

GREEDY ALGORITHM ALMOST DOMINATES IN SMOOTHED CONTEXTUAL BANDITS*

MANISH RAGHAVAN[†], ALEKSANDRS SLIVKINS[‡], JENNIFER WORTMAN VAUGHAN[§],
AND ZHIWEI STEVEN WU[¶]

Abstract. Online learning algorithms, widely used to power search and content optimization on the web, must balance exploration and exploitation, potentially sacrificing the experience of current users in order to gain information that will lead to better decisions in the future. While necessary in the worst case, explicit exploration has a number of disadvantages compared to the greedy algorithm that always “exploits” by choosing an action that currently looks optimal. We ask under what conditions inherent diversity in the data makes explicit exploration unnecessary. We build on a recent line of work on the smoothed analysis of the greedy algorithm in the linear contextual bandits model. We improve on prior results to show that the greedy algorithm almost matches the best possible Bayesian regret rate of any other algorithm on the same problem instance whenever the diversity conditions hold. The key technical finding is that data collected by the greedy algorithm suffices to simulate a run of any other algorithm. Further, we prove that Bayesian regret of the greedy algorithm is at most $\tilde{O}(T^{1/3})$ in the worst case, where T is the time horizon.

Key words. Multi-armed bandits, linear bandits, greedy algorithm, smoothed analysis, data diversity, Bayesian regret

1. Introduction. Online learning algorithms are a key tool in web search and content optimization, adaptively learning what users want to see. In a typical application, each time a user arrives, the algorithm chooses among various content presentation options (e.g., news articles to display), the chosen content is presented to the user, and an outcome (e.g., a click) is observed. Such algorithms must balance *exploration* (making potentially suboptimal decisions now for the sake of acquiring information that will improve decisions in the future) and *exploitation* (using information collected in the past to make better decisions now). Exploration could degrade the experience of a current user, but improves user experience in the long run. This exploration-exploitation tradeoff is commonly studied in the online learning framework of *multi-armed bandits* [10, 30, 22].

Exploration is widely used, both in theory and in practice. Yet, it has several important disadvantages. First, exploration is wasteful and risky in the short term. It is undesirable for the current user, as something imposed only for the sake of the future users. Exploration may appear unfair, and may even be unethical or illegal in sensitive application domains such as medical decisions. Second, exploration adds a layer of complexity to algorithm design (e.g., see [21, 5]), and its adoption at scale tends to require substantial systems support and buy-in from management [3, 4]. A system that only exploits would typically be cheaper to design and deploy. Third, exploration runs into incentive issues when actions (e.g., which product to buy) are

*A version of our results have been announced in an extended abstract [27], and fleshed out in the technical report [28]. This paper is streamlined compared to [27, 28], focusing on the greedy algorithm. It has been available on arxiv.org/abs/2005.10624 since May 2020. The current version (since Dec’21) allows the random perturbations of context vectors to be correlated across actions, and adds a lower bound result.

[†]Harvard University (mraghavan@seas.harvard.edu). Most of this research has been done while the author was a graduate student at Cornell and an intern at Microsoft Research NYC.

[‡]Microsoft Research, New York (slivkins@microsoft.com).

[§]Microsoft Research, New York (jenn@microsoft.com).

[¶]Carnegie Mellon University (zstevenwu@cmu.edu). Most of this research has been done while the author was a postdoc at Microsoft Research NYC.

controlled by users. In applications such as recommender systems, an algorithm can only encourage exploration via recommendations and other provided information, but users would be reluctant to follow if it is not in their self-interest.¹

An algorithm without explicit exploration, a.k.a. the *greedy algorithm*, always chooses the action that appears optimal according to current estimates of the problem parameters. Further, the greedy algorithm describes self-interested behavior of users in a recommendation system.² The greedy algorithm it is known to perform poorly in a wide range of problem instances, yet it works remarkably well in some examples. A more detailed characterization for whether and when the greedy algorithm performs well — put differently, whether and when exploration is not at all helpful — is an important concern in the study of the exploration-exploitation tradeoff.

A recent line of work [8, 19] analyzes conditions under which inherent diversity in the data makes explicit exploration unnecessary. They consider the *linear contextual bandits* [24, 12, 1], a standard variant of multi-armed bandits appropriate for content personalization scenarios. In particular, Kannan et al. [19] model data diversity via small perturbations of the context vectors, and focus on regret in expectation over these perturbations. They prove that the greedy algorithm achieves expected regret which scales as $\tilde{O}(\sqrt{T})$ in terms of the time horizon T . This is the best regret rate that can be achieved in the worst case (i.e., for all problem instances), even without data diversity assumptions. However, this result does not resolve how the greedy algorithm compares to other algorithms under data diversity conditions, neither in the worst case nor for particular problem instances.

We expand on this line of work. We prove that under similar diversity conditions, the greedy algorithm almost matches the best possible Bayesian regret of *any* algorithm *on the same problem instance*. Known upper bounds on algorithms’ Bayesian regret range from $\text{polylog}(T)$ for some problem instances to $\tilde{O}(\sqrt{T})$ in the worst case, and each of them carries over to the greedy algorithm. Moreover, we prove that the Bayesian regret of the greedy algorithm scales as $\tilde{O}(T^{1/3})$ in the worst case, as long as there are at most $\text{polylog}(T)$ feasible actions in each round.

The data diversity conditions in [19] and this paper are inspired by the smoothed analysis framework of Spielman and Teng [32], who proved that the expected running time of the simplex algorithm is polynomial for perturbations of any initial problem instance (whereas the worst-case running time has long been known to be exponential). Such disparity implies that very bad problem instances are brittle. We find a similar disparity for the greedy algorithm in our setting.

Our contributions. We consider a Bayesian version of linear contextual bandits in which the latent weight vector θ is drawn from a known prior. In each round, an algorithm is presented several actions to choose from, each represented by a *context vector*. The expected reward of an action is a linear product of θ and the corresponding context vector. The tuple of context vectors is drawn independently from a fixed distribution. In the spirit of smoothed analysis, we assume that this distribution has a small amount of jitter. Formally, in each round t the tuple of context vectors is drawn from some fixed distribution, and then a small *perturbation* $\varepsilon_{a,t}$ is added to the context vector for each action a . The basic version adopted in [19] is that each $\varepsilon_{a,t}$ is an independent spherical Gaussian distribution; we call it the *action-*

¹Making exploration compatible with users’ incentives is possible, at least in theory, albeit costly; see [31] for an overview of related research.

²The formal model is as follows: users sequentially choose among available actions, after fully observing what happened with the previous users.

independent perturbation. We allow a more general perturbation model, spelled out in Section 3, which can be correlated across actions, but independent across rounds and coordinates. We are interested in Bayesian regret, i.e., regret in expectation over the Bayesian prior. Following the literature, we are primarily interested in the dependence on the time horizon T .

