F. Therefore, TS with exact posterior sampling in the linear PM corresponds to TSPM-Gaussian. In the linear PM, the unknown parameter p^* is in \mathbb{R}^M rather than in \mathcal{P}_M , and therefore we consider the prior $\pi(p) = \mathcal{N}(0, \lambda I_M)$ over \mathbb{R}^M , where the posterior distribution becomes $\mathcal{N}(B_t^{-1} b_t, B_t^{-1})$.

There are a few works that analyze TS for the PM because of its difficulty. For example in [Vanchinathan et al. \(2014\)](#), an analysis of the TS-based algorithm (BPM-TS) is not given despite the fact that its performance is better than the algorithm based on a confidence ellipsoid (BPM-LEAST). [Zimmert and Lattimore \(2019\)](#) considered the theoretical aspect of a variant of TS for the linear PM in view of the Bayes regret, but this algorithm is based on the knowledge on the time horizon and different from the family of TS used in practice. More specifically, their algorithm considers the posterior distribution for *regret* (not pseudo-regret), and an action is chosen according to the posterior probability that each arm minimizes the *cumulative* regret. Thus, the time horizon also needs to be known.

Types of Regret Bounds. We focus on the (a) *problem-dependent* (b) *expected pseudo-regret*. (a) In the literature, a *minimax* (or *problem-independent*) regret bound has mainly been considered, for example, to classify difficulties of the PM problem ([Bartók et al., 2010](#); [Bartók et al., 2011](#)). On the other hand, a *problem-dependent* regret bound often reflects the empirical performance more clearly than the minimax regret ([Bartók et al., 2012](#); [Vanchinathan et al., 2014](#); [Komiyama et al., 2015](#)). For this reason, we consider this problem-dependent regret bound. (b) In complicated settings of bandit problems, a *high-probability regret bound* has mainly been considered ([Abbasi-Yadkori et al., 2011](#); [Agrawal and Goyal, 2013b](#)), which bounds the pseudo-regret with high probability $1 - \delta$. Though such a bound can be transformed to an expected regret bound, this type of analysis often sacrifices the tightness since a linear regret might be suffered with small probability δ . This is why the analysis in [Vanchinathan et al. \(2014\)](#) for BPM-LEAST finally yielded an $\tilde{O}(\sqrt{T})$ expected regret bound whereas their high-probability bound is $O(\log T)$.

4.1 Regret Upper Bound

In the following theorem, we show that logarithmic problem-dependent expected regret is achievable by the TSPM-Gaussian algorithm.

Theorem 3 (Regret upper bound). *Consider any finite stochastic linear partial monitoring game. Assume that the game is strongly locally observable and $\Delta_i = (L_i - L_1)^\top p^* > 0$ for any $i \neq 1$. Then, the regret of TSPM-Gaussian satisfies for sufficiently large T that*

$$\mathbb{E}[\text{Reg}(T)] = O\left(\frac{AN^2 M \max_{i \in [N]} \Delta_i}{\Lambda^2} \log T\right), \quad (4)$$

where $\Lambda := \min_{i \neq 1} \Lambda_i$ for $\Lambda_i = \Delta_i / \|z_{1,i}\|$ with $z_{1,i}$ defined after Definition 5.

Remark. In the proof of Theorem 3, it is sufficient to assume that $L_1 - L_i \in \text{Im } S_1^\top \oplus \text{Im } S_i^\top$ for $i \in [N]$, which is weaker than the strong local observability, though it is still sometimes stronger than the local observability condition.

The proof of Theorem 3 is given in Appendix F. This result is the first problem-dependent bound of TS for PM, which also becomes the first logarithmic regret bound of TS for linear bandits.

The norm of $z_{j,k}$ in Λ intuitively indicates the difficulty of the problem. Whereas we can estimate $(S_j p, S_k p)$ with noise through taking actions j and k , the actual interest is the gap of the losses $p^\top (L_j - L_k) = (S_j p, S_k p)^\top z_{j,k}$. Thus, if $\|z_{j,k}\|$ is large, the gap estimation becomes difficult since the noise is enhanced through $z_{j,k}$.

Unfortunately, the derived bound in Theorem 3 has quadratic dependence on N , which seems to be not tight. This quadratic dependence comes from the difficulty of the *expected* regret analysis. In general, we evaluate the regret before and after the convergence of the statistics separately. Whereas the latter one usually becomes dominant, the main difficulty comes from the analysis of the former one, which might become large with low probability (Agrawal and Goyal, 2012; Kaufmann et al., 2012; Agrawal and Goyal, 2013a).

In our analysis, we were not able to bound the former one within a non-dominant order, though it is still logarithmic in T . In fact, our analysis shows that the regret after convergence is $O(\sum_{i \neq 1} \Delta_i \frac{A}{\Lambda^2} \log T)$ as shown in Lemma 18 in Appendix F, which will become the regret with high probability. In particular, if we consider the classic bandit problem as a PM game, we can confirm that the derived bound after convergence becomes the best possible bound

$$O\left(\sum_{i \neq 1} \frac{\log T}{\Delta_i}\right)$$

by considering Λ_i depending on each suboptimal arm i as the difficulty measure instead of Λ . Still, deriving a regret bound for the term before convergence within an non-dominant order is an important future work.

4.2 Technical Difficulties of the Analysis

The main difficulty of this regret analysis is that PM requires to consider the statistics of *all* actions when the number of selections $N_i(t)$ of some action i is evaluated. This is in stark contrast to the analysis of the classic bandit problems, where it becomes sufficient to evaluate statistics of action i and the best action 1. This makes the analysis remarkably complicated in TS, where we need to separately consider the randomness caused by the feedback and TS.

To overcome this difficulty, we handle the effect of actions of no interest in two different novel ways depending on each decomposed regret. The first one is to evaluate the worst-case effect of these actions based on an argument (Lemma 10) related to the law of the iterated logarithm (LIL), which is sometimes used in the best-arm identification literature to improve the performance (Jamieson et al., 2014). The second one is to bound the action-selection probability of TS using an argument of (super-)martingale (Theorem 16), which is of independent interest. Whereas such a technique is often used for the construction of confidence bounds (Abbasi-Yadkori et al., 2011), we reveal that it is also useful for evaluation of the regret of TS.

We only focused on the Gaussian noise $\epsilon_t \sim \mathcal{N}(0, I_M)$, rather than the more general sub-Gaussian noise. This restriction to the Gaussian noise comes from the essential difficulty of the problem-dependent analysis of TS, where lower bounds for some probabilities are needed whereas the sub-Gaussian assumption is suited for obtaining upper bounds. To the best of our knowledge, the problem-dependent regret analysis for TS on the sub-Gaussian case has never been investigated even for the multi-armed bandit setting, which is quite simple compared to that of PM. In the literature of the problem-dependent regret analysis, the noise distribution is restricted to distributions with explicitly given forms, e.g., Bernoulli, Gaussian, or more generally a one-dimensional canonical exponential family (Kaufmann et al., 2012; Agrawal and Goyal, 2013a; Korda et al., 2013). Their analysis relies on the specific characteristic of the distribution to bound the problem-dependent regret.

5 Experiments

In this section, we numerically compare the performance of TSPM and TSPM-Gaussian against existing methods, which are RandomPM (the algorithm which selects action randomly), FeedExp3 (Piccolboni and Schindelhauer, 2001), and BPM-TS (Vanchinathan et al., 2014). Recently, Lattimore and Szepesvári (2019) considered the sampling-based algorithm called Mario sampling for easy games. Mario sampling coincides with TS (except for the difference between pseudo-regret and regret with known time horizon) mentioned in the last section when any pair of actions is a neighbor. As shown in Appendix G, this property is indeed satisfied for dp-easy games defined in the following. Therefore, the performance is essentially the same between TSPM with $R = 1$ and Mario sampling. To compare the performance, we consider a dynamic pricing problem, which is a typical example of PM games. We conducted experiments on the discrete setting because the experiments for PM has been mainly focused on the discrete setting.




Figure 1: Regret-round plots of algorithms. The solid lines indicate the average over 100 independent trials. The thin fillings are the standard error.

In the dynamic pricing game, the player corresponds to a seller, and the opponent corresponds to a buyer. At each round, the seller sells an item for a specific price $i(t)$, and the buyer comes with an evaluation price $j(t)$ for the item, where the selling price and the evaluation price correspond to the action and outcome, respectively. The buyer buys the item if the selling price $i(t)$ is smaller than or equal to $j(t)$ and not otherwise. The seller can only know if the buyer bought the item (denoted as feedback 0) or did not buy the item (denoted as 1). The seller aims to minimize the cumulative “loss”, and there are two types of definitions for the loss, where each induced game falls into the easy and hard games. We call them *dp-easy* and *dp-hard* games, respectively.

In both cases, the seller incurs the constant loss $c > 0$ when the item is not bought due to the loss of opportunity to sell the item. In contrast, when the item is not bought, the loss incurred to the seller is different between these settings. The seller in the *dp-easy* game *does not* take the buyer’s evaluation price into account. In other words, the seller gains the selling price $i(t)$ as a reward (equivalently incurs $-i(t)$ as a loss). Therefore, the loss for the selling price $i(t)$ and the evaluation $j(t)$ is

$$\ell_{i(t),j(t)} = -i(t)\mathbb{1}[i(t) \leq j(t)] + c\mathbb{1}[i(t) > j(t)].$$

This setting can be regarded as a generalized version of the online posted price mechanism, which was addressed in, e.g., [Blum et al. \(2004\)](#) and [Cesa-Bianchi et al. \(2006\)](#), and an example of strongly locally observable games.

On the other hand, the seller in *dp-hard* game *does* take the buyer’s evaluation price into account when the item is bought. In other words, the seller incurs the difference between the opponent evaluation and the selling price $j(t) - i(t)$ as a loss because the seller could have made more profit if the seller had sold at the price $j(t)$. Therefore, the loss incurred at time t is

$$\ell_{i(t),j(t)} = (j(t) - i(t))\mathbb{1}[i(t) \leq j(t)] + c\mathbb{1}[i(t) > j(t)].$$

This setting is also addressed in [Cesa-Bianchi et al. \(2006\)](#), and belongs to the class of hard games. Note that our algorithm can also be applied to a hard game, though there is no theoretical guarantee.

Setup. In the both *dp-easy* and *dp-hard* games, we fixed $N = M \in \{3, 5, 7\}$ and $c = 2$. We fixed the time horizon T to 10000 and simulated 100 times. For FeedExp3 and BPM-TS, the setup of hyperparameters follows their original papers. For TSPM, we set $\lambda = 0.001$, and R was selected from $\{0.01, 1.0\}$. Here, recall that TSPM with $R = 1$ and $R = 0$ correspond to the exact sampling and TSPM-Gaussian, respectively, and a smaller value of R gives the higher acceptance probability in the accept-reject sampling. Therefore, using small R makes the algorithm time-efficient, although it can worsen the performance since it over-explores the tail of the posterior distributions. To stabilize




Figure 2: The number of rejected times by the accept-reject sampling. The solid lines indicate the average over 100 independent trials after taking moving average with window size 100.

sampling from the proposal distribution in Algorithm 3, we used an initialization that takes each action $n = 10A$ times. The detailed settings of the experiments with more results are given in Appendix H.

Results. Figure 1 is the empirical comparison of the proposed algorithms against the benchmark methods. This result shows that, in all cases, the TSPM with exact sampling gives the best performance. TSPM-Gaussian also outperforms BPM-TS even though both of them use Gaussian distributions as posteriors. Besides, the experimental results suggest that our algorithm performs reasonably well even for a hard game. It can be observed that the proposed methods outperform BPM-TS more significantly for a larger number of outcomes. Further discussion for this observation is given in Appendix D.

Figure 2 shows the number of rejections at each time step in the accept-reject sampling. We counted the number of times that either Line 4 in Algorithm 2 or Line 4 in Algorithm 3 was not satisfied. In the accept-reject sampling, it is desirable that the frequency of rejection does not increase as the time-step t and does not increase rapidly with the number of outcomes. We can see that the former one is indeed satisfied. For the latter property, the frequency of rejection becomes unfortunately large when exact sampling ($R = 1$) is conducted. Still, we can substantially improve this frequency by setting R to be a small value or zero, which still keeps regret tremendously better than that of BPM with almost the same time-efficiency as BPM-TS.

6 Conclusion and Discussion

This paper investigated Thompson sampling (TS) for stochastic partial monitoring from the algorithmic and theoretical viewpoints. We provided a new algorithm that enables exact sampling from the posterior distribution, and numerically showed that the proposed algorithm outperforms existing methods. Besides, we provided an upper bound for the problem-dependent logarithmic expected pseudo-regret for the linearized version of the partial monitoring. To our knowledge, this bound is the first logarithmic problem-dependent expected pseudo-regret bound of a TS-based algorithm for linear bandit problems and strongly locally observable partial monitoring games.

There are several remaining questions. As mentioned in Section 4, Kirschner et al. (2020) considered linear partial monitoring with the feedback structure $y(t) = S_{i(t)}p^* + \epsilon_t$, where $(\epsilon_t)_{t=1}^T$ is a sequence of independent sub-Gaussian noise vector in \mathbb{R}^M . This setting is the generalization of our linear setting, where $(\epsilon_t)_{t=1}^T$ are i.i.d. Gaussian vectors. Therefore, a natural question that arises is whether we can extend our analysis on TSPM-Gaussian to the sub-Gaussian case, although we believe it would be not straightforward as discussed in Section 4. It is also an important open problem to derive a regret bound on TSPM using the exact posterior sampling for the discrete partial monitoring. Although we conjecture that the algorithm also achieves logarithmic regret for the setting, there still remain some difficulties in the analysis. In particular, we have to handle the KL divergence in $f_t(p)$ and consider the restriction of the support of the opponent’s strategy to \mathcal{P}_M , which make the analysis much more complicated. Besides, it is worth noting that the theoretical analysis of TS for hard games has never been theoretically investigated. We believe that in general TS suffers linear regret in the minimax sense due to its greediness. However, we conjecture that TS can achieve the sub-linear regret for some specific instances of hard games in the sense of the problem-dependent regret, as empirically observed in the experiments. Finally, it is an important open problem to derive the minimax regret for anytime TS-based algorithms. This needs more detailed analysis on $o(\log T)$ terms in the regret bound, which were dropped in our main result.

Broader Impact

Application. Partial monitoring (PM) includes various online decision-making problems such as multi-armed bandits, linear bandits, dynamic pricing, and label efficient prediction. Not only can PM handle them, the dueling bandits, combinatorial bandits, transductive bandits, and many other problems can be seen as a partial monitoring game, as discussed in [Kirschner et al. \(2020\)](#). Therefore, our analysis of Thompson sampling (TS) for PM games pushes the application of TS to a more wide range of online decision-making problems forward. Moreover, PM has the potential that novel online-decision making problems are newly discovered, where we have to handle the limited feedback in an online fashion.

Practical Use. The obvious advantage of using TS is that the users can easily apply the algorithm to their problems. They do not have to solve mathematical optimization problems, which are often required to solve when using non-sampling-based algorithms ([Bartók et al., 2012](#); [Komiyama et al., 2015](#)). For the negative side, the theoretical analysis for the regret upper bound might make the users become overconfident when the users use their algorithms. For example, they might use the TSPM algorithm to the linear PM game with heavy-tailed noise, such as sub-exponential noise, without noticing it. Nevertheless, this is not an TS-specific problem, but one that can be found in many theoretical studies, and TS is still one of the most promising policies.

Acknowledgements

The authors would like to thank the meta-reviewer and reviewers for a lot of helpful comments. The authors would like to thank Kento Nozawa and Ikko Yamane for maintaining servers for our experiments, and Kenny Song for helpful discussion on the writing. TT was supported by Toyota-Dwango AI Scholarship, and RIKEN Junior Research Associate Program for the final part of the project. JH was supported by KAKENHI 18K17998, and MS was supported by KAKENHI 17H00757.