We focus on a batched version of the greedy algorithm, in which new data arrives to the algorithm’s optimization routine in small batches, rather than every round. This property is essential for our analysis, and easy to implement in practice. As a restriction, it is well-motivated from a practical perspective: in high-volume applications data usually arrives to the “learner” only after a substantial delay [3, 4].

Our main result is that the greedy algorithm matches the Bayesian regret of *any algorithm* up to $\text{polylog}(T)$ factors *for each problem instance*, i.e., fixing the Bayesian prior and the context distribution. This holds for two natural versions of the batched greedy algorithm, Bayesian and frequentist, henceforth called BatchBayesGreedy and BatchFreqGreedy. For BatchBayesGreedy, the chosen action maximizes expected reward according to the Bayesian posterior. The regret bound holds for any Bayesian prior. BatchFreqGreedy estimates θ using ordinary least squares regression and chooses the best action according to this estimate. The regret bound and comes with an extra additive polylogarithmic factor, but is stronger in that the algorithm does not need to know the prior. This result requires a Gaussian prior, which can, however, be very concentrated.

The key insight is that the data collected with perturbed contexts can be used to simulate a run of any other algorithm **ALG**, with the number of rounds scaled down by some factor Y . (This simulation arises only as a technique in the analysis.) It follows that BatchBayesGreedy at each round t knows at least as much as **ALG** after t/Y rounds, so its selection is at least as good as that of **ALG**. To handle the frequentist algorithm, we consider a hypothetical algorithm that receives the same data as BatchFreqGreedy, but chooses actions like BatchBayesGreedy. We analyze this hypothetical algorithm using a similar ‘simulation argument’, and then upper-bound the difference in Bayesian regret compared to BatchFreqGreedy.

Next, we argue that our problem remains difficult despite perturbations. Specifically, we prove that any algorithm achieves Bayesian regret no better than $\tilde{\Omega}(\sqrt{T})$ in the worst case. This holds even if there are at most two feasible actions in each round, and even if perturbation size can be an absolute constant. For this lower bound, perturbations on both actions are completely correlated (i.e., identical).

Finally, we consider action-independent perturbation, and analyze Bayesian regret in the worst case over all Bayesian instances. We prove that LinUCB algorithm [24, 12, 1], a standard algorithm for linear contextual bandits, achieves Bayesian regret $\tilde{O}(K^{2/3} T^{1/3})$ if there are at most K feasible actions in each round. Consequently, a similar regret bound holds for BatchBayesGreedy and BatchFreqGreedy. The $\tilde{O}(T^{1/3})$ regret rate is a mathematical curiosity, as we are not aware of any published regret bounds between \sqrt{T} and $\text{polylog}(T)$; however, it is unclear if this regret rate is optimal. Regardless, we conclude that action-independent perturbation is substantially “easier” compared to the general case, in light of the lower bound stated above.

Map of the paper. The paper continues with related work (Section 2), model and preliminaries (Section 3), precise statements of the results (Sections 4), and a detailed discussion of the techniques (Section 5). The analysis is spelled out in Sections 6-8, for, resp., the greedy algorithms, the lower bound, and LinUCB algorithm. Some tools are moved to the appendix so as not to interrupt the flow.

2. Related Work. The greedy algorithm works well in some examples, and badly in some others. This has been a folklore knowledge for decades, and it has been confirmed in extensive recent experiments [9]. One way to formalize a *negative* result is to consider a Bayesian prior over 2-armed bandit instances. Then, with positive-constant probability over the prior, the greedy algorithm fails to explore the best arm, and therefore incurs a positive-constant regret in each round (see Chapter 11 in [30]). This is a very general result, as it holds for any Bayesian prior.

Positive results on the greedy algorithm. Most related to ours are papers by Kannan et al. [19] and Bastani et al. [8].³ Both study the greedy algorithm in linear contextual bandits with data diversity conditions. In particular, Kannan et al. [19] introduce the perturbation model adopted in our paper, focusing on the special case of action-independent perturbations. We provide a detailed comparison below.

We substantially improve over the $\tilde{O}(\sqrt{T})$ regret bound from Kannan et al. [19]: our main result applies per-instance rather than only in the worst-case, and allows perturbations to be correlated across actions. Going back to action-independent perturbations, as in [19], we also improve the worst-case bound on Bayesian regret to $\tilde{O}(T^{1/3})$ when there are only $\text{polylog}(T)$ feasible actions in each round. However, these improvements come at the cost of some additional assumptions. First, we consider Bayesian regret, whereas their regret bound holds for each realization of θ . Second, they allow the context vectors to be chosen by an adversary before the perturbation is applied. Third, they extend their analysis to a somewhat more general model, in which there is a separate latent weight vector for every action (which amounts to a different model of perturbations). However, this extension relies on the greedy algorithm being initialized with a substantial amount of data.

Bastani et al. [8] show that the greedy algorithm achieves logarithmic regret in a version of linear contextual bandits that is incomparable to ours in several important ways. First, the actions share a common context vector in each round, but are parameterized by different latent vectors. Then, playing a given arm reveals no information about the other arms, which makes their problem more difficult compared to ours. To compensate for this difficulty, they posit a strong assumption on data diversity: essentially, that the distribution of contexts is approximately symmetric around the origin. It follows that for any pair of arms, each arm is better than the other for a constant fraction of rounds. In contrast, our model allows the context distribution to be arbitrary, subject to a relatively small perturbation; in particular, the same action could be the best action in all rounds. Third, a version of Tsybakov’s *margin condition* is assumed, which is known to substantially reduce regret rates in bandit problems (see, e.g., [29]). Instead, we assume Gaussian perturbations, allowing us to make a finer-grained simulation argument that the greedy algorithm is instance-optimal.

Acemoglu et al. [2] and Immorlica et al. [17]⁴ analyze the greedy algorithm from the economics perspective, providing positive results for “greedy” self-interested behavior of users in a recommendation system. Acemoglu et al. [2] study heterogeneous users with private types. In our terms, it is a version of contextual bandits in which the current context is not observed in the future rounds. Among other results, they prove that the greedy algorithm works well in this setting, under strong heterogeneity assumptions incomparable with yours. In particular, they postulate that a user arriv-

³An early version of Bastani et al. [8] (v2, Jun’17) is prior work relative to this paper. In particular, it focuses on the special case of two actions. Subsequent versions are concurrent work.

⁴The early version of Acemoglu et al. [2] (from Nov’17) is prior work relative to this paper; subsequent versions are concurrent work. Immorlica et al. [17] is subsequent work.

ing in each round inherently prefers each arm with (at least) a constant probability. Immorlica et al. [17] constructs a data disclosure policy which reveals to each user the history for a predetermined subset of prior users. In our terms, they consider a bandit problem with constantly many arms, and a greedy algorithm operating on limited data as prescribed by this policy. They prove that such algorithm attains regret rates that are near-optimal for any bandit algorithm.