References

- Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. In *Advances in Neural Information Processing Systems 24*, pages 2312–2320, 2011.
- Shipra Agrawal and Navin Goyal. Analysis of Thompson sampling for the multi-armed bandit problem. In *the 25th Annual Conference on Learning Theory*, volume 23, pages 39.1–39.26, 2012.
- Shipra Agrawal and Navin Goyal. Further optimal regret bounds for Thompson sampling. In *the Sixteenth International Conference on Artificial Intelligence and Statistics*, volume 31, pages 99–107, 2013a.
- Shipra Agrawal and Navin Goyal. Thompson sampling for contextual bandits with linear payoffs. In *the 30th International Conference on Machine Learning*, pages 127–135, 2013b.
- Yoann Altmann, Steve McLaughlin, and Nicolas Dobigeon. Sampling from a multivariate gaussian distribution truncated on a simplex: A review. In *2014 IEEE Workshop on Statistical Signal Processing (SSP)*, pages 113–116, 2014.
- Gábor Bartók, Dávid Pál, and Csaba Szepesvári. Toward a classification of finite partial-monitoring games. In *Algorithmic Learning Theory*, pages 224–238, 2010.
- Gábor Bartók, Navid Zolghadr, and Csaba Szepesvári. An adaptive algorithm for finite stochastic partial monitoring. In *the 29th International Conference on Machine Learning*, pages 1–20, 2012.
- Gábor Bartók, Dávid Pál, and Csaba Szepesvári. Minimax regret of finite partial-monitoring games in stochastic environments. In *the 24th Annual Conference on Learning Theory*, volume 19, pages 133–154, 2011.
- Avrim Blum, Vijay Kumar, Atri Rudra, and Felix Wu. Online learning in online auctions. *Theoretical Computer Science*, 324(2):137–146, 2004.
- George Casella, Christian P. Robert, and Martin T. Wells. *Generalized accept-reject sampling schemes*, volume 45 of *Lecture Notes–Monograph Series*, pages 342–347. Institute of Mathematical Statistics, 2004.
- Nicolò Cesa-Bianchi, Gábor Lugosi, and Gilles Stoltz. Minimizing regret with label efficient prediction. *IEEE Transactions on Information Theory*, 51(6):2152–2162, 2005.

- Nicolò Cesa-Bianchi, Gábor Lugosi, and Gilles Stoltz. Regret minimization under partial monitoring. *Mathematics of Operations Research*, 31(3):562–580, 2006.
- Olivier Chapelle and Lihong Li. An empirical evaluation of Thompson sampling. In *Advances in Neural Information Processing Systems 24*, pages 2249–2257, 2011.
- Varsha Dani, Thomas P. Hayes, and Sham M. Kakade. Stochastic linear optimization under bandit feedback. In *21st Annual Conference on Learning Theory*, pages 355–366, 2008.
- Junya Honda and Akimichi Takemura. Optimality of Thompson sampling for Gaussian bandits depends on priors. In *the Seventeenth International Conference on Artificial Intelligence and Statistics*, volume 33, pages 375–383, 2014.
- Kevin Jamieson, Matthew Malloy, Robert Nowak, and Sébastien Bubeck. lin^* : An optimal exploration algorithm for multi-armed bandits. In *The 27th Conference on Learning Theory*, pages 423–439, 2014.
- Emilie Kaufmann, Nathaniel Korda, and Rémi Munos. Thompson sampling: An asymptotically optimal finite-time analysis. In *Algorithmic Learning Theory*, pages 199–213, 2012.
- Johannes Kirschner, Tor Lattimore, and Andreas Krause. Information directed sampling for linear partial monitoring. *arXiv preprint arXiv:2002.11182*, 2020.
- Robert Kleinberg and Tom Leighton. The value of knowing a demand curve: Bounds on regret for online posted-price auctions. In *the 44th Annual IEEE Symposium on Foundations of Computer Science*, pages 594–605, 2003.
- Junpei Komiyama, Junya Honda, and Hiroshi Nakagawa. Regret lower bound and optimal algorithm in finite stochastic partial monitoring. In *Advances in Neural Information Processing Systems 28*, pages 1792–1800, 2015.
- Nathaniel Korda, Emilie Kaufmann, and Remi Munos. Thompson sampling for 1-dimensional exponential family bandits. In *Advances in Neural Information Processing Systems 26*, pages 1448–1456, 2013.
- T. L. Lai and Herbert Robbins. Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics*, 6(1):4–22, 1985.
- Tor Lattimore and Csaba Szepesvári. An information-theoretic approach to minimax regret in partial monitoring. In *the 32nd Annual Conference on Learning Theory*, volume 99, pages 2111–2139, 2019.
- Xiuyuan Lu and Benjamin Van Roy. Ensemble sampling. In *Advances in Neural Information Processing Systems 30*, pages 3258–3266, 2017.
- My Phan, Yasin Abbasi-Yadkori, and Justin Domke. Thompson sampling and approximate inference. In *Advances in Neural Information Processing Systems 32*, pages 8804–8813, 2019.
- Antonio Piccolboni and Christian Schindelhauer. Discrete prediction games with arbitrary feedback and loss. In *COLT/EuroCOLT*, pages 208–223, 2001.
- Daniel J. 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.
- Aldo Rustichini. Minimizing regret: The general case. *Games and Economic Behavior*, 29(1): 224–243, 1999.
- 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, 12 1933.
- Hastagir P Vanchinathan, Gábor Bartók, and Andreas Krause. Efficient partial monitoring with prior information. In *Advances in Neural Information Processing Systems 27*, pages 1691–1699, 2014.
- Julian Zimmert and Tor Lattimore. Connections between mirror descent, thompson sampling and the information ratio. In *Advances in Neural Information Processing Systems 32*, pages 11973–11982, 2019.
- Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *the Twentieth International Conference on Machine Learning*, pages 928–935. AAAI Press, 2003.

A Notation

Table 1 summarizes the symbols used in this paper.

Table 1: List of symbols used in this paper.

| Symbol | Meaning |
|---|---|
| \mathcal{P}_n | $(n - 1)$ -dimensional probability simplex |
| $\ \cdot\ $ | Euclidian norm for vector and operator norm for matrix |
| $\ \cdot\ _p$ | p -norm |
| $\ \cdot\ _A$ | norm induced by positive semidefinite matrix A |
| $\mathcal{D}_{\text{KL}}(p\ q)$ | KL divergence from q to p |
| $B_r^n(p)$ | n -dimensional Euclidian ball of radius r at point $p \in \mathbb{R}^N$ |
| $N, M \in \mathbb{N}$ | the number of actions and outcomes |
| Σ | set of feedback symbols |
| A | the number of feedback symbols |
| $p^* \in \mathcal{P}_M$ | opponent's strategy |
| T | time horizon |
| $\mathbf{L} = (\ell_{i,j}) \in \mathbb{R}^{N \times M}$ | loss matrix |
| $\mathbf{H} = (h_{i,j}) \in \Sigma^{N \times M}$ | feedback matrix |
| $S_i \in \{0, 1\}^{A \times M}$ ($i = 1, \dots, N$) | signal matrix |
| $i(t)$ | action taken at time t |
| $N_i(t)$ | the number of times the action i is taken before time $t \in [T]$ |
| $j(t)$ | outcome taken by opponent at time t |
| $y(t)$ | feedback observed at time t |
| $F_t(p)$ | unnormalized posterior distribution in (1) |
| $f_t(p)$ | probability density function corresponding to $F_t(p)$ |
| $G_t(p)$ | unnormalized proposal distribution for $F_t(p)$ in (2) |
| $g_t(p)$ | probability density function corresponding to $G_t(p)$ |
| $q_i^{(t)} \in \mathcal{P}_M$ | empirical feedback distribution of action i by time t |
| $q_{i,n} \in \mathcal{P}_M$ | empirical feedback distribution of action i after the action is taken n times |
| $\mathcal{C}_i \subset \mathcal{P}_M$ | cell of action i |

B Posterior Distribution and Proposal Distribution in Section 3

In this appendix, we discuss representation of the posterior distribution and its relation with the proposal distribution.

Proposition 4. $F_t(p)$ in (1) is proportional to the posterior distribution of the opponent's strategy, and $F_t(p) \leq G_t(p)$ for all $p \in \mathcal{P}_M$.

Proof. The posterior distribution of the opponent's strategy parameter $\pi(p \mid \{i(s), y(s)\}_{s=1}^t)$ is rewritten as

$$\begin{aligned}
 \pi(p \mid \{i(s), y(s)\}_{s=1}^t) &\propto \pi(p, \{i(s), y(s)\}_{s=1}^t) \\
 &\propto \pi(p) \prod_{s=1}^t \mathbb{P}\{y(s) \mid i(s), p\} \\
 &= \pi(p) \prod_{i=1}^N \prod_{y=1}^A (S_{i,y} p)^{n_{iy}} \\
 &\propto \pi(p) \prod_{i=1}^N \exp\left\{-n_i \mathcal{D}_{\text{KL}}\left(q_i^{(t)} \parallel S_i p\right)\right\}, \tag{5}
 \end{aligned}$$

where $S_{i,y}$ is the i -th row of the signal matrix S_i , and note that $q_i^{(t)}$ is the empirical feedback distribution of action i at time t , that is, $q_i^{(t)} = [n_{i1}/n_i, \dots, n_{iA}/n_i]^\top \in \mathcal{P}_A$ for $n_{iy} = \sum_{s=1}^t \mathbb{1}[i(s) = i, y(s) = y]$ and $n_i = \sum_{y=1}^A n_{iy}$.

Next, we show that $F_t(p) \leq G_t(p)$ holds for all $p \in \mathcal{P}_M$. Using the Pinsker's inequality, the unnormalized posterior distribution $F_t(p)$ can be bounded from above as

$$\begin{aligned}
F_t(p) &= \pi(p) \prod_{i=1}^N \exp\left\{-n_i \mathcal{D}_{\text{KL}}\left(q_i^{(t)} \parallel S_i p\right)\right\} \\
&\leq \pi(p) \prod_{i=1}^N \exp\left\{-\frac{1}{2} n_i \|q_i^{(t)} - S_i p\|_1^2\right\} \quad (\text{by Pinsker's ineq.}) \\
&= \pi(p) \exp\left\{-\frac{1}{2} \sum_{i=1}^N n_i \|q_i^{(t)} - S_i p\|_1^2\right\} \\
&\leq \pi(p) \exp\left\{-\frac{1}{2} \sum_{i=1}^N n_i \|q_i^{(t)} - S_i p\|^2\right\} \quad (\text{by } \|q_i^{(t)} - S_i p\|_1 \geq \|q_i^{(t)} - S_i p\|) \\
&= G_t(p). \tag{6}
\end{aligned}$$

□

Remark. The unnormalized density $G_t(p)$ is indeed Gaussian. Recalling that B_t and b_t are defined in (3) as

$$B_t = \sum_{i=1}^N n_i S_i^\top S_i = \sum_{s=1}^t S_{i(s)}^\top S_{i(s)} = B_{t-1} + S_{i(t)}^\top S_{i(t)}, \quad b_t = \sum_{i=1}^N n_i S_i^\top q_i^{(t)} = b_{t-1} + S_{i(t)}^\top e_{y(t)}, \tag{7}$$

we have

$$\begin{aligned}
\sum_{i=1}^N n_i \|q_i^{(t)} - S_i p\|^2 &= \sum_{i=1}^N n_i (q_i^{(t)} - S_i p)^\top (q_i^{(t)} - S_i p) \\
&= p^\top \underbrace{\left(\sum_{i=1}^N n_i S_i^\top S_i\right)}_{B_t} p - 2 \underbrace{\left(\sum_{i=1}^N n_i S_i^\top q_i^{(t)}\right)^\top}_{b_t} p + \underbrace{\sum_{i=1}^N n_i \|q_i^{(t)}\|^2}_{c_t} \\
&= p^\top B_t p - 2 b_t^\top p + c_t \\
&= (p - B_t^{-1} b_t)^\top B_t (p - B_t^{-1} b_t) + c_t - b_t^\top B_t^{-1} b_t. \tag{8}
\end{aligned}$$

Therefore, we have

$$\exp\left\{-\frac{1}{2} \sum_{i=1}^N n_i \|q_i^{(t)} - S_i p\|^2\right\} \propto \exp\left\{-\frac{1}{2} (p - B_t^{-1} b_t)^\top B_t (p - B_t^{-1} b_t)\right\}. \tag{9}$$

C Proof of Proposition 2

We will see that the the procedure sampling \tilde{p}_t from $g_t(p)$ and Algorithm 3 are equivalent. First, we derive the Gaussian density of $g_t(p)$ projected onto $\{p \in \mathbb{R}^M : \sum_{i=1}^M p_i = 1\}$.

For simplicity, we omit the subscript t and write, e.g., B instead of B_t . We define $p = [p^{(\alpha)\top}, p_M]^\top \in \mathbb{R}^{M-1} \times \mathbb{R}$. Let $h = B^{-1}b$, and define $h = [h^{(\alpha)\top}, h_M]^\top \in \mathbb{R}^{M-1} \times \mathbb{R}$. Let $B = \begin{bmatrix} C & d \\ d^\top & f \end{bmatrix}$, where $C \in \mathbb{R}^{(M-1) \times (M-1)}$, $d \in \mathbb{R}^{M-1}$, and $f \in \mathbb{R}$. Also, let $b = [b^{(\alpha)\top}, b^{(M)}]^\top \in \mathbb{R}^{M-1} \times \mathbb{R}$.

Using the decomposition

$$(p - B^{-1}b)^\top B (p - B^{-1}b) = \underbrace{p^\top B p}_{(a)} - 2 \underbrace{h^\top B p}_{(b)} + h^\top B h, \tag{10}$$

we rewrite each term by restricting the domain of p so that it satisfies the condition $\sum_{i=1}^M p_i = 1$. Now the first term (a) is rewritten as

$$\begin{aligned} \text{(a)} &= p^{(\alpha)\top} C p^{(\alpha)} + 2p^{(\alpha)\top} d p_M + f p_M^2 \\ &= p^{(\alpha)\top} C p^{(\alpha)} + \underbrace{2p^{(\alpha)\top} d \left(1 - \sum_{i=1}^{M-1} p_i\right)}_{\text{(a1)}} + \underbrace{f \left(1 - \sum_{i=1}^{M-1} p_i\right)^2}_{\text{(a2)}}. \end{aligned} \quad (11)$$

The term (a1) is rewritten as

$$\begin{aligned} \text{(a1)} &= p^{(\alpha)\top} d - p^{(\alpha)\top} d \sum_{i=1}^{M-1} p_i \\ &= p^{(\alpha)\top} d - p^{(\alpha)\top} d \mathbf{1}_{M-1}^\top p^{(\alpha)} \\ &= p^{(\alpha)\top} d - p^{(\alpha)\top} D p^{(\alpha)} \quad \left(D = \frac{1}{2} (d \mathbf{1}_{M-1}^\top + \mathbf{1}_{M-1} d^\top) \right), \end{aligned} \quad (12)$$

and the term (a2) is rewritten as

$$\begin{aligned} \text{(a2)} &= \left(1 - \sum_{i=1}^{M-1} p_i\right)^2 \\ &= 1 - 2 \sum_{i=1}^{M-1} p_i + \left(\sum_{i=1}^{M-1} p_i\right)^2 \\ &= 1 - 2 \mathbf{1}_{M-1}^\top p^{(\alpha)} + p^{(\alpha)\top} \mathbf{1}_{M-1} \mathbf{1}_{M-1}^\top p^{(\alpha)}. \end{aligned} \quad (13)$$

Therefore,

$$\text{(a)} = p^{(\alpha)\top} \underbrace{(C - 2D + f \mathbf{1}_{M-1} \mathbf{1}_{M-1}^\top)}_{\tilde{B}} p^{(\alpha)} - 2(f \mathbf{1}_{M-1} - d)^\top p^{(\alpha)} + f. \quad (14)$$

With regard to the term (b), we have

$$\begin{aligned} \text{(b)} &= b^\top p \\ &= b^{(\alpha)\top} p^{(\alpha)\top} + b^{(M)} p_M \\ &= (b^{(\alpha)} - b^{(M)} \mathbf{1}_{M-1})^\top p^{(\alpha)} + b^{(M)}. \end{aligned} \quad (15)$$

Therefore,

$$\begin{aligned} &(p - B^{-1}b)^\top B (p - B^{-1}b) \\ &= p^{(\alpha)\top} \tilde{B} p^{(\alpha)} - 2 \underbrace{(f \mathbf{1}_{M-1} - d + b^{(\alpha)} - b^{(M)} \mathbf{1}_{M-1})^\top}_{\tilde{b}} p^{(\alpha)} + f - 2b^{(M)} + h^\top B h \\ &= (p^{(\alpha)} - \tilde{B}^{-1} \tilde{b})^\top \tilde{B} (p^{(\alpha)} - \tilde{B}^{-1} \tilde{b}) + f - 2b^{(M)} - \tilde{b}^\top \tilde{B}^{-1} \tilde{b} + b^\top B^{-1} b \quad (\text{by } h^\top B h = b^\top B^{-1} b). \end{aligned} \quad (16)$$

From the above argument, the density $\mathcal{N}(\tilde{B}^{-1} \tilde{b}, \tilde{B}^{-1})$ is the Gaussian distribution of $g_t(p)$ on $\{p \in \mathbb{R}^M : \sum_{i=1}^M p_i = 1\}$. Therefore, the $p = [p^{(\alpha)\top}, 1 - \sum_{i=1}^{M-1} (p^{(\alpha)})_i]^\top$ for $p^{(\alpha)} \sim \mathcal{N}(\tilde{B}^{-1} \tilde{b}, \tilde{B}^{-1})$ is supported over $\{p \in \mathbb{R}^M : \sum_{i=1}^M p_i = 1\}$.

If the sample $p^{(\alpha)}$ from $\mathcal{N}(\tilde{B}^{-1} \tilde{b}, \tilde{B}^{-1})$ is in \mathcal{P}_{M-1} , then we can obtain the last element $p^{(M)}$ by $p^{(M)} = 1 - \sum_{i=1}^{M-1} (p^{(\alpha)})_i$. Otherwise, the probability that $p^{(\alpha)}$ is the first $M-1$ elements of the sample from $g_t(p)$ is zero, and hence, $[p^{(\alpha)\top}, p^{(M)}]^\top$ cannot be a sample from $g_t(p)$. Therefore, sampling \tilde{p}_t from $g_t(p)$ and Algorithm 3 are equivalent.

D Relation between TSPM-Gaussian and BPM-TS

In this appendix, we discuss the relation between TSPM-Gaussian and BPM-TS (Vanchinathan et al., 2014).

Underlying Feedback Structure. Here, we discuss the underlying feedback structure behind TSPM-Gaussian and BPM-TS.

We first consider the underlying feedback structure behind BPM-TS. In the following, we see that the feedback structure

$$y(t) = S_{i(t)}p + S_{i(t)}\epsilon, \quad \epsilon \sim \mathcal{N}(0, I_M) \quad (17)$$

induces the posterior distribution in BPM-TS. Under this feedback structure, we have $y(t) \sim \mathcal{N}(S_{i(t)}p, S_{i(t)}S_{i(t)}^\top)$.

When we take the prior distribution $\pi(p)$ as $\mathcal{N}(0, \sigma_0^2 I_M)$, the posterior distribution for the opponent's strategy parameter can be written as

$$\begin{aligned} & \pi\left(p \mid \{i(s), y(s)\}_{s=1}^t\right) \\ & \propto \pi(p) \prod_{s=1}^t \pi(y(s) \mid i(s), p) \\ & = \pi(p) \prod_{s=1}^t \mathbb{P}_{y \sim \mathcal{N}(S_{i(s)}p, S_{i(s)}S_{i(s)}^\top)} \{y = y(s)\} \\ & = \exp\left(-\frac{p^\top p}{2\sigma_0^2}\right) \prod_{s=1}^t \exp\left(-\frac{1}{2}(y(s) - S_{i(s)}p)^\top (S_{i(s)}S_{i(s)}^\top)^{-1}(y(s) - S_{i(s)}p)\right) \\ & = \exp\left\{-\frac{1}{2}\left(p^\top \left(\frac{1}{\sigma_0^2}I_M + \sum_{s=1}^t S_{i(s)}^\top (S_{i(s)}S_{i(s)}^\top)^{-1} S_{i(s)}\right)p\right.\right. \\ & \quad \left.\left.- 2\left(\sum_{s=1}^t y(s)^\top (S_{i(s)}S_{i(s)}^\top)^{-1} S_{i(s)}p\right) + (\text{a term independent of } p)\right\} \\ & \propto \exp\left\{-\frac{1}{2}(p^\top B_t^{\text{BPM}}p - 2b_t^{\text{BPM}\top}p)\right\} \\ & \propto \exp\left\{-\frac{1}{2}(p - B_t^{\text{BPM}^{-1}}b_t^{\text{BPM}})^\top B_t^{\text{BPM}}(p - B_t^{\text{BPM}^{-1}}b_t^{\text{BPM}})\right\}, \end{aligned} \quad (18)$$

where

$$B_t^{\text{BPM}} = \frac{1}{\sigma_0^2}I_M + \sum_{s=1}^t S_{i(s)}^\top (S_{i(s)}S_{i(s)}^\top)^{-1} S_{i(s)} = B_{t-1}^{\text{BPM}} + S_{i(t)}^\top (S_{i(t)}S_{i(t)}^\top)^{-1} S_{i(t)}, \quad (19)$$

$$b_t^{\text{BPM}} = \sum_{s=1}^t S_{i(s)}^\top (S_{i(s)}S_{i(s)}^\top)^{-1} y(s) = b_{t-1}^{\text{BPM}} + S_{i(t)}^\top (S_{i(t)}S_{i(t)}^\top)^{-1} y(t). \quad (20)$$

Therefore, the posterior distribution $\pi\left(p \mid \{i(s), y(s)\}_{s=1}^t\right)$ is

$$\frac{1}{\sqrt{(2\pi)^M |B_t^{\text{BPM}^{-1}}|}} \exp\left\{-\frac{1}{2}(p - B_t^{\text{BPM}^{-1}}b_t^{\text{BPM}})^\top B_t^{\text{BPM}}(p - B_t^{\text{BPM}^{-1}}b_t^{\text{BPM}})\right\}. \quad (21)$$

and this distribution indeed corresponds to the posterior distribution in BPM-TS (Vanchinathan et al., 2014) with $B_t^{\text{BPM}} = \Sigma_t^{-1}$.

Using the same argument, we can confirm that the feedback structure

$$y_t = S_i p + \epsilon, \quad \epsilon \sim \mathcal{N}(0, I_M). \quad (22)$$

induces

$$\bar{g}_t(p) := \frac{1}{\sqrt{(2\pi)^M |B_t^{-1}|}} \exp\left(-\frac{1}{2} \|p - B_t^{-1} b_t\|_{B_t}^2\right), \quad (23)$$

which corresponds to the posterior distribution for TSPM in linear partial monitoring.

Covariances in TSPM-Gaussian and BPM-TS. In the linear partial monitoring, TSPM assumes noise with covariance I_M , which is compatible with the fact that the discrete setting can be regarded as linear PM with I_M -sub-Gaussian noise. On the other hand, BPM-TS assumes covariance $S_i S_i^\top$, and in general $I_M \preceq S_i S_i^\top$ holds. Therefore, BPM-TS assumes unnecessarily larger covariance, which makes learning slow down.

E Preliminaries for Regret Analysis

In this appendix, we give some technical lemmas, which are used for the derivation of the regret bound in Appendix F. Here, we write $X \succeq Y$ to denote $X - Y \succeq 0$. For $a, b \in \mathbb{R}$, let $a \wedge b$ be a if $a \leq b$ otherwise b , and $a \vee b$ be b if $a \leq b$ otherwise a . We use $h(a) := \mathbb{P}_{X \sim \chi_M^2} \{X \geq a\}$ to evaluate the behavior of the posterior samples, where χ_M^2 is the chi-squared distribution with M degree of freedom.

E.1 Basic Lemmas

Fact 5 (Moment generating function of squared-Gaussian distribution). *Let X be the random variable following the standard normal distribution. Then, the moment generating function of X^2 is $\mathbb{E}[\exp(\xi X^2)] = (1 - 2\xi)^{-1/2}$ for $\xi < 1/2$.*

Lemma 6 (Chernoff bound for chi-squared random variable). *Let X be the random variable following the chi-squared distribution with k degree of freedom. Then, for any $a \geq 0$ and $0 \leq \xi < 1/2$,*

$$\mathbb{P}\{X \geq a\} \leq e^{-\xi a} (1 - 2\xi)^{-\frac{k}{2}}. \quad (24)$$

Proof. By Markov's inequality, the LHS can be bounded as

$$\begin{aligned} \mathbb{P}\{X \geq a\} &= \mathbb{P}\left\{\sum_{i=1}^k X_i^2 \geq a\right\} \quad (X_1, \dots, X_k \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)) \\ &= \mathbb{P}\left\{\exp\left(\xi \sum_{i=1}^k X_i^2\right) \geq \exp(\xi a)\right\} \\ &\leq e^{-\xi a} \left(\mathbb{E}\left[e^{\xi X_1^2}\right]\right)^k \quad (\text{by Markov's ineq.}) \\ &= e^{-\xi a} (1 - 2\xi)^{-\frac{k}{2}} \quad (\text{by Fact 5}), \end{aligned} \quad (25)$$

which completes the proof. \square

E.2 Property of Strong Local Observability

Recall that $\Delta_i = (L_i - L_1)^\top p^* > 0$ for $i \in [N]$, which is the difference of the expected loss of actions i and 1. For this define

$$\epsilon := \left(\frac{1}{2\sqrt{A}} \min_{i \neq 1} \frac{\Delta_i}{\|z_{1,i}\|}\right) \wedge \left(\min_{p \in \mathcal{C}_c^i} \frac{4}{3} \|p - p^*\|\right), \quad (26)$$

which is used throughout the proof of this appendix and Appendix F. The following lemma provides the key property of the strong local observability condition.

Lemma 7. *For any partial monitoring game with strong local observability and $p \in \mathbb{R}^M$, any of the conditions 1–3 in the following is not satisfied:*

1. $L_1^\top p > L_k^\top p$ (Worse action k looks better under p .)
2. $\|S_1 p - S_1 p^*\| \leq \epsilon$

$$3. \|S_k p - S_k p^*\| \leq \epsilon .$$

Proof. We prove by contradiction. Assume that there exists $p \in \mathbb{R}^M$ such that conditions 1–3 are simultaneously satisfied.

Now, by the conditions 2 and 3, we have

$$\begin{aligned} |S_1 p - S_1 p^*| &\preceq \epsilon \mathbf{1}_A , \\ |S_k p - S_k p^*| &\preceq \epsilon \mathbf{1}_A . \end{aligned} \quad (27)$$

Here, $|\cdot|$ is the element-wise absolute value, and \preceq means that the inequality \leq holds for each element. Therefore,

$$\left| \begin{pmatrix} S_1 \\ S_k \end{pmatrix} (p - p^*) \right| \preceq \epsilon \mathbf{1}_{2A} . \quad (28)$$

On the other hand, by the strong local observability condition, for any $k \neq 1$, there exists $z_{1,k} \neq 0 \in \mathbb{R}^{2A}$ such that

$$(L_1 - L_k)^\top = z_{1,k}^\top \begin{pmatrix} S_1 \\ S_k \end{pmatrix} . \quad (29)$$

Now, we have

$$\begin{aligned} &z_{1,k}^\top \begin{pmatrix} S_1 \\ S_k \end{pmatrix} (p - p^*) \\ &\leq \|z_{1,k}\| \left\| \begin{pmatrix} S_1 \\ S_k \end{pmatrix} (p - p^*) \right\| \quad (\text{by Cauchy-Schwarz ineq.}) \\ &\leq \sqrt{2A}\epsilon \|z_{1,k}\| \quad (\text{by Eq. (28)}) , \end{aligned} \quad (30)$$

and

$$\begin{aligned} &z_{1,k}^\top \begin{pmatrix} S_1 \\ S_k \end{pmatrix} (p - p^*) \\ &= (L_1 - L_k)^\top (p - p^*) \quad (\text{by Eq. (29)}) \\ &= (L_1 - L_k)^\top p + (L_k - L_1)^\top p^* \\ &\geq \Delta_k \quad (\text{by Condition 1 \& def. of } \Delta_k) . \end{aligned} \quad (31)$$

Therefore, from (30) and (31), we have

$$\Delta_k \leq \sqrt{2A}\epsilon \|z_{1,k}\| . \quad (32)$$

This inequality does not hold for all $k \neq 1$ for the predefined value of ϵ , since we have

$$\epsilon \leq \frac{1}{2\sqrt{A}} \min_{k \neq 1} \frac{\Delta_k}{\|z_{1,k}\|} . \quad (33)$$

Therefore, the proof is completed by contradiction. \square

Remark. The similar result holds when the optimal action 1 is replaced with action $j \neq k$ such that $\Delta_{j,k} := (L_j - L_k)^\top p^* > 0$ by taking ϵ satisfying

$$\epsilon \leq \frac{1}{2\sqrt{A}} \min_{j \neq k: \Delta_{j,k} > 0} \frac{\Delta_{j,k}}{\|z_{j,k}\|} . \quad (34)$$

From Lemma 7, we have the following corollary.

Corollary 8. For any $p \in \mathbb{R}^M$ satisfying $p \in \mathcal{C}_i$ and $\|S_1 p - S_1 p^*\| \leq \epsilon$, we have

$$\|S_i p - S_i p^*\| > \epsilon . \quad (35)$$

Proof. Note that $p \in \mathcal{C}_i$ is equivalent to $(L_1 - L_i)^\top p^* > 0$ for any $i \neq 1$. Therefore, the result directly follows from Lemma 7. \square

The next lemma is the property of Mahalanobis distance corresponding to $\bar{g}_t(p)$.

Lemma 9. Define $\mathcal{J}_i := \{p \in \mathbb{R}^M : \|S_i p - S_i p^*\| > \epsilon\}$. Assume that $N_i(t) \geq n_i$, $\|S_i \hat{p}_t - S_i p^*\| \leq \epsilon/4$. Then, for any $0 \leq \xi < 1/2$

$$h\left(\inf_{p \in \mathcal{J}_i} \|B_t^{1/2}(p - \hat{p}_t)\|^2\right) \leq \exp\left(-\frac{9}{16}\xi\epsilon^2 n_i\right)(1 - 2\xi)^{-M/2}. \quad (36)$$

Proof. To bound the LHS of the above inequality, we bound $\|B_t^{1/2}(p - \hat{p}_t)\|^2$ from below for $p \in \mathcal{J}_i$. Using the triangle inequality and the assumptions, we have

$$\begin{aligned} \|S_i(p - \hat{p}_t)\| &\geq \|S_i p - S_i p^*\| - \|S_i \hat{p}_t - S_i p^*\| \\ &> \epsilon - \epsilon/4 > 0. \end{aligned} \quad (37)$$

Therefore, we have

$$\begin{aligned} \|B_t^{1/2}(p - \hat{p}_t)\|^2 &\geq \sum_{k \in [N]} N_k(t) \|S_k(p - \hat{p}_t)\|^2 \quad (\text{by def. of } B_t) \\ &\geq n_i \|S_i(p - \hat{p}_t)\|^2 \quad (N_i(t) \geq n_i) \\ &> \frac{9}{16}\epsilon^2 n_i \quad (\text{by Eq. (37)}) . \end{aligned} \quad (38)$$

By the Chernoff bound for a chi-squared random variable in Lemma 6, we now have

$$h(a) \leq e^{-\xi a}(1 - 2\xi)^{-M/2}, \quad (39)$$

for any $a \geq 0$ and $0 \leq \xi < 1/2$. Hence, using the fact that $\|B_t^{1/2}(p - \hat{p}_t)\|^2$ follows the chi-squared distribution with M degree of freedom, we have

$$\begin{aligned} h\left(\inf_{p \in \mathcal{J}_i} \|B_t^{1/2}(p - \hat{p}_t)\|^2\right) &\leq h\left(\frac{9}{16}\epsilon^2 n_i\right) \\ &\leq \exp\left(-\frac{9}{16}\xi\epsilon^2 n_i\right)(1 - 2\xi)^{-M/2}, \end{aligned} \quad (40)$$

which completes the proof. \square

E.3 Statistics of Uninterested Actions

For any $k \neq i$ and $n_k \in [T]$, define

$$Z_{n_k} := n_k \|q_{k,n_k} - S_k p^*\|^2, \quad (41)$$

$$Z_{\setminus i} := \sum_{k \neq i} \max_{n_k \in [T]} Z_{n_k}. \quad (42)$$

In this section, we bound $\mathbb{E}[Z_{\setminus i}]$ from above. Note that $Z_{\setminus i}$ is independent of the randomness of Thompson sampling.