Technical aspects. Any contextual bandit algorithm can be simulated using data collected by any other contextual bandit algorithm which independently randomizes actions in each round [25, 16]. Essentially, the required number of samples is inversely proportional to the smallest sampling probability across arms. While not very complicated technically, this approach works for contextual bandits (linear or not) without any additional assumptions. However, this approach fails in our setting because our data is collected by a *deterministic* algorithm. Instead, our simulation uses a different approach, which relies on random perturbation of contexts.

The work on “batched bandit algorithms” [26] assumes that the rounds are partitioned into “batches” so that the algorithm cannot use the data from the current batch. The goal is to achieve efficient exploration despite this restriction. In contrast, we focus on the greedy algorithm rather than exploration, and invoke the batch property as a voluntary feature which helps in the analysis.

Linear contextual bandits. The problem was introduced in [24], motivated by personalized news recommendations. The non-contextual version stems from [6]. Both versions have been studied extensively, see books [10, 22] for background.

Algorithm LinUCB, which we discuss in Section 8, implements ‘optimism under uncertainty’, a common paradigm for problems with explore-exploit tradeoff. The algorithm was defined in [24], and analyzed in [12, 1]. (A non-contextual version of LinUCB was introduced earlier in [6], and analyzed in [14].) The details of the algorithm differ subtly between the papers; we focus on the version from [1].

LinUCB achieves regret $\tilde{O}(d\sqrt{T})$, where d is the dimension, for any number of actions [1]. Any algorithm suffers regret $\Omega(d\sqrt{T})$ in the worst case [14]. LinUCB has been observed to perform well even when the problems are not linear [20].

3. Our Model and Preliminaries. We consider the model of *linear contextual bandits* [24, 12, 1]. A learner operates over T timesteps (a.k.a. rounds), where T is a known time horizon.⁵ Each round t proceeds as follows. There are at most K actions available, a.k.a. *arms*. Denote the action set as $A_t \subset \{1, \dots, K\}$. Each action $a \in A_t$ is associated with a *context vector* $x_{a,t} \in \mathbb{R}^d$, which may contain features of the action and/or the round. We assume that the tuple of context vectors $(x_{a,t} : a \in A_t)$ is drawn independently from a fixed distribution D . The learner observes this tuple, selects an action $a_t \in A_t$, and observes reward r_t . We assume that r_t is drawn independently from some distribution determined by the chosen context vector $x_{a_t,t}$, and the expected reward is linear in this vector. More precisely, we let $r_{a,t}$ be the reward of each action $a \in A_t$ if this action is chosen in round t (so that $r_t = r_{a_t,t}$), and posit an unknown vector $\theta \in \mathbb{R}^d$ such that

$$\mathbb{E}[r_{a,t} \mid x_{a,t}] = \theta^\top x_{a,t} \quad \text{for any round } t \text{ and action } a \in A_t.$$

So far, it is a standard *frequentist* formulation of linear contextual bandits. It is determined by time horizon T , dimension d , number K of feasible actions per round, context distribution D , and the latent vector θ .

⁵For intuition, each round typically corresponds to an interaction with a new user.

We consider a natural *Bayesian* version, where θ drawn from a known Bayesian prior \mathcal{P} . Thus, a problem instance consists of parameters T, d, K , context distribution D , and Bayesian prior \mathcal{P} . The prior can be arbitrary unless specified otherwise.

The learner strives to maximize the expected total reward over T rounds, or $\sum_{t=1}^T \mathbb{E} [\theta^\top x_{a,t}]$. We focus on *regret*, a standard performance measure which compares the learner to the all-knowing benchmark: a hypothetical algorithm that knows the best action in each round. Formally, we define the best context vector in round t as

$$x_t^* \in \arg \max_{x \in \{x_{a,t} : a \in A_t\}} \theta^\top x,$$

i.e., a context vector which achieves the highest expected reward. Next,

$$(3.1) \quad \text{Regret}(T) = \sum_{t=1}^T \theta^\top x_t^* - \theta^\top x_{a_t,t}.$$

Expected regret is defined as the expectation of (3.1) over the context vectors, the rewards, and the algorithm's random seed. We are mainly interested in *Bayesian regret*, where the expectation is taken over all of the above and the prior over θ .

Data diversity. We model data diversity via the following process, called *perturbed context generation*. Fix round t , and recall that A_t denotes the set of available actions. First, a tuple $(\mu_{a,t} \in \mathbb{R}^d : a \in A_t)$ of *mean context vectors* is drawn independently from some fixed distribution D_μ over $(\mathbb{R}^d)^{|A_t|}$. Then for each action $a \in A_t$, the context vector is $x_{a,t} = \mu_{a,t} + \varepsilon_{a,t}$, where $\varepsilon_{a,t} \in \mathbb{R}^d$ is a zero-mean perturbation vector. Marginally, each perturbation vector $\varepsilon_{a,t}$ is distributed as $\mathcal{N}(0, \rho_{a,t} \cdot I)$, a spherical Gaussian distribution over \mathbb{R}^d with zero mean and per-coordinate standard deviation $\rho_{a,t} > 0$. We consider two basic versions for correlation across actions:

- *action-independent perturbation*: each perturbation vector $\varepsilon_{a,t}$ is an independent draw from $\mathcal{N}(0, \rho I)$.
- *fully-action-correlated perturbation*: $\varepsilon_{a,t} = \varepsilon_t$ for all arms $a \in A_t$, where the (common) perturbation vector ε_t is an independent draw from $\mathcal{N}(0, \rho I)$.

Our guarantees deteriorate if *perturbation size* ρ is very small.

We allow a more general model of action-correlation which interpolates between these two extremes defined above. We have a set covering \mathcal{F}_t of A_t , i.e., a family \mathcal{F}_t of subsets of A_t whose union equals A_t . For each subset $S \in \mathcal{F}_t$, we have a *base perturbation* $\varepsilon_{S,t}^{\text{base}} \in \mathbb{R}^d$, which is an independent draw from $\mathcal{N}(0, \rho_{S,t} \cdot I)$, for some $\rho_{S,t} > 0$. We sum up the base perturbations over all relevant subsets $S \in \mathcal{F}_t$:

$$(3.2) \quad \varepsilon_{a,t} = \sum_{S \in \mathcal{F}_t : a \in S} \varepsilon_{S,t}^{\text{base}}.$$

The paradigmatic case is that $\mathcal{F}_t = \mathcal{F}$ for all rounds t , but we allow it to change over time. Likewise, the paradigmatic case is that $\rho_{a,t} = \rho$ for all arms a and rounds t , but we allow $\rho_{S,t}$ can vary for different subsets S and rounds t . In the latter case, we summarize the dependence on the perturbations via the *perturbation size*

$$\rho := \min_{t \in [T], a \in A_t} \max_{S \in \mathcal{F}_t : a \in S} \rho_{S,t}.$$

Note that we “use” the largest relevant perturbation for a given arm-round pair.