Lemma 10 (Upper bound for the expectation of $Z_{\setminus i}$).

$$\mathbb{E}[Z_{\setminus i}] \leq 4N \left(\log T + \frac{A}{2} \log 2 + 1 \right). \quad (43)$$

Proof. Recall that in linear partial monitoring, the feedback $y(t) \in \mathbb{R}^A$ for action k is given as

$$y_t = S_k p^* + \epsilon, \quad \epsilon \sim \mathcal{N}(0, I_A) \quad (44)$$

at round $t \in [T]$, Therefore, $y(t) - S_k p^* \sim \mathcal{N}(0, I_A)$. Since $q_{k,n_k} = \frac{1}{n_k} \sum_{s \in [T]: i(s)=k} y(s)$ for any $n_k \in [T]$, we have

$$q_{k,n_k} - S_k p^* = \frac{1}{n_k} \sum_{s \in [T]: i(s)=k} (y(s) - S_k p^*) \sim \mathcal{N}(0, I_A/n_k). \quad (45)$$

Therefore,

$$\sqrt{n_k}(q_{k,n_k} - S_k p^*) \sim \mathcal{N}(0, I_A), \quad (46)$$

and thus

$$n_k \|q_{k,n_k} - S_k p^*\|^2 = \|\sqrt{n_k}(q_{k,n_k} - S_k p^*)\|^2 \sim \chi_A^2. \quad (47)$$

Therefore, for any $0 \leq \xi < 1/2$,

$$\begin{aligned} \mathbb{E} \left[\max_{n_k \in [T]} Z_{n_k} \right] &= \int_0^\infty \mathbb{P} \left\{ \max_{n_k \in [T]} Z_{n_k} \geq x \right\} dx \\ &\leq \int_0^\infty [1 \wedge T \cdot \mathbb{P}\{Z_1 \geq x\}] dx \quad (\text{by the union bound}) \\ &\leq \int_0^\infty [1 \wedge T \cdot e^{-\xi x} (1 - 2\xi)^{-\frac{A}{2}}] dx \quad (\text{by } Z_1 \sim \chi(\xi^2 A) \text{ and Lemma 6}) \\ &= \int_0^{x^*} dx + \int_{x^*}^\infty T \cdot e^{-\xi x} (1 - 2\xi)^{-\frac{A}{2}} dx \\ &\leq x^* + T \cdot \int_{x^*}^\infty e^{-\xi x} (1 - 2\xi)^{-\frac{A}{2}} dx \\ &= x^* + T(1 - 2\xi)^{-\frac{A}{2}} \left[-\frac{1}{\xi} e^{-\xi x} \right]_{x^*}^\infty \\ &= \frac{1}{\xi} \left\{ \log T - \frac{A}{2} \log(1 - 2\xi) + 1 \right\}, \end{aligned} \quad (48)$$

where $x^* := \frac{1}{\xi} \left\{ \log T - \frac{A}{2} \log(1 - 2\xi) \right\}$. Therefore, taking $\xi = 1/4$, we have

$$\begin{aligned} \mathbb{E} [Z_{\setminus i}] &= \mathbb{E} \left[\sum_{k \neq i} \max_{n_k \in [T]} Z_{n_k} \right] \\ &\leq \sum_{k \neq i} \mathbb{E} \left[\max_{n_k \in [T]} Z_{n_k} \right] \\ &\leq (N - 1) \frac{1}{\xi} \left\{ \log T - \frac{A}{2} \log(1 - 2\xi) + 1 \right\} \\ &\leq 4N \left(\log T + \frac{A}{2} \log 2 + 1 \right), \end{aligned} \quad (49)$$

which completes the proof. \square

E.4 Mahalanobis Distance Process

Discussions in this section are essentially very similar to [Abbasi-Yadkori et al. \(2011, Lemma 11\)](#), but their results are not directly applicable and we give the full derivation for self-containedness. To maximize the applicability here we only assume sub-Gaussian noise rather than a Gaussian one.

Let ϵ_t be zero-mean 1-sub-Gaussian random variable, which satisfies

$$\mathbb{E} \left[e^{\lambda^\top \epsilon_t} \right] \leq e^{-\frac{\|\lambda\|^2}{2}} \quad (50)$$

for any $\lambda \in \mathbb{R}^M$.

Lemma 11. For any vector $v \in \mathbb{R}^M$ and positive definite matrix $V \in \mathbb{R}^{M \times M}$ such that $V \succ I$,

$$\mathbb{E}_{\epsilon_t} \left[e^{\frac{\|\epsilon_t + v\|_V^2}{2}} \right] \leq \frac{\sqrt{|V|}}{\sqrt{|V - I|}} e^{\frac{1}{2} v^\top (V - I)^{-1} v}. \quad (51)$$

Proof. For any $x \in \mathbb{R}^M$

$$\mathbb{E}_{\lambda \sim \mathcal{N}(0, V^{-1})} \left[e^{\lambda^\top x} \right] = e^{\frac{\|x\|_{V^{-1}}^2}{2}}. \quad (52)$$

Therefore, by letting $x = \epsilon_t + v$ we see that

$$e^{\frac{\|\epsilon_t + v\|_V^2}{2}} = \mathbb{E}_{\lambda \sim \mathcal{N}(0, V^{-1})} \left[e^{\lambda^\top (\epsilon_t + v)} \right]. \quad (53)$$

As a result, by the definition of sub-Gaussian random variables, we have

$$\begin{aligned} \mathbb{E}_{\epsilon_t} \left[e^{\frac{\|\epsilon_t + v\|_V^2}{2}} \right] &= \mathbb{E}_{\lambda \sim \mathcal{N}(0, V^{-1})} \left[\mathbb{E}_{\epsilon_t} \left[e^{\lambda^\top (\epsilon_t + v)} \right] \right] \\ &= \mathbb{E}_{\lambda \sim \mathcal{N}(0, V^{-1})} \left[e^{\lambda^\top v} \mathbb{E}_{\epsilon_t} \left[e^{\lambda^\top \epsilon_t} \right] \right] \\ &\leq \mathbb{E}_{\lambda \sim \mathcal{N}(0, V^{-1})} \left[e^{\lambda^\top v} e^{\|\lambda\|^2/2} \right] \\ &= \frac{1}{(2\pi)^{d/2} \sqrt{|V^{-1}|}} \int e^{\lambda^\top v} e^{\|\lambda\|^2/2} e^{-\|\lambda\|_V^2/2} d\lambda \\ &= \frac{1}{(2\pi)^{d/2} \sqrt{|V^{-1}|}} \int e^{-\frac{1}{2}(\lambda^\top (V-I)\lambda - 2v^\top \lambda)} d\lambda \\ &= \frac{\sqrt{|V-I|}}{(2\pi)^{d/2} \sqrt{|V^{-1}| |V-I|}} \int e^{-\frac{1}{2}((\lambda - (V-I)^{-1}v)^\top (V-I)(\lambda - (V-I)^{-1}v) - v^\top (V-I)^{-1}v)} d\lambda \\ &= \frac{\sqrt{|V|}}{\sqrt{|V-I|}} e^{\frac{1}{2}v^\top (V-I)^{-1}v}. \end{aligned} \quad (54)$$

□

Lemma 12.

$$\mathbb{E} \left[\exp \left(\frac{1}{2} \left(\|\hat{p}_t - p^*\|_{B_t}^2 - \|\hat{p}_{t-1} - p^*\|_{B_{t-1}}^2 \right) \right) \middle| \hat{p}_{t-1}, B_{t-1}, S_{i(t-1)} \right] \leq \sqrt{\frac{|B_t|}{|B_{t-1}|}}. \quad (55)$$

Proof. Let $Z_t := -\lambda p^* + \sum_{s=1}^t S_{i(s)}^\top \epsilon_s$, and we have

- $B_t = \lambda I + \sum_{s=1}^t S_{i(s)}^\top S_{i(s)}$,
- $b_t = \sum_{s=1}^t S_{i(s)}^\top y(s) = B_t p^* + Z_t$,
- $\hat{p}_t = B_t^{-1} b_t = p^* + B_t^{-1} Z_t$.

In the following we omit the conditioning on $(\hat{p}_{t-1}, B_{t-1}, S_{i(t-1)})$ for notational simplicity.

Let us define $C_t := S_{i(t)} B_{t-1} S_{i(t)}^\top$ and $d_t := S_{i(t)} B_{t-1}^{-1} Z_{t-1} = S_{i(t)} (\hat{p}_t - p^*)$. Then, using the Sherman-Morrison-Woodbury formula we have

$$\begin{aligned}
& \|\hat{p}_t - p^*\|_{B_t}^2 - \|\hat{p}_{t-1} - p^*\|_{B_{t-1}}^2 \\
&= Z_t^\top B_t^{-1} Z_t - Z_{t-1}^\top B_{t-1}^{-1} Z_{t-1} \\
&= (Z_{t-1}^\top + \epsilon_t^\top S_{i(t)})(B_{t-1}^{-1} - B_{t-1}^{-1} S_{i(t)}^\top (I + S_{i(t)} B_{t-1}^{-1} S_{i(t)}^\top)^{-1} S_{i(t)} B_{t-1}^{-1})(Z_{t-1} + S_{i(t)}^\top \epsilon_t) - Z_{t-1}^\top B_{t-1}^{-1} Z_{t-1} \\
&= (Z_{t-1}^\top + \epsilon_t^\top S_{i(t)}) B_{t-1}^{-1} (Z_{t-1} + S_{i(t)}^\top \epsilon_t) - Z_{t-1}^\top B_{t-1}^{-1} Z_{t-1} \\
&\quad - (Z_{t-1}^\top + \epsilon_t^\top S_{i(t)}) B_{t-1}^{-1} S_{i(t)}^\top (I + S_{i(t)} B_{t-1}^{-1} S_{i(t)}^\top)^{-1} S_{i(t)} B_{t-1}^{-1} (Z_{t-1} + S_{i(t)}^\top \epsilon_t) \\
&= \epsilon_t^\top S_{i(t)} B_{t-1}^{-1} S_{i(t)}^\top \epsilon_t + 2 Z_{t-1}^\top B_{t-1}^{-1} S_{i(t)}^\top \epsilon_t \\
&\quad - (Z_{t-1}^\top + \epsilon_t^\top S_{i(t)}) B_{t-1}^{-1} S_{i(t)}^\top (I + S_{i(t)} B_{t-1}^{-1} S_{i(t)}^\top)^{-1} S_{i(t)} B_{t-1}^{-1} (Z_{t-1} + S_{i(t)}^\top \epsilon_t) \\
&= \epsilon_t^\top C_t \epsilon_t + 2 d_t^\top \epsilon_t - (d_t^\top + \epsilon_t^\top C_t)(I + C_t)^{-1} (d_t + C_t \epsilon_t) \\
&= \epsilon_t^\top C_t (I - (I + C_t)^{-1} C_t) \epsilon_t + 2 d_t^\top (I - (I + C_t)^{-1} C_t) \epsilon_t - d_t^\top (I + C_t)^{-1} d_t \\
&= \epsilon_t^\top C_t (I + C_t)^{-1} \epsilon_t + 2 d_t^\top (I + C_t)^{-1} \epsilon_t - d_t^\top (I + C_t)^{-1} d_t \\
&= \|\epsilon_t + C_t^{-1} d_t\|_{C_t(I+C_t)^{-1}}^2 - d_t^\top (I + C_t)^{-1} C_t^{-1} d_t - d_t^\top (I + C_t)^{-1} d_t \\
&= \|\epsilon_t + C_t^{-1} d_t\|_{C_t(I+C_t)^{-1}}^2 - d_t^\top (I + C_t)^{-1} (I + C_t^{-1}) d_t. \tag{56}
\end{aligned}$$

Therefore, Lemma 11 with $V := (C_t(I + C_t)^{-1})^{-1} = (I + C_t)C_t^{-1}$, $v := C_t^{-1}d_t$ yields

$$\begin{aligned}
& \mathbb{E} \left[\exp \left(\frac{1}{2} \left(\|\hat{p}_t - p^*\|_{B_t}^2 - \|\hat{p}_{t-1} - p^*\|_{B_{t-1}}^2 \right) \right) \right] \\
& \leq \frac{\sqrt{|(I + C_t)C_t^{-1}|}}{\sqrt{|(I + C_t)C_t^{-1} - I|}} e^{\frac{1}{2} d_t^\top C_t^{-1} ((I + C_t)C_t^{-1} - I)^{-1} C_t^{-1} d_t} e^{-\frac{1}{2} d_t^\top (I + C_t)^{-1} (I + C_t^{-1}) d_t} \\
& \leq \frac{\sqrt{|(I + C_t)C_t^{-1}|}}{\sqrt{|C_t^{-1}|}} e^{\frac{1}{2} d_t^\top C_t^{-1} (C_t^{-1})^{-1} C_t^{-1} d_t} e^{-\frac{1}{2} d_t^\top (I + C_t)^{-1} (I + C_t^{-1}) d_t} \\
& = \sqrt{|(I + C_t)|} \\
& = \sqrt{\frac{|B_t|}{|B_{t-1}|}}, \tag{57}
\end{aligned}$$

where see, e.g., [Abbasi-Yadkori et al. \(2011, Lemma 11\)](#) for the last equality. \square

E.5 Norms under Perturbations

In the following two lemmas, we give some analysis of norms under perturbations.

Lemma 13. *Let A be a positive definite matrix. Let $a \in \mathbb{R}^d$ and $\epsilon > 0$ be such that $\epsilon < \|a\|/3$. Then*

$$\min_{x: \|x\| \leq 2\epsilon} \max_{x': \|x'\| \leq \epsilon} \{(a + x + x')^\top A(a + x + x')\} = \min_{x'': \|x''\| \leq \epsilon} \{(a + x'')^\top A(a + x'')\}. \tag{58}$$

Proof. By considering the Lagrangian multiplier we see that any stationary point of the function $(a + x'')^\top A(a + x'')$ over $\{(x, x') : \|x\| \leq 2\epsilon, \|x'\| \leq \epsilon\}$ satisfies

$$\begin{aligned}
A(a + x + x') - \lambda_1 x &= 0, \\
A(a + x + x') - \lambda_2 x' &= 0, \\
x^\top x &= 4\epsilon^2, \\
x'^\top x' &= \epsilon^2, \tag{59}
\end{aligned}$$

and therefore $\lambda_1 x = \lambda_2 x'$. Considering the last two conditions of (59) we have $\lambda_2 = \pm 2\lambda_1$, implying that

$$x' = -(3A - 2\lambda_1 I)Aa \tag{60}$$

or

$$x' = (A - 2\lambda_1 I)Aa \quad (61)$$

for λ_1 satisfying $x'^\top x' = \epsilon^2$.

Note that it holds for any positive definite matrix B that

$$\frac{d^2}{d\lambda^2} a(B + \lambda I)^{-2} a = a(B + \lambda I)^{-4} a = \|(B + \lambda I)^{-2} a\|^2, \quad (62)$$

which is positive almost everywhere, meaning that $a(B + \lambda I)^{-2} a$ is strictly convex with respect to $\lambda \in \mathbb{R}$. Therefore, there exists at most two λ_1 's satisfying (60) and $x'^\top x' = \epsilon^2$, and there exists at most two λ_1 's satisfying (61) and $x'^\top x' = \epsilon^2$. In summary, there at most four stationary points of $(a + x'')^\top A(a + x'')$ over $\{(x, x') : \|x\| \leq 2\epsilon, \|x'\| \leq \epsilon\}$.

On the other hand, two optimization problems

$$\min_{x: \|x\| \leq 2\epsilon} \min_{x': \|x'\| \leq \epsilon} \{(a + x + x')^\top A(a + x + x')\} = \min_{x'': \|x''\| \leq 3\epsilon} \{(a + x'')^\top A(a + x'')\} \quad (63)$$

and

$$\max_{x: \|x\| \leq 2\epsilon} \max_{x': \|x'\| \leq \epsilon} \{(a + x + x')^\top A(a + x + x')\} = \max_{x'': \|x''\| \leq 3\epsilon} \{(a + x'')^\top A(a + x'')\} \quad (64)$$

can be easily solved by an elementary calculation and the optimal values are equal to those corresponding to (60).

Therefore, the optimal solutions of the two minimax problems

$$\max_{x: \|x\| \leq 2\epsilon} \min_{x': \|x'\| \leq \epsilon} \{(a + x + x')^\top A(a + x + x')\} \quad (65)$$

and

$$\min_{x: \|x\| \leq 2\epsilon} \max_{x': \|x'\| \leq \epsilon} \{(a + x + x')^\top A(a + x + x')\} \quad (66)$$

correspond to two points corresponding to (61).