We make several technical assumptions. First, the distribution D_μ is such that each context vector has bounded 2-norm, i.e., $\|\mu_{a,t}\|_2 \leq 1$. It can be arbitrary otherwise. Second, the perturbation size needs to be sufficiently small compared to the

dimension d , $\rho \leq 1/\sqrt{d}$. Third, the realized reward $r_{a,t}$ for each action a and round t is $r_{a,t} = x_{a,t}^\top \theta + \eta_{a,t}$, the mean reward $x_{a,t}^\top \theta$ plus standard Gaussian noise $\eta_{a,t} \sim \mathcal{N}(0, 1)$.⁶

Batched greedy algorithms. We write x_t for $x_{a_t,t}$, the context vector chosen at time t . The history up to round t is the tuple $h_t = ((x_1, r_1), \dots, (x_t, r_t))$.

For the batch version of the greedy algorithm, time is divided in batches of Y consecutive rounds each. When forming its estimate of the optimal action at round t , the algorithm may only use the history up to the last round of the previous batch, denoted t_0 . We consider both Bayesian and frequentist versions, called BatchBayesGreedy and BatchFreqGreedy.

BatchBayesGreedy forms a posterior over θ using prior \mathcal{P} and history h_{t_0} . In round t it chooses the action that maximizes reward in expectation over this posterior. This is equivalent to choosing

$$(3.3) \quad a_t = \arg \max_a x_{a,t}^\top \theta_t^{\text{bay}}, \quad \text{where } \theta_t^{\text{bay}} := \mathbb{E}[\theta \mid h_{t_0}].$$

BatchFreqGreedy does not rely on any knowledge of the prior. It chooses the best action according to the least squares estimate of θ , denoted θ_t^{fre} , computed with respect to history h_{t_0} :

$$(3.4) \quad a_t = \arg \max_a x_{a,t}^\top \theta_t^{\text{fre}}, \quad \text{where } \theta_t^{\text{fre}} = \arg \min_{\theta'} \sum_{\tau=1}^{t_0} ((\theta')^\top x_\tau - r_\tau)^2.$$

Empirical covariance matrix. Fix round t . Let $X_t \in \mathbb{R}^{t \times d}$ be the *context matrix*, a matrix whose rows are vectors $x_1, \dots, x_t \in \mathbb{R}^d$. A $d \times d$ matrix

$$Z_t := \sum_{\tau=1}^t x_\tau x_\tau^\top = X_t^\top X_t,$$

called the *empirical covariance matrix*, is an important concept in some of the prior work on linear contextual bandits (e.g., [1, 19]), as well as in this paper.

A note on notation. We adopt a common (albeit slightly non-standard) convention that $\tilde{O}(\cdot)$ hides $\text{polylog}(T)$ factors, regardless of the expression in brackets. In particular, the expression in brackets is always interpreted as a function of T .

4. Statement of the Results. We prove that in expectation over the random perturbations, both greedy algorithms favorably compare to any other algorithm. For any specific problem instance, both algorithms match the Bayesian regret of any algorithm on that particular instance up to polylogarithmic factors. We state the theorem in terms of the main relevant parameters T , K , d , Y , and ρ .

THEOREM 4.1. *With perturbed context generation, there is some $Y_0 = \text{polylog}(d, T)/\rho^2$ such that with batch duration $Y \geq Y_0$, the following holds. Fix any bandit algorithm, and let $R_0(T)$ be its Bayesian regret on a particular problem instance. Then on that same instance,*

- (a) *BatchBayesGreedy has Bayesian regret at most $Y \cdot R_0(T/Y) + \tilde{O}(1/T)$,*
- (b) *Suppose prior \mathcal{P} is a multivariate Gaussian distribution with invertible covariance matrix Σ , and the eigenvalues of Σ are at least ρ^{-4}/T . Then BatchFreqGreedy has Bayesian regret at most $Y \cdot R_0(T/Y) + \tilde{O}(C_\Sigma \sqrt{d}/\rho^2)$, where C_Σ is determined by Σ .*

⁶Our analysis can be easily extended to handle reward noise of fixed variance, i.e., $\eta_{a,t} \sim \mathcal{N}(0, \sigma^2)$. BatchFreqGreedy would not need to know σ . BatchBayesGreedy would need to know either Σ and σ or just Σ/σ^2 .

REMARK 4.2. The dependence on the covariance matrix Σ in Theorem 4.1(b) is $C_\Sigma = \sqrt{\lambda_{\max}} + 1/\sqrt{\lambda_{\min}}$, where λ_{\max} , λ_{\min} are, resp., the largest and smallest eigenvalues of Σ . (This comes from Theorem 6.11.) The dependence on λ_{\min} captures the deterioration in Bayesian regret if the prior is very concentrated. For example, if prior \mathcal{P} is independent over the components of θ , with variance $\sigma^2 \leq 1$ in each component, then $\lambda_{\max} \leq 1$ and $\lambda_{\min} = \sigma$, so that $C_\Sigma = 1 + 1/\sqrt{\sigma}$.

Next, we prove that $\tilde{\Omega}(\sqrt{T})$ lower bound on Bayesian regret holds even under perturbed context generation. We posit the most difficult regime: $K = d = 2$ and constant ρ . The lower bound focuses on fully-action-correlated perturbation.

THEOREM 4.3. Consider perturbed context generation with fully-action-correlated perturbation. Any algorithm achieves Bayesian regret no better than $\tilde{\Omega}(\sqrt{T})$ for some problem instance. This holds even if there are only $d = 2$ dimensions, $K = 2$ feasible actions in each round, and perturbation size ρ is an absolute constant.

Finally, we focus on worst-case Bayesian regret. We consider action-independent perturbation, and posit that prior \mathcal{P} is a multivariate Gaussian with mean vector $\bar{\theta}$ and invertible covariance matrix Σ .

THEOREM 4.4. Consider perturbed context generation with action-independent perturbation. Assume that all eigenvalues of the covariance matrix Σ are at most 1,⁷ and the mean vector satisfies $\|\bar{\theta}\|_2 \geq 1 + \sqrt{3 \log T}$. Then

- (a) LinUCB algorithm, with appropriate parameter settings, has Bayesian regret $\mathbf{reg} := \tilde{O}(d^2 K^{2/3} / \rho^2) \times T^{1/3}$.
- (b) (Follows from Theorem 4.1.) Let Y_0 from Theorem 4.1 be the batch duration. Then BatchBayesGreedy has Bayesian regret at most $\tilde{O}(\mathbf{reg})$. Likewise, BatchFreqGreedy has Bayesian regret at most $\tilde{O}\left(\mathbf{reg} + \sqrt{d} \rho^{-2} / \sqrt{\lambda_{\min}(\Sigma)}\right)$, provided that $\lambda_{\min}(\Sigma) \geq \rho^{-4}/T$.

REMARK 4.5. The assumption $\|\bar{\theta}\|_2 \geq 1 + \sqrt{3 \log T}$ in Theorem 4.4 can be replaced with an assumption that the dimension d is sufficiently large: $d \geq \log T / \log \log T$.

5. Overview: Key Techniques. The key idea is to show that, with perturbed context generation, BatchBayesGreedy collects data that is informative enough to “simulate” the history of contexts and rewards from the run of any other algorithm **ALG** over fewer rounds. This implies that it remains competitive with **ALG** since it has at least as much information and makes myopically optimal decisions.

Let us formulate what we mean by “simulation”. We want to use the data collected from a single batch B in order to simulate the reward for any one context x , and we want to accomplish this without knowing the latent vector θ . More formally, we use the tuple $h_B = ((x_t, r_t) : t \in B)$, which we call the *batch history*. We are interested in the randomized function $\mathbf{Rew}_\theta(\cdot)$ that takes a context vector x and outputs an independent random sample from $\mathcal{N}(\theta^\top x, 1)$; this is the realized reward for an action with context vector x . So, the function $\mathbf{Rew}_\theta(\cdot)$ is what we want to simulate using batch history h_B . To do so, we construct a fixed function g such that $g(x, h_B)$ is distributed identically to $\mathbf{Rew}_\theta(x)$, for any fixed context vector x ; the randomness in $g(x, h_B)$ comes from h_B .

This definition needs to be refined, so as to simulate independent noise in rewards. Indeed, randomness in h_B comes from several sources: context arrivals, algorithm’s

⁷In particular, if the prior \mathcal{P} is independent across the coordinates of θ , then the variance in each coordinate is at most 1.

decisions, realization of θ , and observed rewards. Relying on the first three sources introduces dangerous correlations. To rule them out, we require our simulation $g(x, h_B)$ to have the same distribution as $\text{Rew}_\theta(x)$, even if we condition on the context vectors x_t previously chosen by the algorithm during this batch, i.e., on the tuple $(x_t : t \in B)$.

DEFINITION 5.1. *Consider batch B in the execution of BatchBayesGreedy. Batch history h_B can simulate $\text{Rew}_\theta(\cdot)$ up to radius $R > 0$ if there exists a function $g : \{\text{context vectors}\} \times \{\text{batch histories } h_B\} \rightarrow \mathbb{R}$ such that $g(x, h_B)$ is distributed identically to $\text{Rew}_\theta(x)$, conditional on the tuple $(x_t : t \in B)$, for all θ and all context vectors $x \in \mathbb{R}^d$ with $\|x\|_2 \leq R$.*

Let us comment on how it may be possible to simulate $\text{Rew}_\theta(x)$. For intuition, suppose that $x = \frac{1}{2}x_1 + \frac{1}{2}x_2$. Then $(\frac{1}{2}r_1 + \frac{1}{2}r_2 + \xi)$ is distributed as $\mathcal{N}(\theta^\top x, 1)$ if ξ is drawn independently from $\mathcal{N}(0, \frac{1}{2})$. Thus, we can define $g(x, h) = \frac{1}{2}r_1 + \frac{1}{2}r_2 + \xi$ in Definition 5.1. We generalize this idea and show that a batch history can simulate Rew_θ with high probability as long as the batch size Y is sufficiently large.

LEMMA 5.2. *With perturbed context generation, there is some $Y_0 = \text{polylog}(d, T)/\rho^2$ and $R = O(\rho\sqrt{d \log(TKd)})$ such that with probability at least $1 - T^{-2}$ any batch history from BatchBayesGreedy can simulate $\text{Rew}_\theta(\cdot)$ up to radius R , as long as $Y \geq Y_0$.*

To prove this, we ensure that the data collected in batch B are sufficiently diverse. To define “sufficiently diverse”, let the *batch context matrix*, denoted X_B , be a matrix which comprises the context vectors $(x_t : t \in B)$. Namely, X_B is the $Y \times d$ matrix whose rows are vectors x_t , $t \in B$, in the order of increasing t . Similarly to the “empirical covariance matrix”, we define the *batch covariance matrix* as

$$(5.1) \quad Z_B := X_B^\top X_B = \sum_{t \in B} x_t x_t^\top.$$

We think of data diversity in terms of the minimal eigenvalue of Z_B : the larger it is, the more diverse is the data. And we prove that the minimal eigenvalue of Z_B is sufficiently large whenever $Y \geq Y_0$.

If the batch history of an algorithm can simulate Rew_θ , the algorithm has enough information to simulate the outcome of a fresh round of any other algorithm **ALG**. We use a coupling argument in which we couple a run of BatchBayesGreedy with a slowed-down run of **ALG**, and prove that the former accumulates at least as much information as the latter, and therefore the Bayesian-greedy action choice is, in expectation, at least as good as that of **ALG**. This yields the regret bounds for BatchBayesGreedy in Theorems 4.1 and 4.4.

We use the same technique to handle BatchFreqGreedy. To treat both greedy algorithms at once, we define a template that unifies them. A bandit algorithm is called *batch-greedy-style* if it divides the timeline in batches of Y consecutive rounds each, in each round t chooses some estimate θ_t of θ , based only on the data from the previous batches, and then chooses the best action according to this estimate, so that $a_t = \arg \max_a \theta_t^\top x_{a,t}$. Lemma 5.2 extends to any batch-greedy-style algorithm.

The analysis of BatchFreqGreedy requires an additional step. We consider a hypothetical batch-greedy-style algorithm which separates data collection and reward collection: it receives feedback based on the actions of BatchFreqGreedy, but collects rewards based on the (batched) Bayesian-greedy selection rule. We analyze this hypothetical algorithm using Lemma 5.2, and then argue that its Bayesian regret cannot be much smaller than that of BatchFreqGreedy. Intuitively, this is because the two algorithms form very similar estimates of θ , differing only in the fact that the hypo-

thetical algorithm uses the prior \mathcal{P} as well as the data. Due to this similarity, we show that the numerical difference between the two estimates at time t is at most $\tilde{O}(1/t)$, even though either estimate is typically $\Omega(1/\sqrt{t})$ away from θ . This adds up to a maximal difference of $O(\log T)$ in Bayesian regret between the two algorithms, and completes our regret bounds for BatchFreqGreedy.