We can see again from an elementary calculation that the optimal solutions for two optimization problems

$$\begin{aligned} & \min_{x'': \|x''\| \leq \epsilon} \{(a + x'')^\top A(a + x'')\} \\ & \max_{x'': \|x''\| \leq \epsilon} \{(a + x'')^\top A(a + x'')\} \end{aligned} \quad (67)$$

have the same necessary and sufficient conditions as (61) and we complete the proof by noticing that (65) is less than (66). \square

Lemma 14. Let $A \succeq nS_1^\top S_1$ be a positive-definite matrix with minimum eigenvalue at least $\lambda > 0$. Then, for any $\hat{p} \in \mathbb{R}^d$ and $\epsilon > 0$ satisfying $\epsilon < \|\hat{p} - p^*\|/3$,

$$\|\hat{p} - p^*\|_A^2 - \inf_{p: \|p - p^*\| \leq 2\epsilon} \sup_{p': \|p' - p\| \leq \epsilon} \|p' - \hat{p}\|_A^2 \geq \epsilon \sqrt{n\lambda} \|S_1(\hat{p} - p^*)\|. \quad (68)$$

Proof. Let $a = \hat{p} - p^*$. By Lemma 13, we have

$$\begin{aligned} & \inf_{p: \|p - p^*\| \leq 2\epsilon} \sup_{p': \|p' - p\| \leq \epsilon} \|p' - \hat{p}\|_A^2 \\ &= \inf_{x: \|x\| \leq 2\epsilon} \sup_{x': \|x'\| \leq \epsilon} \|a + x + x'\|_A^2 \\ &= \inf_{x: \|x\| \leq \epsilon} \|a + x\|_A^2. \end{aligned} \quad (69)$$

Now define $\mathcal{S}_{\epsilon', A} = \{x : \|x\|_A \leq \epsilon'\}$. Then, we see that $\mathcal{S}_{\epsilon\sqrt{\lambda}, A} \subset \{x : \|x\| \leq \epsilon\}$. Therefore, an elementary calculation using the Lagrange multiplier technique shows

$$\begin{aligned} \inf_{x: \|x\| \leq \epsilon} \|p' - \hat{p}\|_A^2 &\leq \inf_{x \in \mathcal{S}_{\epsilon\sqrt{\lambda}, A}} \|p' - \hat{p}\|_A^2 \\ &= \left(\|a\|_A - \epsilon\sqrt{\lambda} \right)^2. \end{aligned} \quad (70)$$

As a result, we see that

$$\begin{aligned} \|p^* - \hat{p}\|_A^2 - \inf_{p: \|p-p^*\| \leq 2\epsilon} \sup_{p': \|p'-p\| \leq \epsilon} \|p' - \hat{p}\|_A^2 &\geq \|a\|_A^2 - \left(\|a\|_A - \epsilon\sqrt{\lambda} \right)^2 \\ &= \epsilon\sqrt{\lambda} \left(\|a\|_A + \|a\|_A - \epsilon\sqrt{\lambda} \right) \\ &\geq \epsilon\sqrt{\lambda} \left(\|a\|_A + \|a\| \sqrt{\lambda} - \epsilon\sqrt{\lambda} \right) \\ &= \epsilon\sqrt{\lambda} \left(\|a\|_A + \sqrt{\lambda}(\|a\| - \epsilon) \right) \\ &\geq \epsilon\sqrt{\lambda} \|a\|_A \\ &\geq \epsilon\sqrt{n\lambda} \|S_1 a\|. \end{aligned} \quad (71)$$

□

For the subsets of \mathbb{R}^n , \mathcal{X} and \mathcal{Y} , let $\mathcal{X} + \mathcal{Y} := \{x + y : x \in \mathcal{X}, y \in \mathcal{Y}\}$ be the Minkowski sum, and let $B_r^n(p)$ be the n -dimensional Euclidian ball of radius r at point $p \in \mathbb{R}^n$ (the superscript n can be omitted when it is clear from context). We also let ϵ' be

$$\epsilon' := \frac{\epsilon}{\left(16 \max_{i \in [N]} \|S_i\| \right) \vee \left(\frac{1}{\sqrt{A}} \max_{i \in [N]} \frac{\|L_i - L_1\|}{\|z_{1,i}\|} \right)}, \quad (72)$$

which is also used throughout the proof of this appendix and Appendix F as ϵ in (26).

Theorem 15. *Let $\epsilon'' \in (0, \epsilon)$ be a constant for ϵ defined in (26). Let $\hat{p} \in \mathcal{C}_k + B_{\epsilon'}^d(0)$ be satisfying $\|S_k(\hat{p} - p^*)\| \leq \epsilon''$. Then, there exists $\delta > 0$ satisfying for any $n \geq 0$ and $A \succeq nS_1^\top S_1 + \lambda I$ that*

$$\|p^* - \hat{p}\|_A^2 - \inf_{p: \|p-p^*\| \leq 2\epsilon} \sup_{p': \|p'-p\| \leq \epsilon} \|p' - \hat{p}\|_A^2 \geq \epsilon\delta\sqrt{\lambda n}. \quad (73)$$

Proof. Recall that $\epsilon'' < \epsilon \leq \min_{p \in \mathcal{C}_k} \|p - p^*\| / 3$. It is enough from Lemma 14 to prove that

$$\delta := \min_{\hat{p} \in \{p \in \mathcal{C}_k + B_{\epsilon'}^d(0) : \|S_k(p - p^*)\| \leq \epsilon''\}} \|S_1(\hat{p} - p^*)\| \quad (74)$$

is positive.

We prove by contradiction and the proof is basically same as that of Lemma 7 but more general in the sense that the condition on \hat{p} is not $\hat{p} \in \mathcal{C}_k$ but $\hat{p} \in \mathcal{C}_k + B_{\epsilon'}^d(0)$. Assume that $\delta = 0$, that is, there exists $\hat{p} \in \mathcal{C}_k + B_{\epsilon'}^d(0)$ satisfying $\|S_k(p - p^*)\| \leq \epsilon''$ and $\|S_1(\hat{p} - p^*)\| = 0$. Note that $\|S_1(\hat{p} - p^*)\| = 0$ implies $\|S_1(\hat{p} - p^*)\| \leq \epsilon''$. Therefore, we now have following conditions on \hat{p} :

- $\hat{p} \in \mathcal{C}_k + B_{\epsilon'}^d(0)$
- $\|S_1(\hat{p} - p^*)\| \leq \epsilon''$
- $\|S_k(\hat{p} - p^*)\| \leq \epsilon''$.

Following the same argument as the proof of Lemma 7, we have

$$z_{1,k}^\top \begin{pmatrix} S_1 \\ S_k \end{pmatrix} (\hat{p} - p^*) \leq \sqrt{2A}\epsilon'' \|z_{1,k}\|. \quad (75)$$

On the other hand, since $\hat{p} \in \mathcal{C}_k + B_{\epsilon'}^d(0)$ we can take $\bar{p} \in \mathcal{C}_k$ such that $\|\hat{p} - \bar{p}\| \leq \epsilon'$. Hence,

$$\begin{aligned} z_{1,k}^\top \begin{pmatrix} S_1 \\ S_k \end{pmatrix} (\hat{p} - p^*) &= (L_1 - L_k)^\top (\hat{p} - p^*) \\ &= -(L_k - L_1)^\top (\hat{p} - \bar{p}) + (L_1 - L_k)^\top \bar{p} + (L_k - L_1)^\top p^* \\ &\geq -(L_k - L_1)^\top (\hat{p} - p^*) + \Delta_k. \quad (\text{by } \bar{p} \in \mathcal{C}_k \text{ and def. of } \Delta_k) \end{aligned} \quad (76)$$

From (75) and (76), we have

$$\Delta_k - (L_k - L_1)^\top (\hat{p} - p^*) \leq \sqrt{2A}\epsilon'' \|z_{1,k}\|. \quad (77)$$

Now, the left hand side of (77) is bounded from below as

$$\begin{aligned} \Delta_k - (L_k - L_1)^\top (\hat{p} - \bar{p}) &\geq \Delta_k - \|L_k - L_1\| \|\hat{p} - \bar{p}\| \\ &\geq \Delta_k - \|L_k - L_1\| \epsilon' \\ &= \Delta_k - \|L_k - L_1\| \frac{\epsilon}{\frac{1}{\sqrt{A}} \max_i \frac{\|L_1 - L_i\|}{\|z_{1,i}\|}} \\ &= \Delta_k - \|L_k - L_1\| \frac{\frac{1}{2\sqrt{A}} \min_i \frac{\Delta_i}{\|z_{1,i}\|}}{\frac{1}{\sqrt{A}} \max_i \frac{\|L_1 - L_i\|}{\|z_{1,i}\|}} \\ &\geq \Delta_k - \Delta_k/2. \end{aligned} \quad (78)$$

On the other hand, using the definition of ϵ'' , the right hand side of (77) is bounded from above as

$$\sqrt{2A}\epsilon'' \|z_{1,k}\| < \Delta_k/2. \quad (79)$$

Therefore, the proof is completed by contradiction. \square

E.6 Exit Time Analysis

We next consider the exit time. Let \mathcal{A}_t be an event deterministic given \mathcal{F}_t , and \mathcal{B}_t be a random event such that if \mathcal{B}_t occurred then $\mathcal{A}_{t'}$ never occurs for $t' = t+1, t+2, \dots$. Let $P_t, t = 1, 2, \dots, T$, be a stochastic process satisfying $P_t \leq \mathbb{P}\{\mathcal{B}_t | \mathcal{F}_t\}$ a.s. and P_t^{-1} is a supermartingale with respect to the filtration induced by \mathcal{F}_t .

Theorem 16. *Let τ be the stopping time defined as*

$$\tau = \begin{cases} \min\{t \in [T] : \mathcal{A}_t\} & \text{if } \mathcal{A}_t \text{ occurs for some } t \in [T]. \\ T+1 & \text{otherwise.} \end{cases} \quad (80)$$

Then we almost surely have

$$\mathbb{E} \left[\sum_{t=1}^T \mathbb{1}[\mathcal{A}_t] \middle| \mathcal{F}_\tau \right] \leq \begin{cases} P_\tau^{-1} & \tau \leq T, \\ 0 & \tau = T+1. \end{cases} \quad (81)$$

We prove this theorem based on the following lemma.

Lemma 17. *Let $(Q_i)_{i=1}^\infty \subset [0, 1]$ be an arbitrary stochastic process such that $(Q_i^{-1})_{i=1}^\infty$ is a supermartingale with respect to a filtration $(\mathcal{G}_i)_{i=1}^\infty$. Then, for any $\mathcal{G}_0 \subset \mathcal{G}_1$,*

$$\mathbb{E} \left[\sum_{i=1}^T \prod_{j=1}^i (1 - Q_j) \middle| \mathcal{G}_0 \right] \leq \mathbb{E} [Q_1^{-1} | \mathcal{G}_0] - 1 \quad \text{a.s.} \quad (82)$$

Proof. Let

$$\begin{aligned} N_k((Q_i, \mathcal{G}_i)_{i=1}^\infty, \mathcal{G}_0) &= \mathbb{E} \left[\sum_{i=1}^k \prod_{j=1}^i (1 - Q_j) \middle| \mathcal{G}_0 \right] \\ \bar{N}_k((Q_i, \mathcal{G}_i)_{i=1}^\infty, \mathcal{G}_0) &= \mathbb{E} \left[\sum_{i=1}^\infty \prod_{j=1}^i (1 - Q_j) \middle| \mathcal{G}_0 \right] \quad \text{where } Q_j = Q_k \text{ for } j > k. \end{aligned} \quad (83)$$

We show $\bar{N}_k((Q_i, \mathcal{G}_i)_{i=1}^\infty, \mathcal{G}_0) \leq \mathbb{E}[Q_1^{-1} | \mathcal{G}_0] - 1$ a.s. for any $(Q_i, \mathcal{G}_i)_{i=1}^\infty, \mathcal{G}_0 \subset \mathcal{G}_1$ and $k \in \mathbb{N}$ by induction. First, for $k = 1$ the statement holds since

$$\begin{aligned} \bar{N}_1((Q_i, \mathcal{G}_i)_{i=1}^\infty, \mathcal{G}_0) &= \mathbb{E} \left[\sum_{i=1}^{\infty} \prod_{j=1}^i (1 - Q_j) \middle| \mathcal{G}_0 \right] \\ &= \mathbb{E} [Q_1^{-1} - 1 | \mathcal{G}_0] \\ &= \mathbb{E} [Q_1^{-1} | \mathcal{G}_0] - 1 \end{aligned} \quad (84)$$

Next, assume that the statement holds for all $(Q_i, \mathcal{G}_i)_{i=1}^k, \mathcal{G}_0 \subset \mathcal{G}_1$ and $k \leq k_0$. Then, we almost surely have

$$\begin{aligned} \bar{N}_{k_0+1}((Q_i, \mathcal{G}_i)_{i=1}^\infty, \mathcal{G}_0) &= \mathbb{E} \left[(1 - Q_1) \mathbb{E} \left[1 + \sum_{i=2}^{\infty} \prod_{j=2}^i (1 - Q_j) \middle| \mathcal{G}_1 \right] \middle| \mathcal{G}_0 \right] \\ &= \mathbb{E} [(1 - Q_1)(1 + \bar{N}_{k_0}((Q_i, \mathcal{G}_i)_{i=2}^\infty, \mathcal{G}_1)) | \mathcal{G}_0] \\ &\leq \mathbb{E} [(1 - Q_1)\mathbb{E}[Q_2^{-1} | \mathcal{G}_1] | \mathcal{G}_0] \quad (\text{assumption of the induction}) \\ &\leq \mathbb{E} [Q_1^{-1} | \mathcal{G}_0] - 1 \quad (Q_i^{-1} \text{ is a supermartingale.}) \end{aligned} \quad (85)$$

We obtain the lemma from

$$\mathbb{E} \left[\sum_{i=1}^k \prod_{j=1}^i (1 - Q_j) \middle| \mathcal{G}_0 \right] = N_k((Q_i, \mathcal{G}_i)_{i=1}^\infty, \mathcal{G}_0) \leq \bar{N}_k((Q_i, \mathcal{G}_i)_{i=1}^\infty, \mathcal{G}_0) \quad \text{a.s.} \quad (86)$$

□

Proof of Theorem 16. The statement is obvious for the case $\tau = T + 1$ and we consider the other case in the following.

Let τ_i be the time of the i -th occurrence of \mathcal{A}_t . More formally, we define τ_i as the stopping time $\tau_1 = \tau$ and

$$\tau_{i+1} = \begin{cases} \min \left\{ t \in [T] : \sum_{t'=1}^t \mathbb{1}[\mathcal{A}_{t'}] = i + 1 \right\} & \sum_{t'=1}^T \mathbb{1}[\mathcal{A}_{t'}] \geq i + 1, \\ \tau_i + 1 & \text{otherwise.} \end{cases} \quad (87)$$

Then $(P'_i) = (P_{\tau_i})$ is a stochastic process measurable by the filtration induced by $(\mathcal{F}'_i) = (\mathcal{F}_{\tau_i})$. By Lemma 17 we obtain

$$\begin{aligned} \mathbb{E} \left[\sum_{t=1}^T \mathbb{1}[\mathcal{A}_t] \middle| \mathcal{F}_\tau \right] &= \mathbb{E} \left[\sum_{n=1}^T \mathbb{1} \left[\sum_{t=1}^T \mathbb{1}[\mathcal{A}_t] \geq n \right] \middle| \mathcal{F}_\tau \right] \\ &\leq 1 + \mathbb{E} \left[\sum_{n=2}^T \mathbb{1} \left[\sum_{t=1}^T \mathbb{1}[\mathcal{A}_t] \geq n \right] \middle| \mathcal{F}_\tau \right] \\ &\leq 1 + \mathbb{E} \left[\sum_{i=1}^T \prod_{j=1}^i (1 - P'_j) \middle| \mathcal{F}'_1 \right] \\ &\leq 1 + \mathbb{E} [(P'_1)^{-1} | \mathcal{F}'_1] - 1 \\ &= P_\tau^{-1}. \end{aligned} \quad (88)$$

□

F Regret Analysis of TSPM Algorithm

In this appendix, we give the proof of Theorem 3. Note that the cells are defined for the decomposition of \mathbb{R}^M , not \mathcal{P}_M . In other words, the cell \mathcal{C}_i is here defined as $\mathcal{C}_i = \{p \in \mathbb{R}^M : \text{action } i \text{ is optimal}\}$.

For the linear setting, the empirical feedback distribution $q_i^{(t)}$ and $q_{i,n}$ are defined as

$$q_i^{(t)} := \frac{1}{N_i(t)} \sum_{s \in [t-1]: i(s)=i} y(s), \quad (89)$$

$$q_{i,n} := \text{the value of } q_i^{(t)} \text{ after taking action } i \text{ for } n \text{ times.} \quad (90)$$

Recall that $\hat{p}_t = B_t^{-1} b_t$, which is the mode of $\bar{g}_t(p)$.