6. Analysis: Greedy Algorithms. We present the proofs for all results on greedy algorithms. This section is structured as follows. In Section 6.1, we quantify the diversity of data collected by batch-greedy-style algorithms, assuming perturbed context generation. In Section 6.2, we show that a sufficiently “diverse” batch history suffices to simulate the reward for any given context vector, in the sense of Definition 5.1. Jointly, these two subsections imply that any batch history generated by a batch-greedy-style algorithm can simulate rewards with high probability, as long as the batch size is sufficiently large. Section 6.3 builds on this foundation to derive regret bounds for BatchBayesGreedy. The crux is that the history collected by BatchBayesGreedy suffices to simulate a “slowed-down” run of any other algorithm. This analysis extends to a version of BatchFreqGreedy equipped with a Bayesian-greedy prediction rule (and tracks the performance of the prediction rule). Finally, Section 6.4 derives the regret bounds for BatchFreqGreedy, by comparing the prediction-rule version of BatchFreqGreedy with BatchFreqGreedy itself.

Preliminaries. We assume perturbed context generation in this section, without further mention. We use definitions for batch greedy algorithms from Section 5: batch-greedy-style algorithm, batch history, batch context matrix, and batch covariance matrix. Throughout, we will use the following parameters as a shorthand:

$$\begin{aligned}\delta_R &= T^{-2} \\ \hat{R} &= \rho\sqrt{2\log(2TKd/\delta_R)} \\ R &= 1 + \hat{R}\sqrt{d}.\end{aligned}$$

Recall that ρ denotes perturbation size, and d is the dimension. The meaning of \hat{R} and R is that they are high-probability upper bounds on the perturbations and the contexts, respectively. More formally, by Lemma A.5 we have:

$$(6.1) \quad \Pr \left[\|\varepsilon_{a,t}\|_\infty \leq \hat{R} : \text{ for all arms } a \text{ and all rounds } t \right] \geq 1 - \delta_R$$

$$(6.2) \quad \Pr [\|x_{a,t}\|_2 \leq R : \text{ for all arms } a \text{ and all rounds } t] \geq 1 - \delta_R$$

6.1. Data Diversity under Perturbations. We are interested in the diversity of data collected by batch-greedy-style algorithms, assuming perturbed context generation. Informally, the observed contexts x_1, x_2, \dots should cover all directions in order to enable good estimation of the latent vector θ . Following Kannan et al. [19], we quantify data diversity via the minimal eigenvalue of the empirical covariance matrix Z_t . More precisely, we are interested in proving that $\lambda_{\min}(Z_t)$ is sufficiently large. We adapt some tools from [19], extending them from to action-correlated perturbations, and then derive some improvements for batch-greedy-style algorithms.

6.1.1. Tools from [19]. Kannan et al. [19] prove that for action-independent perturbations, $\lambda_{\min}(Z_t)$ grows linearly in time t , assuming t is sufficiently large.

LEMMA 6.1 (implicit in [19]). *Consider action-independent perturbations. Fix any batch-greedy-style algorithm. Consider round $t \geq \tau_0$, where $\tau_0 = 160 \frac{R^2}{\rho^2} \log \frac{2d}{\delta}$.*

$\log T$. Then for any realization of θ , with probability $1 - \delta$

$$\lambda_{\min}(Z_t) \geq \frac{\rho^2 t}{32 \log T}.$$

Proof. The claimed conclusion follows from an argument inside the proof of Lemma B.1 from [19], plugging in $\lambda_0 = \frac{\rho^2}{2 \log T}$. This argument applies for any $t \geq \tau'_0$, where $\tau'_0 = \max\left(32 \log \frac{2}{\delta}, 160 \frac{R^2}{\rho^2} \log \frac{2d}{\delta} \cdot \log T\right)$. We observe that $\tau'_0 = \tau_0$ since $R \geq \rho$. \square

Rather than use Lemma 6.1 directly, we extract a key portion in its proof, encapsulate it as a standalone lemma, and extend it to our model of action-correlated perturbations. Specifically, recalling that $Z_t := \sum_{\tau=1}^t x_\tau x_\tau^\top$, we zero in on the expected contribution of a single round t .

LEMMA 6.2 (implicit in [19] for action-independent perturbations). *Fix any batch-greedy-style algorithm, and the latent vector θ . Fix round t . Assume $T \geq 4K$. Condition on the event that all perturbations $\varepsilon_{a,t}$ are at most \hat{R} , denote it with \mathcal{E} . Fix round t . Then with probability at least $1/4$,*

$$\lambda_{\min}(\mathbb{E}[x_t x_t^\top \mid h_{t-1}, \mathcal{E}]) \geq \frac{\rho^2}{2 \log T}.$$

The proof of this lemma is assembled from several pieces in the analysis in [19], which extend naturally to our perturbation model. Qualitatively, our goal is as follows: we need to argue that the context vector of the chosen arm at each round has sufficient variance in expectation (or equivalently, $\lambda_{\min}(\mathbb{E}[x_t x_t^\top])$ is large) for us to “learn” about all components of θ . Because of the perturbations, each context vector independently has high variance; however, we need to show that this remains true even conditioning on an arm being chosen. Note that this conditioning should intuitively *reduce* variance, since it selects for arms that have been perturbed in the direction of $\hat{\theta}_t$, all else equal. One way to view this conditioning is to consider the rewards of the best arm a and the second-best arm a' : if a' has a **high** expected reward, then the perturbation applied to a must have a large component in the direction of $\hat{\theta}_t$ in order for a to be chosen over a' , and so conditioned on the realized context vector $x_{a',t}$, we would expect $x_{a,t}$ to have little variance in the direction of $\hat{\theta}_t$. On the other hand, if a' has **low** expected reward, then the perturbation applied to a is less constrained, allowing it to have more variance. Our analysis will argue that the latter case is sufficiently common: we’ll define $\hat{c}_{a,t}$ to be the expected reward of the second-best arm a' , and argue that $\hat{c}_{a,t}$ is “low” with constant probability, and as a result, $x_{a,t}$ has high variance in expectation.

Proof of Lemma 6.2. Let $\hat{\theta}_t$ be the algorithm’s estimate for θ at time t . For ease of exposition, assume that $\hat{\theta}_t = \left[\|\hat{\theta}_t\|_2 \ 0 \ \dots \ 0\right]^\top$. This is w.l.o.g. because we can just rotate the space.

Fix arm a , and let $\varepsilon_{a,t}^*$ be a largest-size perturbation affecting this arm. Formally, let $S = S_a^*$ be a subset $a \in S \in \mathcal{F}_t$ which maximizes perturbation size $\rho_{S,t}$, and let $\varepsilon_{a,t}^* = \varepsilon_{S,t}^{\text{base}}$. Let $\mu_{a,t}^* := x_{a,t} - \varepsilon_{a,t}^*$ be the result of applying all base perturbations to arm a , except $\varepsilon_{a,t}^*$. Further, let $\varepsilon_{a,t}^{-*}$ be the tuple of all *other* base perturbations, including those not affecting arm a :

$$\varepsilon_{a,t}^{-*} := (\varepsilon_{S,t}^{\text{base}} : S \in \mathcal{F}_t \setminus \{S_a^*\}).$$

Similarly to [19], define the “estimated best arm” among those not affected by $\varepsilon_{a,t}^*$:

$$\hat{c}_{a,t} = \max_{\text{arms } a' \in A_t \setminus S_a^*} \hat{\theta}_t^\top x_{a',t}.$$

Let us say that round t is “good” (meaning the expected reward for any other arm isn’t too large) for arm a if

$$(6.3) \quad \hat{c}_{a,t} \leq \hat{\theta}_t^\top \mu_{a,t}^* + \rho \sqrt{2 \log T} \|\hat{\theta}_t\|_2.$$

Our argument from here will take two steps:

1. Given that arm a was selected, $\hat{c}_{a,t}$ is good with constant probability
2. Given that $\hat{c}_{a,t}$ is good, $x_{a,t}$ has sufficient variance (equivalently, $\lambda_{\min}(\mathbb{E}[x_{a,t}x_{a,t}^\top])$ is large).

First, we argue that for each arm $a \in A_t$,

$$(6.4) \quad \Pr[\hat{c}_{a,t} \text{ is good for } a \mid a_t = a, \mathcal{E}] \geq \frac{1}{4}.$$

To do so, we adapt the proof of Lemma 3.4 from [19] to handle our general perturbation model. Let $A_a = \{t : \Pr[a_t \text{ is good}] \geq \frac{1}{2}\}$, and let $B_a = \{t : \Pr_\varepsilon[a_t = a] \geq \frac{2}{T}\}$. Let S_a be the rounds at which a was chosen, i.e., $S_a = \{t : a_t = a\}$. We’ll argue that if $t \in B_a \cap S_a$, then $t \in A_a$. As a result, $\{t \in S_a : t \notin A_a\} \subseteq \{t \in S_a : t \notin B_a\}$. Since our goal is to upper-bound $\{t \in S_a : t \notin A_a\}$, it suffices to upper-bound $\{t \in S_a : t \notin B_a\}$. Consider some $t \in B_a$. Then, taking all probabilities over all of the perturbations, and denoting the right-hand side of (6.3) with Λ , we have:

$$\begin{aligned} \Pr[\hat{c}_{a,t} \text{ is not good for } a \mid a_t = a] &= \Pr[\hat{c}_{a,t} > \Lambda \mid a_t = a] \\ &= \frac{\Pr[\hat{c}_{a,t} > \Lambda \text{ and } a_t = a]}{\Pr[a_t = a]} \\ (t \in B_a) \quad &\leq \frac{\Pr[\hat{c}_{a,t} > \hat{\theta}_t^\top \mu_{a,t}^* + \rho \sqrt{2 \log T} \|\hat{\theta}_t\|_2 \text{ and } a_t = a]}{2/T} \\ &= \frac{\Pr[\hat{c}_{a,t} > \Lambda \text{ and } \hat{c}_{a,t} < \hat{\theta}_t^\top (\mu_{a,t}^* + \varepsilon_{a,t}^*)]}{2/T} \\ &= \frac{\Pr[\Lambda < \hat{c}_{a,t} < \hat{\theta}_t^\top (\mu_{a,t}^* + \varepsilon_{a,t}^*)]}{2/T} \\ &\leq \frac{\Pr[\Lambda < \hat{\theta}_t^\top (\mu_{a,t}^* + \varepsilon_{a,t}^*)]}{2/T} \\ &= \frac{\Pr[\rho \sqrt{2 \log T} \|\hat{\theta}_t\|_2 < \hat{\theta}_t^\top \varepsilon_{a,t}^*]}{2/T} \\ &\leq 1/2, \end{aligned}$$

where the last step follows from standard tail bounds on a Gaussian. Thus, if $t \in B_a$, then $\Pr[\hat{c}_{a,t} \text{ is good for } a \mid a_t = a] \geq \frac{1}{2}$.

Finally, let $\mathbb{P}_{\text{base}}[\cdot]$ be the probability taken over the randomness in all base perturbations for round t (including those not affecting arm a). Let C_t be the set of arms at round t with probability at most $2/T$ of being chosen over the randomness of the perturbation, i.e., $C_t = \{a \mid \mathbb{P}_{\text{base}}[a_t = a] \leq 2/T\}$. Then,

$$\mathbb{P}_{\text{base}}[t \notin B_a \mid a_t = a] = \mathbb{P}_{\text{base}}[a_t \in C_t] \leq \sum_{a \in C_t} \mathbb{P}_{\text{base}}[a_t = a] \leq \frac{2}{T} |C_t| \leq \frac{2K}{T} \leq 1/2.$$

Since by assumption $T \geq 4K$, (6.4) follows.

Second, we argue that for each arm $a \in A_t$,

$$(6.5) \quad \lambda_{\min}(\mathbb{E}[x_{a,t}x_{a,t}^\top \mid a_t = a \text{ and } \hat{c}_{a,t} \text{ is good}]) \geq \frac{\rho^2}{2 \log T}.$$

To prove (6.5), we adapt the proof of Lemma 3.2 from [19] to handle action-correlated perturbations. Fix arm $a \in A_t$. For brevity, let's use notation for the matrix $M_{a,t} = x_{a,t}x_{a,t}^\top$ and the event $\mathcal{E}_t = \{a_t = a \text{ and } \hat{c}_{a,t} \text{ is good}\}$. Let $\mathbb{E}_{\text{base}}[\cdot]$ be the expectation over the randomness in all base perturbations for round t (including those not affecting arm a). Then

$$\begin{aligned} \lambda_{\min}(\mathbb{E}_{\text{base}}[M_{a,t} \mid \mathcal{E}_t]) &= \lambda_{\min}(\mathbb{E}_{\text{base}}[\mathbb{E}_{\text{base}}[M_{a,t} \mid \mathcal{E}_t, \varepsilon_{a,t}^{-*}] \mid \mathcal{E}_t]) \\ &\geq \mathbb{E}_{\text{base}}[\lambda_{\min}(\mathbb{E}_{\text{base}}[M_{a,t} \mid \mathcal{E}_t, \varepsilon_{a,t}^{-*}]) \mid \mathcal{E}_t] \end{aligned}$$

by superadditivity of the minimum eigenvalue.