F.1 Regret Decomposition

Here, we break the regret into several terms. For any $i \in [N]$, we define events

$$\mathcal{A}_i(t) := \left\{ \|S_i \hat{p}_t - S_i p^*\| \leq \frac{\epsilon}{4} \right\}, \quad (91)$$

$$\tilde{\mathcal{A}}_i(t) := \{ \|S_i \tilde{p}_t - S_i p^*\| \leq \epsilon \}. \quad (92)$$

We first decompose the regret as

$$\begin{aligned} \text{Reg}(T) &= \sum_{t=1}^T \Delta_{i(t)} \\ &\leq \sum_{t=1}^T \left(\Delta_{i(t)} \mathbb{1}[\tilde{\mathcal{A}}_1(t)] + \max_{j \in [N]} \Delta_j \mathbb{1}[\tilde{\mathcal{A}}_1^c(t)] \right) \\ &= \sum_{i \neq 1} \sum_{t=1}^T \Delta_i \mathbb{1}[i(t) = i, \tilde{\mathcal{A}}_1(t)] + \max_{j \in [N]} \Delta_j \sum_{t=1}^T \mathbb{1}[\tilde{\mathcal{A}}_1^c(t)] \\ &\leq \sum_{i \neq 1} \Delta_i \sum_{t=1}^T \left(\underbrace{\mathbb{1}[i(t) = i, \tilde{\mathcal{A}}_1(t), \mathcal{A}_i(t)]}_{(A)} + \underbrace{\mathbb{1}[i(t) = i, \mathcal{A}_i^c(t)]}_{(B)} \right) + \max_{j \in [N]} \Delta_j \sum_{t=1}^T \mathbb{1}[\tilde{\mathcal{A}}_1^c(t)]. \end{aligned} \quad (93)$$

To decompose the last term, we define the following notation. We define for any $i \in [N]$

$$P_i(t) := \mathbb{P}\{\tilde{p}_t \in \mathcal{C}_i \mid \mathcal{F}_t\}. \quad (94)$$

We also define

$$\mathcal{C}_{i,t} := \mathcal{C}_i \cap B_{\epsilon'}(\hat{p}_t), \quad (95)$$

where ϵ' is defined in (72), and

$$\bar{i}_t := \arg \max_{i \in [N]} \mathbb{P}\{\tilde{p}_t \in \mathcal{C}_{i,t} \mid \mathcal{F}_t\}. \quad (96)$$

We define \bar{p}_t as an arbitrary point in $\mathcal{C}_{\bar{i}_t,t}$. Then, we define

$$\bar{\mathcal{A}}_i(t) := \left\{ \|S_i \bar{p}_t - S_i p^*\| \leq \frac{\epsilon}{8} \right\}. \quad (97)$$

Using these notations, the last term in (93) can be decomposed as

$$\begin{aligned} \mathbb{1}[\tilde{\mathcal{A}}_1^c(t)] &\leq \sum_{k=1}^N \mathbb{1}[\bar{p}_t \in \mathcal{C}_k, \tilde{\mathcal{A}}_1^c(t)] \\ &= \sum_{k=1}^N \mathbb{1}[\bar{p}_t \in \mathcal{C}_k, \bar{\mathcal{A}}_k^c(t), \tilde{\mathcal{A}}_1^c(t)] + \sum_{k=1}^N \mathbb{1}[\bar{p}_t \in \mathcal{C}_k, \bar{\mathcal{A}}_k(t), \tilde{\mathcal{A}}_1^c(t)] \\ &\leq \underbrace{\sum_{k=1}^N \mathbb{1}[\bar{p}_t \in \mathcal{C}_k, \bar{\mathcal{A}}_k^c(t)]}_{(C)} + \underbrace{\mathbb{1}[\bar{p}_t \in \mathcal{C}_1, \bar{\mathcal{A}}_1(t), \tilde{\mathcal{A}}_1^c(t)]}_{(D)} + \underbrace{\sum_{k=2}^N \mathbb{1}[\bar{p}_t \in \mathcal{C}_k, \bar{\mathcal{A}}_k(t)]}_{(E)}. \end{aligned} \quad (98)$$

We will bound the expectation of each term in the following and complete the proof of Theorem 3 as

$$\begin{aligned}
\mathbb{E} [\text{Reg}(T)] &= \sum_{i \neq 1} \Delta_i \left(O\left(\frac{1}{\epsilon^2} \log T\right) + O\left(\frac{N}{\epsilon^2} \log T\right) \right) \\
&\quad + \max_{j \in [N]} \Delta_j \left(\sum_{k=1}^N O\left(\frac{NM}{\epsilon^2} \log T\right) + O(1) + \sum_{k=2}^N O(1) \right) \\
&= O\left(\max \left\{ \frac{N \sum_{i \in [N]} \Delta_i}{\epsilon^2}, \frac{N^2 M \max_{i \in [N]} \Delta_i}{\epsilon^2} \right\} \log T \right) \\
&= O\left(\frac{AN^2 M \max_{i \in [N]} \Delta_i}{\Lambda^2} \log T \right), \tag{99}
\end{aligned}$$

where the last transformation follows from the definition of ϵ in (26).

F.2 Analysis for Case (A)

Lemma 18. *For any $i \neq 1$,*

$$\mathbb{E} \left[\sum_{t=1}^T \mathbb{1} \left[i(t) = i, \tilde{\mathcal{A}}_1(t), \mathcal{A}_i(t) \right] \right] \leq \frac{64}{9\epsilon^2} \log T + 2^{M/2}. \tag{100}$$

To prove Lemma 18, we prove the following lemma using Corollary 8 and Lemma 9.

Lemma 19. *For any $0 \leq \xi < 1/2$,*

$$\mathbb{P}\{\tilde{p}_t \in \mathcal{V}_i \mid \mathcal{A}_i(t), N_i(t) > n_i\} \leq \exp\left(-\frac{9}{16}\xi\epsilon^2 n_i\right) (1 - 2\xi)^{-M/2}, \tag{101}$$

where $\mathcal{V}_i := \{p \in \mathcal{C}_i : \|S_1 p - S_1 p^*\| \leq \epsilon\}$.

Proof. Since $\tilde{p}_t \sim \mathcal{N}(\hat{p}_t, B_t^{-1})$ for $\hat{p}_t = B_t^{-1} b_t$, the squared Mahalanobis distance $\|B_t^{1/2}(\tilde{p}_t - \hat{p}_t)\|^2$ follows the chi-squared distribution with M degree of freedom. Therefore, we have

$$\mathbb{P}\{\tilde{p}_t \in \mathcal{V}_i \mid \mathcal{A}_i(t), N_i(t) > n_i\} \leq h\left(\inf_{p \in \mathcal{V}_i} \|B_t^{1/2}(p - \hat{p}_t)\|^2\right), \tag{102}$$

where $h(a) = \mathbb{P}_{X \sim \chi_M^2} \{X \geq a\}$. To use Lemma 9, we check the condition of Lemma 9 is indeed satisfied. First, it is obvious that the assumptions $N_i(t) \geq n_i$ and $\|S_i \hat{p}_t - S_i p^*\| < \epsilon/4$ are satisfied. Besides, $p \in \mathcal{V}_i$ implies $p \in \mathcal{T}_i = \{p \in \mathbb{R}^M : \|S_i p - S_i p^*\| \geq \epsilon\}$ from Corollary 8. Thus, applying Lemma 9 concludes the proof. \square

Proof of Lemma 18. For any $n_i > 0$,

$$\begin{aligned}
&\sum_{t=1}^T \mathbb{1} \left[i(t) = i, \tilde{\mathcal{A}}_1(t), \mathcal{A}_i(t) \right] \\
&= \sum_{t=1}^T \mathbb{1} \left[i(t) = i, \tilde{\mathcal{A}}_1(t), \mathcal{A}_i(t), N_i(t) \leq n_i \right] + \sum_{t=1}^T \mathbb{1} \left[i(t) = i, \tilde{\mathcal{A}}_1(t), \mathcal{A}_i(t), N_i(t) > n_i \right] \\
&\leq n_i + \sum_{t=1}^T \mathbb{1} \left[i(t) = i, \tilde{\mathcal{A}}_1(t), \mathcal{A}_i(t), N_i(t) > n_i \right]. \tag{103}
\end{aligned}$$

The second term is bounded from above as

$$\begin{aligned}
& \mathbb{E} \left[\sum_{t=1}^T \mathbb{1} \left[i(t) = i, \tilde{\mathcal{A}}_1(t), \mathcal{A}_i(t), N_i(t) > n_i \right] \right] \\
&= \sum_{t=1}^T \mathbb{P} \left\{ i(t) = i, \tilde{\mathcal{A}}_1(t), \mathcal{A}_i(t), N_i(t) > n_i \right\} \\
&\leq \sum_{t=1}^T \mathbb{P} \left\{ i(t) = i, \tilde{\mathcal{A}}_1(t) \mid \mathcal{A}_i(t), N_i(t) > n_i \right\} \\
&= \sum_{t=1}^T \mathbb{P} \left\{ i(t) = i, \tilde{\mathcal{A}}_1(t), \tilde{p}_t \in \mathcal{C}_i \mid \mathcal{A}_i(t), N_i(t) > n_i \right\} \quad (i(t) = i \text{ implies } \tilde{p}_t \in \mathcal{C}_i) \\
&\leq \sum_{t=1}^T \mathbb{P} \{ \tilde{p}_t \in \mathcal{V}_i \mid \mathcal{A}_i(t), N_i(t) > n_i \}. \tag{104}
\end{aligned}$$

To obtain an upper bound for $\mathbb{P} \{ \tilde{p}_t \in \mathcal{V}_i \mid \mathcal{A}_i(t), N_i(t) > n_i \}$, we use Lemma 19. By taking $n_i = \frac{16}{9} \frac{1}{\xi \epsilon^2} \log T$ with $\xi = 1/4$, we have

$$\begin{aligned}
\mathbb{E} \left[\sum_{t=1}^T \mathbb{1} \left[i(t) = i, \tilde{\mathcal{A}}_1(t), \mathcal{A}_i(t) \right] \right] &\leq n_i + \sum_{t=1}^T \mathbb{P} \{ \tilde{p}_t \in \mathcal{V}_i \mid \mathcal{A}_i(t), N_i(t) > n_i \} \\
&\leq n_i + \sum_{t=1}^T \exp \left(-\frac{9}{16} \xi \epsilon^2 n_i \right) (1 - 2\xi)^{-M/2} \quad (\text{by Lemma 19}) \\
&= \frac{16}{9} \frac{1}{\xi \epsilon^2} \log T + (1 - 2\xi)^{-M/2} \\
&= \frac{64}{9\epsilon^2} \log T + 2^{M/2}, \tag{105}
\end{aligned}$$

which completes the proof. \square

F.3 Analysis for Case (B)

Lemma 20. For any $i \neq 1$,

$$\mathbb{E} \left[\sum_{t=1}^T \mathbb{1} [i(t) = i, \mathcal{A}_i^c(t)] \right] \leq \frac{256N (\log T + \frac{A}{2} \log 2 + 1)}{\epsilon^2} + \frac{16A^2}{\epsilon^2} \tag{106}$$

The regret in this case can intuitively be bounded because as the round proceeds the event $i(t) = i$ makes $S_i \hat{p}_t$ close to $S_i p^*$, which implies that the expected number of times the event $\mathcal{A}_i^c(t)$ occurs is not large.

Before going to the analysis of Lemma 20, we prove useful inequalities between $\|q_i^{(t)} - S_i p^*\|$, $\|q_i^{(t)} - S_i \hat{p}_t\|$, and $\|S_i \hat{p}_t - S_i p^*\|$.

Lemma 21. Assume $N_i(t) > 0$. Then,

$$\|q_i^{(t)} - S_i \hat{p}_t\|^2 \leq \frac{Z_{\setminus i}}{N_i(t)} + \|q_i^{(t)} - S_i p^*\|^2. \tag{107}$$

Proof. Recall that \hat{p}_t is the maximizer of $\bar{g}_t(p)$, and we have

$$\hat{p}_t = \arg \max_{p \in \mathbb{R}^M} \bar{g}_t(p) = \arg \max_{p \in \mathbb{R}^M} \prod_{i=1}^N \exp \left\{ -\frac{1}{2} N_i(t) \|q_i^{(t)} - S_i p\|^2 \right\} = \arg \min_{p \in \mathbb{R}^M} \sum_{i=1}^N N_i(t) \|q_i^{(t)} - S_i p\|^2. \tag{108}$$

Using this and the definition of $Z_{\setminus i}$, we have

$$\begin{aligned}
N_i(t)\|q_i^{(t)} - S_i\hat{p}_t\|^2 &\leq \sum_{k \in [N]} N_k(t)\|q_k^{(t)} - S_k\hat{p}_t\|^2 \\
&\leq \sum_{k \in [N]} N_k(t)\|q_k^{(t)} - S_k p^*\|^2 \\
&\leq Z_{\setminus i} + N_i(t)\|q_i^{(t)} - S_i p^*\|^2.
\end{aligned} \tag{109}$$

Dividing by $N_i(t)$ on the both sides completes the proof. \square

Lemma 22. *Assume that $\mathcal{A}_i^c(t)$ and $N_i(t) > 0$ hold. Then,*

$$\|q_i^{(t)} - S_i p^*\| > \frac{1}{2} \left(\frac{\epsilon}{4} - \sqrt{\frac{Z_{\setminus i}}{N_i(t)}} \right). \tag{110}$$

Proof. By the triangle inequality,

$$\begin{aligned}
\|q_i^{(t)} - S_i p^*\| &\geq \|S_i\hat{p}_t - S_i p^*\| - \|q_i^{(t)} - S_i\hat{p}_t\| \\
&> \frac{\epsilon}{4} - \sqrt{\frac{Z_{\setminus i}}{N_i(t)} + \|q_i^{(t)} - S_i p^*\|^2} \quad (\text{by } \mathcal{A}_i^c(t) \text{ and Lemma 21}) \\
&\geq \frac{\epsilon}{4} - \sqrt{\frac{Z_{\setminus i}}{N_i(t)}} - \|q_i^{(t)} - S_i p^*\| \quad (\text{by } \sqrt{x+y} \leq \sqrt{x} + \sqrt{y} \text{ for } x, y \geq 0),
\end{aligned} \tag{111}$$

which is equivalent to (110). \square

Proof of Lemma 20. We first bound the expectation conditioned on $Z_{\setminus i}$, and then take the expectation for $Z_{\setminus i}$. Now,

$$\begin{aligned}
&\mathbb{E} \left[\sum_{t=1}^T \mathbb{1}[i(t) = i, \mathcal{A}_i^c(t)] \mid Z_{\setminus i} \right] \\
&= \mathbb{E} \left[\sum_{t=1}^T \mathbb{1} \left[i(t) = i, \mathcal{A}_i^c(t), N_i(t) \leq \frac{64Z_{\setminus i}}{\epsilon^2} \right] \mid Z_{\setminus i} \right] \\
&\quad + \mathbb{E} \left[\sum_{t=1}^T \mathbb{1} \left[i(t) = i, \mathcal{A}_i^c(t), N_i(t) > \frac{64Z_{\setminus i}}{\epsilon^2} \right] \mid Z_{\setminus i} \right] \\
&\leq \frac{64Z_{\setminus i}}{\epsilon^2} + \mathbb{E} \left[\sum_{t=1}^T \mathbb{1} \left[i(t) = i, \mathcal{A}_i^c(t), N_i(t) > \frac{64Z_{\setminus i}}{\epsilon^2} \right] \mid Z_{\setminus i} \right] \quad (i(t) = i \text{ for all } t \in [T]).
\end{aligned} \tag{112}$$

The first term becomes $256N(\log T + \frac{A}{2} \log 2 + 1)/\epsilon^2$ by taking expectation over $Z_{\setminus i}$ using Lemma 10. Then, we bound the second term. From Lemma 22, $\mathcal{A}_i^c(t)$ and $N_i(t) > \frac{64Z_{\setminus i}}{\epsilon^2}$ im-

ply $\|q_i^{(t)} - S_i p^*\| > \epsilon/16$. Therefore,

$$\begin{aligned}
& \mathbb{E} \left[\sum_{t=1}^T \mathbb{1} \left[i(t) = i, \mathcal{A}_i^c(t), N_i(t) > \frac{64Z_{\setminus i}}{\epsilon^2} \right] \middle| Z_{\setminus i} \right] \\
& \leq \mathbb{E} \left[\sum_{t=1}^T \mathbb{1} \left[i(t) = i, \|q_i^{(t)} - S_i p^*\| > \frac{\epsilon}{16} \right] \right] \\
& \leq \mathbb{E} \left[\sum_{t=1}^T \mathbb{1} \left[i(t) = i, \bigcup_{y \in [A]} |(q_i^{(t)})_y - (S_i)_y p^*| > \frac{\epsilon}{16\sqrt{A}} \right] \right] \\
& \leq \mathbb{E} \left[\sum_{y=1}^A \sum_{t=1}^T \mathbb{1} \left[i(t) = i, |(q_i^{(t)})_y - (S_i)_y p^*| > \frac{\epsilon}{16\sqrt{A}} \right] \right] \\
& \leq \mathbb{E} \left[\sum_{y=1}^A \sum_{n_i=1}^T \sum_{t=1}^T \mathbb{1} \left[i(t) = i, N_i(t) = n_i, |(q_i^{(t)})_y - (S_i)_y p^*| > \frac{\epsilon}{16\sqrt{A}} \right] \right] \\
& = \mathbb{E} \left[\sum_{y=1}^A \sum_{n_i=1}^T \mathbb{1} \left[\bigcup_{t=1}^T \left\{ i(t) = i, N_i(t) = n_i, |(q_i^{(t)})_y - (S_i)_y p^*| > \frac{\epsilon}{16\sqrt{A}} \right\} \right] \right] \\
& \quad \text{(The event } \{i(t) = i, N_i(t) = n_i\} \text{ occurs at most once for fixed } n_i.) \\
& \leq \sum_{y=1}^A \sum_{n_i=1}^T \mathbb{P} \left\{ |(q_{i,n_i})_y - (S_i)_y p^*| > \frac{\epsilon}{4\sqrt{A}} \right\} \\
& \leq \sum_{y=1}^A \sum_{n_i=1}^T 2 \exp \left(-2n_i \left(\frac{\epsilon}{4\sqrt{A}} \right)^2 \right) \quad \text{(by Hoeffding's ineq.)} \\
& \leq 2A \sum_{n_i=1}^{\infty} \exp \left(-\frac{n_i \epsilon^2}{8A} \right) \\
& = 2A \frac{1}{\exp \left(\frac{\epsilon^2}{8A} \right) - 1} \\
& \leq 2A \frac{1}{\frac{\epsilon^2}{8A}} \quad \text{(by } e^x \geq 1 + x \text{)} \\
& = \frac{16A^2}{\epsilon^2}. \tag{113}
\end{aligned}$$

By summing up the above argument, the proof is completed. \square

F.4 Analysis for Case (C)

Before going to the analysis of cases (C), (D), and (E), we recall some notations. Recall that

$$P_i(t) = \mathbb{P}\{\tilde{p}_t \in \mathcal{C}_i \mid \mathcal{F}_t\}, \tag{114}$$

$\mathcal{C}_{i,t} = \mathcal{C}_i \cap B_{\epsilon'}(\hat{p}_t)$, $\bar{i}_t = \arg \max_{i \in [N]} \mathbb{P}\{\tilde{p}_t \in \mathcal{C}_{i,t} \mid \mathcal{F}_t\}$, and \bar{p}_t is an arbitrary point in $\mathcal{C}_{\bar{i}_t,t}$. Also recall that

$$\bar{\mathcal{A}}_i(t) = \left\{ \|S_i \bar{p}_t - S_i p^*\| \leq \frac{\epsilon}{8} \right\}. \tag{115}$$

Lemma 23. For any $i \in [N]$,

$$\mathbb{E} \left[\sum_{t=1}^T \mathbb{1}[\bar{p}_t \in \mathcal{C}_i, \bar{\mathcal{A}}_i^c(t)] \right] \leq \frac{N}{p_0} \left(\frac{2^5 M \log T}{\epsilon^2} + e^{\lambda \|p^*\|^2/2} \left(\frac{1}{\lambda T} + \frac{L}{M\lambda} \right)^{M/2} \frac{1}{1 - e^{-\epsilon^2/2^5}} \right). \tag{116}$$

Before proving the above lemma, we give two lemmas.

Lemma 24.

$$\mathbb{P}\{\tilde{p}_t \in \mathcal{C}_{\bar{i}_t} \mid \mathcal{F}_t\} \geq \mathbb{P}\{\tilde{p}_t \in \mathcal{C}_{\bar{i}_t, t} \mid \mathcal{F}_t\} \geq p_0/N, \quad (117)$$

where $p_0 := 1 - h((\lambda\epsilon')^2)$.

Proof. First, we prove

$$\mathbb{P}\left\{\tilde{p}_t \in \bigcup_{i \in [N]} \mathcal{C}_{i, t} \mid \mathcal{F}_t\right\} \geq 1 - h((\lambda\epsilon')^2). \quad (118)$$

This follows from

$$\begin{aligned} \mathbb{P}\left\{\tilde{p}_t \notin \bigcup_{i \in [N]} \mathcal{C}_{i, t} \mid \mathcal{F}_t\right\} &= \mathbb{P}\{\tilde{p}_t \in B_{\epsilon'}(\hat{p}_t) \mid \mathcal{F}_t\} \\ &\leq h\left(\inf_{p \in \{p: \|p - \hat{p}_t\| > \epsilon'\}} \|B_t^{1/2}(p - \hat{p}_t)\|^2\right) \\ &\leq h(\lambda\|p - \hat{p}_t\|^2) \\ &\leq h((\lambda\epsilon')^2). \end{aligned} \quad (119)$$

Using the definition of \bar{i}_t completes the proof. \square

Lemma 25. For any $i \in [N]$, the event $\bar{\mathcal{A}}_i^c(t)$ implies $\|S_i \hat{p}_t - S_i p^*\| \geq \epsilon/16$.

Proof. Using the triangle inequality, we have

$$\begin{aligned} \|S_i \hat{p}_t - S_i p^*\| &\geq \|S_i \bar{p}_t - S_i p^*\| - \|S_i \bar{p}_t - S_i \hat{p}_t\| \\ &\geq \epsilon/8 - \|S_i\| \|\bar{p}_t - \hat{p}_t\| \\ &\geq \epsilon/8 - \|S_i\| \frac{\epsilon}{16 \max_i \|S_i\|} \\ &\geq \epsilon/8 - \epsilon/16 = \epsilon/16. \end{aligned} \quad (120)$$

\square

Proof of Lemma 23. For any n_0 , which is specified later, we have

$$\begin{aligned} &\mathbb{E}\left[\sum_{t=1}^T \mathbb{1}[\bar{p}_t \in \mathcal{C}_i, \bar{\mathcal{A}}_i^c(t)]\right] \\ &= \mathbb{E}\left[\sum_{t=1}^T \mathbb{1}[\bar{p}_t \in \mathcal{C}_i, \bar{\mathcal{A}}_i^c(t), N_i(t) < n_0]\right] + \mathbb{E}\left[\sum_{t=1}^T \mathbb{1}[\bar{p}_t \in \mathcal{C}_i, \bar{\mathcal{A}}_i^c(t), N_i(t) \geq n_0]\right] \end{aligned} \quad (121)$$

The first term can be bounded by $(p_0/N)^{-1} \cdot n_0$ from Lemma 24. The rigorous proof can be obtained by the almost same argument as the following analysis of the second term using Theorem 16.

Then, we will bound the second term. Specifically, we will prove that for $n_0 = \frac{M \log T}{(\epsilon/16)^2}$,

$$\mathbb{E}\left[\sum_{t=1}^T \mathbb{1}[\bar{p}_t \in \mathcal{C}_i, \bar{\mathcal{A}}_i^c(t), N_i(t) \geq n_0]\right] = O(1). \quad (122)$$

First we have

$$\begin{aligned} &\mathbb{E}\left[\sum_{t=1}^T \mathbb{1}[\bar{p}_t \in \mathcal{C}_i, \bar{\mathcal{A}}_i^c(t), N_i(t) \geq n_0]\right] \\ &\leq \sum_{m=n_0}^{\infty} \mathbb{E}\left[\sum_{t=1}^T \mathbb{1}[\bar{p}_t \in \mathcal{C}_i, \bar{\mathcal{A}}_i^c(t), N_i(t) = m]\right]. \end{aligned} \quad (123)$$

Let

$$\tau = \min \{t : \bar{p}_t \in \mathcal{C}_i, \bar{\mathcal{A}}_i^c(t), N_i(t) = m\} \wedge (T + 1) \quad (124)$$

be the first time such that $\bar{p}_t \in \mathcal{C}_i, \bar{\mathcal{A}}_i^c(t)$ and $N_i(t) = m$ occur. Letting $\mathcal{A}_t := \{\bar{p}_t \in \mathcal{C}_i, \bar{\mathcal{A}}_i^c(t), N_i(t) = m\}$, $\mathcal{B}_t := \{i(t) = i\}$ and $P_t := p_0/N$ in Theorem 16, we have

$$\mathbb{E} \left[\sum_{t=1}^T \mathbb{1}[\bar{p}_t \in \mathcal{C}_i, \bar{\mathcal{A}}_i^c(t), N_i(t) = m] \right] \leq \frac{N}{p_0} \mathbb{P}\{\tau \leq T\}. \quad (125)$$

Here $\tau \leq T$ implies that

$$\begin{aligned} \|\hat{p}_\tau - p^*\|_{B_\tau} &= (\hat{p}_\tau - p^*)^\top \left(\lambda I + \sum_{j \in [N]} N_j(\tau) S_j^\top S_j \right) (\hat{p}_\tau - p^*) \\ &\geq m(\hat{p}_\tau - p^*)^\top (S_i^\top S_i) (\hat{p}_\tau - p^*) \\ &= m \|S_i(\hat{p}_\tau - p^*)\|^2 \geq m(\epsilon/16)^2, \end{aligned} \quad (126)$$

where the last inequality follows from Lemma 25. Therefore we have

$$\begin{aligned} \mathbb{E} [\exp(\|\hat{p}_\tau - p^*\|_{B_\tau}^2/2)] &\geq \mathbb{E} [\mathbb{1}\{\tau \leq T\} \exp(\|\hat{p}_\tau - p^*\|_{B_\tau}^2/2)] \\ &\geq \exp(m(\epsilon/16)^2/2) \mathbb{P}\{\tau \leq T\}. \end{aligned} \quad (127)$$

Note that $|B_\tau| \leq |B_T| \leq (1 + TL/M)^M$ for $L = \max_i \sqrt{\text{trace}(S_i^\top S_i)} = \max_i \|S_i\|_F$ by Lemma 10 of Abbasi-Yadkori et al. (2011), where $\|\cdot\|_F$ is the Frobenius norm. Therefore we have

$$\begin{aligned} \mathbb{E} [\exp(\|\hat{p}_\tau - p^*\|_{B_\tau}/2)] &\leq \mathbb{E} \left[\sqrt{|B_\tau|} \cdot \frac{\exp(\|\hat{p}_\tau - p^*\|_{B_\tau}^2/2)}{\sqrt{|B_\tau|}} \right] \\ &\leq (1 + TL/M)^{M/2} \mathbb{E} \left[\frac{\exp(\|\hat{p}_\tau - p^*\|_{B_\tau}^2/2)}{\sqrt{|B_\tau|}} \right] \\ &\leq (1 + TL/M)^{M/2} \mathbb{E} \left[\frac{\exp(\|\hat{p}_0 - p^*\|_{B_0}^2/2)}{\sqrt{|B_0|}} \right] \end{aligned} \quad (128)$$

$$= \left(1 + \frac{TL}{M\lambda}\right)^{M/2} e^{\lambda\|p^*\|^2/2}, \quad (129)$$

where (128) holds since $\frac{\exp(\|\hat{p}_t - p^*\|_{B_t}^2/2)}{\sqrt{|B_t|}}$ is a supermartingale from Lemma 12. Combining (125), (127), and (129), we obtain

$$\begin{aligned} \sum_{m=n_0}^{\infty} \mathbb{E} \left[\sum_{t=1}^T \mathbb{1}[\bar{p}_t \in \mathcal{C}_i, \bar{\mathcal{A}}_i^c(t), N_i(t) = m] \right] &\leq \frac{N}{p_0} \left(\frac{1}{\lambda} + \frac{TL}{M\lambda} \right)^{M/2} e^{\lambda\|p^*\|^2/2} \sum_{m=n_0}^{\infty} e^{-m(\epsilon/16)^2/2} \\ &\leq \frac{N}{p_0} \left(\frac{1}{\lambda} + \frac{TL}{M\lambda} \right)^{M/2} e^{\lambda\|p^*\|^2/2} \frac{e^{-n_0\epsilon^2/2}}{1 - e^{-(\epsilon/16)^2/2}}. \end{aligned} \quad (130)$$

By choosing $n_0 = \frac{M \log T}{(\epsilon/16)^2}$ we obtain the lemma. \square

F.5 Analysis for Case (D)

Lemma 26. For any $i \in [N]$,

$$\mathbb{E} \left[\sum_{t=1}^T \mathbb{1}[\bar{p}_t \in \mathcal{C}_i, \bar{\mathcal{A}}_i(t), \tilde{\mathcal{A}}_i^c(t)] \right] \leq \frac{48}{9} \frac{M + 2}{\epsilon^2} \frac{N}{p_0}. \quad (131)$$

Remark. To prove the regret upper bound, it is enough to prove Lemma 26 only for $i = 1$. However, for the sake of generality, we prove the lemma for any $i \in [N]$.

Before proving Lemma 26, we give two following lemmas.

Lemma 27. For any $i \in [N]$, the event $\bar{\mathcal{A}}_i(t)$ implies $\mathcal{A}_i(t)$.

Proof. Using the triangle inequality, we have

$$\begin{aligned} \|S_i p^* - S_i \hat{p}_t\| &\leq \|S_i p^* - S_i \bar{p}_t\| + \|S_i \bar{p}_t - S_i \hat{p}_t\| \\ &\leq \epsilon/8 + \|S_i\| \cdot \frac{\epsilon}{16 \max_i \|S_i\|} < \epsilon/4, \end{aligned} \quad (132)$$

which completes the proof. \square

Now, Lemma 26 can be intuitively proven because from Lemma 27, $\bar{\mathcal{A}}_i(t)$ implies $\mathcal{A}_i(t)$, and the events $\mathcal{A}_i(t)$ and $\tilde{\mathcal{A}}_i^c(t)$ does not simultaneously occur many times.

Let $t = \sigma_1, \dots, \sigma_m$ be the time of the first m times that the event $\{\bar{p}_t \in \mathcal{C}_i, \mathcal{A}_i(t), N_i(t) = n_i\}$ occurred (not $\{\bar{p}_t \in \mathcal{C}_i, \bar{\mathcal{A}}_i(t), N_i(t) = n_i\}$). In other words, we define

- σ_1 : the first time that $\bar{p}_t \in \mathcal{C}_i, \mathcal{A}_i(t)$ and $N_i(t) = n_i$ occurred
- σ_2 : the second time that $\bar{p}_t \in \mathcal{C}_i, \mathcal{A}_i(t)$ and $N_i(t) = n_i$ occurred
- \dots

Now we prove the following lemma using Lemma 9.

Lemma 28. For any $0 \leq \xi < 1/2$,

$$\mathbb{P}\left\{\tilde{\mathcal{A}}_i^c(t) \mid \mathcal{A}_i(t), \sigma_k = t\right\} \leq \exp\left(-\frac{9}{16}\xi\epsilon^2 n_i\right) (1 - 2\xi)^{-M/2}. \quad (133)$$

Proof. Recall that $\mathcal{T}_i = \{p \in \mathbb{R}^M : \|S_i p - S_i p^*\| > \epsilon\}$. We follow a similar argument as the analysis for Lemma 19. Since $\tilde{p}_t \sim \mathcal{N}(B_t^{-1} b_t, B_t^{-1})$, the squared Mahalanobis distance $\|B_t^{1/2}(p - \hat{p}_t)\|^2$ follows the chi-squared distribution with M degree of freedom. Hence, for $h(a) = \mathbb{P}_{X \sim \chi_M^2}\{X \geq a\}$, we have

$$\mathbb{P}\left\{\tilde{\mathcal{A}}_i^c(t) \mid \mathcal{A}_{i,n_i}, \sigma_k = t\right\} \leq h\left(\inf_{p \in \mathcal{T}_i} \|B_t^{1/2}(p - \hat{p}_t)\|^2\right). \quad (134)$$

Then, Eq. (133) directly follows from Lemma 9. \square

Proof of Lemma 26. From Lemma 27, the event $\bar{\mathcal{A}}_i(t)$ implies $\mathcal{A}_i(t)$. Hence, it is enough to derive the upper bound for

$$\mathbb{E}\left[\sum_{t=1}^T \mathbb{1}\left[\bar{p}_t \in \mathcal{C}_i, \mathcal{A}_i(t), \tilde{\mathcal{A}}_i^c(t)\right]\right] \quad (135)$$

instead of the bound for

$$\mathbb{E}\left[\sum_{t=1}^T \mathbb{1}\left[\bar{p}_t \in \mathcal{C}_i, \bar{\mathcal{A}}_i(t), \tilde{\mathcal{A}}_i^c(t)\right]\right]. \quad (136)$$

Using Lemma 28, we can bound the term for case (D) from above as

$$\begin{aligned}
& \mathbb{E} \left[\sum_{t=1}^T \mathbb{1} \left[\bar{p}_t \in \mathcal{C}_i, \mathcal{A}_i(t), \tilde{\mathcal{A}}_i^c(t) \right] \right] \\
&= \mathbb{E} \left[\sum_{n_i=1}^T \sum_{t=1}^T \mathbb{1} \left[\mathcal{A}_i(t), \tilde{\mathcal{A}}_i^c(t), N_i(t) = n_i \right] \right] \\
&= \sum_{n_i=1}^T \sum_{t=1}^T \mathbb{P} \left\{ \mathcal{A}_i(t), \tilde{\mathcal{A}}_i^c(t), N_i(t) = n_i \right\} \\
&= \sum_{n_i=1}^T \sum_{t=1}^T \sum_{k=1}^T \mathbb{P} \left\{ \mathcal{A}_i(t), \tilde{\mathcal{A}}_i^c(t), \sigma_k = t \right\} \quad (\text{the event } \{\sigma_k = t\} \text{ is exclusive for fixed } n_i) \\
&= \sum_{n_i=1}^T \sum_{t=1}^T \sum_{k=1}^T \mathbb{P} \left\{ \mathcal{A}_i(t), \sigma_k = t \right\} \mathbb{P} \left\{ \tilde{\mathcal{A}}_i^c(t) \mid \mathcal{A}_i(t), \sigma_k = t \right\} \\
&\leq \sum_{n_i=1}^T \sum_{t=1}^T \sum_{k=1}^T \mathbb{P} \left\{ \mathcal{A}_i(t), \sigma_k = t \right\} C e^{-n_i \iota} \quad (\text{by Lemma 28}) \\
&= \sum_{n_i=1}^T C e^{-n_i \iota} \sum_{t=1}^T \sum_{k=1}^T \mathbb{P} \left\{ \mathcal{A}_i(t), \sigma_k = t \right\} \\
&\leq \sum_{n_i=1}^T C e^{-n_i \iota} \sum_{t=1}^T \sum_{k=1}^T \mathbb{P} \left\{ \sigma_k = t \right\} \\
&\leq \sum_{n_i=1}^T C e^{-n_i \iota} \sum_{k=1}^T \mathbb{P} \left\{ \sigma_k \text{ exists} \right\} \\
&\leq \sum_{n_i=1}^T C e^{-n_i \iota} \sum_{k=1}^T \left(1 - \frac{p_0}{N} \right)^{k-1} \quad (\text{by } \tilde{p}_{\sigma_s} \notin \mathcal{C}_i \text{ for } s = 1, \dots, k-1) \\
&\leq 3C \frac{1}{e^\iota - 1} \frac{N}{p_0} \\
&\leq \frac{48}{9} \frac{M+2}{\epsilon^2} \frac{N}{p_0}, \tag{137}
\end{aligned}$$

where $\iota = \frac{9\xi\epsilon^2}{16}$, $C = (1 - 2\xi)^{-\frac{M}{2}}$, and in the last inequality we select the optimal ξ and use $1 + x \leq e^x$. \square

F.6 Analysis for Case (E)

Lemma 29. For any $i \neq 1$,

$$\mathbb{E} \left[\sum_{t=1}^T \mathbb{1} \left[\bar{p}_t \in \mathcal{C}_i, \bar{\mathcal{A}}_i(t) \right] \right] \leq \frac{2^{5M/2+7} \Gamma(M/2+1) e^{\lambda^2 \|p^*\|^2/2}}{\delta^2 \epsilon^{M+2} \lambda^{M/2+1}}, \tag{138}$$

where ϵ is defined in (26) and satisfies $\epsilon \leq \min_{p \in \mathcal{C}_i} \|p - p^*\|/3$, and

$$\delta := \min_{\hat{p}: (L_1 - L_i)^\top \hat{p} \geq 0, \|S_i(\hat{p} - p^*)\| \leq \epsilon/8} \|S_1(\hat{p} - p^*)\|. \tag{139}$$

We prove Lemma 29 using Lemma 12 and Theorem 16.

Remark. The upper bound in (138) goes to infinite when we set $\lambda = 0$, that is, a flat prior is used. However, this is not the essential effect of the prior but just comes from the minimum eigenvalue of B_1 . In fact, we can see from the proof that a similar bound can be obtained for $\lambda = 0$ if we run some deterministic initialization until B_t becomes positive definite.

Proof. We evaluate each term in the summation using Theorem 16 with

$$\begin{aligned}\mathcal{A}_t &= \{\bar{p}_t \in \mathcal{C}_i, \|S_i(\bar{p}_t - p^*)\| \leq \epsilon/8, N_1(t) = n\}, \\ \mathcal{B}_t &= \{\bar{p}_t \in \mathcal{C}_1\}.\end{aligned}\tag{140}$$

for $n \in [T]$. Recall that

$$\bar{g}_t(p) = \frac{1}{\sqrt{(2\pi)^M |B_t^{-1}|}} \exp\left(-\frac{1}{2} \|p - \hat{p}_t\|_{B_t}^2\right)\tag{141}$$

is the probability density function of \hat{p}_t given $\mathcal{F}_t = \{B_t, b_t\}$. Using τ defined in (80), it holds for any $\tau \in [T]$ that

$$\begin{aligned}\mathbb{P}\{\mathcal{B}_\tau | \mathcal{F}_\tau\} &= \mathbb{P}\{\bar{p}_\tau \in \mathcal{C}_1 | \mathcal{F}_\tau\} \\ &= \int_{p \in \mathcal{C}_1} \bar{g}_\tau(p) dp \\ &\geq \int_{p: \|p - p^*\| \leq 3\epsilon} \bar{g}_\tau(p) dp \\ &\geq \sup_{p: \|p - p^*\| \leq 2\epsilon} \int_{p': \|p' - p\| \leq \epsilon} \bar{g}_\tau(p') dp' \\ &\geq \sup_{p: \|p - p^*\| \leq 2\epsilon} \inf_{p': \|p' - p\| \leq \epsilon} \bar{g}_\tau(p') \text{Vol}(\{p'' : \|p'' - p\| \leq \epsilon\}) \\ &= \frac{(\sqrt{\pi}\epsilon)^M}{\Gamma(M/2 + 1)} \sup_{p: \|p - p^*\| \leq 2\epsilon} \inf_{p': \|p' - p\| \leq \epsilon} \bar{g}_\tau(p') \\ &= \frac{(\epsilon/\sqrt{2})^M \sqrt{|B_\tau|}}{\Gamma(M/2 + 1)} \exp\left\{-\frac{1}{2} \left(\inf_{p: \|p - p^*\| \leq 2\epsilon} \sup_{p': \|p' - p\| \leq \epsilon} \|p' - \hat{p}_\tau\|_{B_\tau}^2\right)\right\} \\ &\geq \frac{(\epsilon/\sqrt{2})^M \sqrt{|B_\tau|}}{\Gamma(M/2 + 1)} \exp\left\{-\frac{\|\hat{p}_\tau - p^*\|_{B_\tau}^2 - \epsilon\delta\sqrt{\lambda n}}{2}\right\},\end{aligned}\tag{142}$$

where (142) follows since $\{p : \|p - p^*\| \leq 3\epsilon\} \supset \{p' : \|p' - p_0\| \leq \epsilon\}$ for any p_0 such that $\|p_0 - p^*\| \leq 2\epsilon$, and the last inequality follows from Theorem 15. To apply Theorem 15, we used Lemma 27.

Now we define a stochastic process corresponds to (143) as

$$P_t = \frac{(\epsilon/\sqrt{2})^M \sqrt{|B_t|}}{\Gamma(M/2 + 1)} \exp\left\{-\frac{\|\hat{p}_t - p^*\|_{B_t}^2 - \epsilon\delta\sqrt{\lambda n}}{2}\right\}.\tag{144}$$

Then, by Lemma 12,

$$\begin{aligned}\mathbb{E}[P_{t+1}^{-1} | \mathcal{F}_t] &\leq \frac{\Gamma(M/2 + 1)}{(\epsilon/\sqrt{2})^M} e^{-\epsilon\delta\sqrt{\lambda n}/2} \mathbb{E}\left[\frac{1}{\sqrt{|B_{t+1}|}} \mathbb{E}\left[\exp\left(\frac{\|\hat{p}_t - p^*\|_{B_{t+1}}^2}{2}\right) \middle| \mathcal{F}_t, S_{i(t)}\right] \middle| \mathcal{F}_t\right] \\ &\leq \frac{\Gamma(M/2 + 1)}{(\epsilon/\sqrt{2})^M} e^{-\epsilon\delta\sqrt{\lambda n}/2} \mathbb{E}\left[\frac{1}{\sqrt{|B_t|}} \exp\left(\frac{\|\hat{p}_t - p^*\|_{B_t}^2}{2}\right) \middle| \mathcal{F}_t\right] \\ &= P_t^{-1},\end{aligned}\tag{145}$$

which means that P_t^{-1} is a supermartingale. Therefore we can apply Theorem 16 and obtain

$$\begin{aligned}\mathbb{E}\left[\sum_{t=1}^T \mathbb{1}[\hat{p}_t \in \mathcal{C}_i, \|S_i(\bar{p}_t - p^*)\| \leq \epsilon/8, N_1(t) = n]\right] &\leq \mathbb{E}[\mathbb{1}[\tau \leq T] P_\tau^{-1}] \\ &\leq \mathbb{E}[P_\tau^{-1}] \\ &\leq \mathbb{E}[P_1^{-1}] \\ &= \frac{\Gamma(M/2 + 1) e^{\lambda^2 \|p^*\|^2/2}}{(\epsilon\sqrt{\lambda}/2)^M} e^{-\epsilon\delta\sqrt{\lambda n}/2}.\end{aligned}\tag{146}$$

Finally we have

$$\begin{aligned}
& \mathbb{E} \left[\sum_{t=1}^T \mathbb{1}[\bar{p}_t \in \mathcal{C}_i, \|S_i(\bar{p}_t - p^*)\| \leq \epsilon/8] \right] \\
&= \sum_{n=1}^T \mathbb{E} \left[\sum_{t=1}^T \mathbb{1}[\bar{p}_t \in \mathcal{C}_i, \|S_i(\bar{p}_t - p^*)\| \leq \epsilon/8, N_1(t) = n] \right] \\
&\leq \frac{\Gamma(M/2 + 1)e^{\lambda^2 \|p^*\|^2/2}}{(\epsilon\sqrt{\lambda/2})^M} \sum_{n=1}^{\infty} e^{-\epsilon\delta\sqrt{\lambda n}/2} \\
&\leq \frac{\Gamma(M/2 + 1)e^{\lambda^2 \|p^*\|^2/2}}{(\epsilon\sqrt{\lambda/2})^M} \int_0^{\infty} e^{-\epsilon\delta\sqrt{\lambda x}/2} dx \\
&= \frac{\Gamma(M/2 + 1)e^{\lambda^2 \|p^*\|^2/2}}{(\epsilon\sqrt{\lambda/2})^M} \frac{2}{(\epsilon\delta\sqrt{\lambda/2})^2} \Gamma(2) \\
&= \frac{2^{M/2+3}\Gamma(M/2 + 1)e^{\lambda^2 \|p^*\|^2/2}}{\delta^2 \epsilon^{M+2} \lambda^{M/2+1}}, \tag{147}
\end{aligned}$$

which completes the proof. \square

G Property of Dynamic Pricing Games

In this appendix, we will see a property of dp-easy games.

Proposition 30. *Consider any dp-easy games with $c > -1$. Then, any two actions in the game are neighbors.*

Remark. In section 5, we considered dp-easy games with $c > 0$, but this can be relaxed to $c > -1$ to prove Proposition 30.

Proof. Take any two different actions $j, k \in [N]$ such that $j < k$. From the definition of the loss matrix in dp-easy games, we have $e_j \in \mathcal{C}_j$ and $e_k \in \mathcal{C}_k$.

First, we will find $\alpha \in [0, 1]$ such that

$$\alpha e_j + (1 - \alpha)e_k \in \mathcal{C}_j \cap \mathcal{C}_k. \tag{148}$$

From the definition of the loss matrix, the i -th element of $L(\alpha e_j + (1 - \alpha)e_k) \in \mathcal{P}_M$ is

$$\begin{cases} -i & (1 \leq i \leq j) \\ \alpha c + (1 - \alpha) \cdot (-i) & (j + 1 \leq i \leq k) \\ c & (k < i \leq N) \end{cases}. \tag{149}$$

It is easy to see that the indices which give the minimum value in (149) is j or k . Thus, to achieve the condition (148), the following should be satisfied,

$$-j = \alpha c + (1 - \alpha) \cdot (-k), \tag{150}$$

which is equivalent to

$$\alpha = \frac{k - j}{c + k} (=:\alpha^*). \tag{151}$$

Note that we have $0 \leq \alpha \leq 1$ for any $c > -1$.

Next, we introduce the following definitions.

$$p^{(j,k)} := \alpha^* e_j + (1 - \alpha^*) e_k \in \mathcal{C}_j \cap \mathcal{C}_k, \tag{152}$$

$$\text{Ball}_\epsilon^{(j,k)} := \left\{ p \in \mathcal{P}_M : \|p - p^{(j,k)}\| \leq \epsilon \right\}, \tag{153}$$

$$L^{(x)} := L(p^{(j,k)} + x) \in \mathbb{R}^N. \tag{154}$$

To prove the proposition, it is enough to prove the following: there exists $\epsilon > 0$, $\text{Ball}_\epsilon^{(j,k)} \subset \mathcal{C}_j \cup \mathcal{C}_k$.

To prove this, it is enough to prove that, there exists $\epsilon > 0$,

$$\min_{x \in \mathbb{R}^M: \|x\| \leq \epsilon} \min_{i \in [N] \setminus \{j, k\}} \left((L^{(x)})_i - (L^{(x)})_j \right) \vee \left((L^{(x)})_i - (L^{(x)})_k \right) > 0. \quad (155)$$

We will prove (155) in the following. Take any $i \in [N] \setminus \{j, k\}$ and

$$\epsilon := \min_{i: 1 \leq i < j} \frac{1}{2} \frac{j-i}{\|L_j - L_i\|} \wedge \min_{i: j < i < k} \frac{1}{2} \frac{(1-\alpha^*)(k-i)}{\|L_i - L_k\|} \wedge \min_{i: k < i \leq N} \frac{1}{2} \frac{c+j}{\|L_j - L_i\|}. \quad (156)$$

Note that the ϵ used here is different from the one used in the proof of the regret upper bounds.

Case (A): When $1 \leq i < j$, using Cauchy–Schwarz inequality, we have

$$\begin{aligned} \left((L^{(x)})_i - (L^{(x)})_j \right) \vee \left((L^{(x)})_i - (L^{(x)})_k \right) &\geq (L^{(x)})_i - (L^{(x)})_j \\ &= (-i + L_i^\top x) - (-j + L_j^\top x) \\ &= (j-i) - (L_j - L_i)^\top x \\ &\geq (j-i) - \|L_j - L_k\| \|x\| \\ &\geq (j-i) - \epsilon \|L_j - L_i\| \\ &\geq \frac{1}{2}(j-i) \\ &> 0. \end{aligned} \quad (157)$$

The arguments for cases (B) and (C) follow in the similar manner as case (A).

Case (B): When $j < i < k$, we have

$$\begin{aligned} \left((L^{(x)})_i - (L^{(x)})_j \right) \vee \left((L^{(x)})_i - (L^{(x)})_k \right) &\geq (L^{(x)})_i - (L^{(x)})_k \\ &= \{\alpha^* c + (1-\alpha^*) \cdot (-i) + L_i^\top x\} - \{\alpha^* c + (1-\alpha^*) \cdot (-k) + L_k^\top x\} \\ &= (1-\alpha^*)(k-i) - (L_i - L_k)^\top x \\ &\geq (1-\alpha^*)(k-i) - \epsilon (L_i - L_k)^\top \\ &\geq \frac{1}{2}(1-\alpha^*)(k-i) \\ &> 0. \end{aligned} \quad (158)$$

Case (C): When $k < i \leq N$, we have

$$\begin{aligned} \left((L^{(x)})_i - (L^{(x)})_j \right) \vee \left((L^{(x)})_i - (L^{(x)})_k \right) &\geq (L^{(x)})_i - (L^{(x)})_j \\ &= (c + L_i^\top x) - (-j + L_j^\top x) \\ &\geq c + j - \|L_j - L_i\| \|x\| \\ &\geq c + j - \epsilon \|L_j - L_i\| \\ &\geq \frac{1}{2}(c+j) \\ &> 0. \end{aligned} \quad (159)$$


Summing up the argument for cases (A) to (C), the proof is completed. \square

H Details and Additional Results of Experiments


Here we give the specific values of the opponent's strategy used in Section 5 and show the extended experimental results for performance comparison. Table 2 summarizes the values of opponent's strategy used in this appendix and Section 5. Figure 3 shows the empirical comparison of the proposed algorithms against the benchmark methods, and Figure 4 shows the number of the rejected times. We can see the same tendency as Section 5, that is, TSPM performs the best and the number of rejections does not increase with the time step t .

Table 2: The values of the opponent's strategy.


| # of outcomes M | opponent's strategy p^* |
|-------------------|---------------------------------------|
| 2 | [0.7, 0.3] |
| 3 | [0.5, 0.3, 0.2] |
| 4 | [0.3, 0.3, 0.3, 0.1] |
| 5 | [0.2, 0.3, 0.3, 0.1, 0.1] |
| 6 | [0.2, 0.2, 0.3, 0.1, 0.1, 0.1] |
| 7 | [0.2, 0.2, 0.3, 0.1, 0.1, 0.05, 0.05] |




(a) dp-easy, $N = M = 2$




(b) dp-easy, $N = M = 3$




(c) dp-easy, $N = M = 4$




(d) dp-easy, $N = M = 5$




(e) dp-easy, $N = M = 6$




(f) dp-easy, $N = M = 7$




(g) dp-hard, $N = M = 2$




(h) dp-hard, $N = M = 3$




(i) dp-hard, $N = M = 4$



(j) dp-hard, $N = M = 5$




(k) dp-hard, $N = M = 6$




(l) dp-hard, $N = M = 7$


Figure 3: Regret-round plots of the algorithms. The solid lines indicate the average over 100 independent trials. The thin fillings are the standard error.




(a) dp-easy, $N = M = 2$




(b) dp-easy, $N = M = 3$




(c) dp-easy, $N = M = 4$




(d) dp-easy, $N = M = 5$




(e) dp-easy, $N = M = 6$




(f) dp-easy, $N = M = 7$




(g) dp-hard, $N = M = 2$




(h) dp-hard, $N = M = 3$




(i) dp-hard, $N = M = 4$



(j) dp-hard, $N = M = 5$



(k) dp-hard, $N = M = 6$



(l) dp-hard, $N = M = 7$

Figure 4: The number of rejected times by the accept-reject sampling. The solid lines indicate the average over 100 independent trials after taking moving average with window size 100.