Thus, it suffices to fix $\hat{c}_{a,t}$ and show a lower bound on

$$\begin{aligned} &\lambda_{\min}(\mathbb{E}_{\text{base}}[M_{a,t} \mid \mathcal{E}_t, \varepsilon_{a,t}^{-*}]) \\ &= \lambda_{\min}\left(\mathbb{E}_{\text{base}}\left[M_{a,t} \mid \varepsilon_{a,t}^{-*} \text{ and } \hat{\theta}_t^\top x_{a,t} \geq \hat{c}_{a,t} \text{ and } \hat{c}_{a,t} \leq \hat{\theta}_t^\top \mu_{a,t}^* + \rho\sqrt{2 \log T}\|\hat{\theta}_t\|_2\right]\right) \\ &= \lambda_{\min}\left(\mathbb{E}_{\text{base}}\left[M_{a,t} \mid \hat{\theta}_t^\top \varepsilon_{a,t}^* \geq \hat{c}_{a,t} - \hat{\theta}_t^\top \mu_{a,t}^* \text{ and } \hat{c}_{a,t} - \hat{\theta}_t^\top \mu_{a,t}^* \leq \rho\sqrt{2 \log T}\|\hat{\theta}_t\|_2\right]\right) \\ &= \lambda_{\min}\left(\mathbb{E}_{\text{base}}\left[M_{a,t} \mid \hat{\theta}_t^\top \varepsilon_{a,t}^* \geq b \text{ and } b \leq \rho\sqrt{2 \log T}\|\hat{\theta}_t\|_2\right]\right) \\ &\geq \frac{\rho^2}{2 \log T}, \end{aligned}$$

where $b = \hat{c}_{a,t} - \hat{\theta}_t^\top \mu_{a,t}^*$. The second line follows from the third because conditioned on $\hat{c}_{a,t}, \varepsilon_{a,t}^{-*}$ provides no additional information about $M_{a,t}$. The final inequality is referred to as the “diversity condition” in [19], and they prove that truncated Gaussian noise (recall that we are conditioning on the event that all the perturbations are component-wise bounded by \hat{R}) satisfies this condition in Lemma 3.7 of [19], with parameters $r = \rho\sqrt{2 \log T}$ and $\lambda_0 = \frac{\rho^2}{2 \log T}$.

This completes the proof of (6.5). The lemma follows from (6.4) and (6.5). \square

Let θ_t^{fre} be the BatchFreqGreedy estimate for θ at time t , as defined in (3.4). We are interested in quantifying how the quality of this estimate improves over time. [19] prove, essentially, that the distance between θ_t^{fre} and θ scales as $\sqrt{t}/\lambda_{\min}(Z_t)$.

LEMMA 6.3 ([19]). *Consider any round t in the execution of BatchFreqGreedy. Let t_0 be the last round of the previous batch. For any θ and any $\delta > 0$, with probability $1 - \delta$,*

$$\|\theta - \theta_t^{\text{fre}}\|_2 \leq \frac{\sqrt{t_0 \cdot 2dR \log \frac{d}{\delta}}}{\lambda_{\min}(Z_{t_0})}.$$

6.1.2. Some improvements. We focus on batch covariance matrix Z_B of a given batch in a batch-greedy-style algorithm. We would like to prove that $\lambda_{\min}(Z_B)$ is sufficiently large with high probability, as long as the batch size Y is large enough. The analysis from [19] (a version of Lemma 6.1) would apply, but only as long as the

batch size is least as large as the τ_0 from the statement of Lemma 6.1. We derive a more efficient version, essentially shaving off a factor of 8.⁸

LEMMA 6.4. *Fix a batch-greedy-style algorithm and any batch B in the execution of this algorithm. Fix $\delta > 0$ and assume that the batch size Y is at least*

$$(6.6) \quad Y_0 := \left(\frac{R}{\rho}\right)^2 \frac{8e^2}{(e-1)^2} \left(1 + \log \frac{2d}{\delta}\right) \log(T) + \frac{4e}{e-1} \log \frac{2}{\delta}.$$

Condition on the event that all perturbations in this batch are upper-bounded by \hat{R} , more formally:

$$\mathcal{E}_B = \{\|\varepsilon_{a,t}\|_\infty \leq \hat{R} : \text{ for all arms } a \text{ and all rounds } t \text{ in } B\}.$$

Further, condition on the latent vector θ and the history h before batch B . Then

$$(6.7) \quad \Pr \left[\lambda_{\min}(Z_B) \geq R^2 \mid \mathcal{E}_B, h, \theta \right] \geq 1 - \delta.$$

The probability in (6.7) is over the randomness in context arrivals and rewards in batch B .

The improvement over Lemma 6.1 comes from two sources: we use a tail bound on the sum of geometric random variables instead of a Chernoff bound on a binomial random variable, and we derive a tighter application of the eigenvalue concentration inequality of [33].

Proof. Let t_0 be the last round before batch B . Recalling (5.1), let

$$W_B = \sum_{t=t_0+1}^{t_0+Y} \mathbb{E} [x_t x_t^\top \mid h_{t-1}]$$

be a similar sum over the expected per-round covariance matrices. Assume $Y \geq Y_0$

The proof proceeds in two steps: first we lower-bound $\lambda_{\min}(Z_B)$, and then we show that it implies (6.7). Denoting $m = R^2 \frac{e}{e-1} (1 + \log \frac{2d}{\delta})$, we claim that

$$(6.8) \quad \Pr [\lambda_{\min}(W_B) < m \mid \mathcal{E}_B, h] \leq \frac{\delta}{2}.$$

To prove this, observe that W_B 's minimum eigenvalue increases by at least $\lambda_0 = \rho^2/(2 \log T)$ with probability at least $1/4$ each round by Lemma 6.2, where the randomness is over the history, i.e., the sequence of (context, reward) pairs. If we want it to go up to m , this should take $4m/\lambda_0$ rounds in expectation. However, we need it to go to m with high probability. Notice that this is dominated by the sum of m/λ_0 geometric random variables with parameter $\frac{1}{4}$. We'll use the following bound from [18]: for $X = \sum_{i=1}^n X_i$ where $X_i \sim \text{Geom}(p)$ and any $c \geq 1$,

$$\Pr[X \geq c\mathbb{E}[X]] \leq \exp(-n(c-1-\log c)).$$

Because we want the minimum eigenvalue of W_B to be m , we need $n = m/\lambda_0$, so $\mathbb{E}[X] = 4m/\lambda_0$. Choose $c = \left(1 + \frac{\lambda_0}{m} \log \frac{2}{\delta}\right) \frac{e}{e-1}$. By Corollary A.13,

$$c - 1 - \log c \geq \frac{e-1}{e} \cdot c - 1 = \frac{\lambda_0}{m} \log \frac{2}{\delta}.$$

⁸Essentially, the factor of 160 in Lemma 6.1 is replaced with factor $\frac{8e^2}{(e-1)^2} < 20.022$ in (6.6).

Therefore,

$$\Pr[X \geq c\mathbb{E}[X]] \leq \exp\left(-n \cdot \frac{\lambda_0}{m} \log \frac{2}{\delta}\right) = \left(\frac{\delta}{2}\right)^{n \cdot \lambda_0/m} = \frac{\delta}{2}$$

Thus, with probability $1 - \frac{\delta}{2}$, $\lambda_{\min}(W_B) \geq m$ as long as the batch size Y is at least

$$\frac{e}{e-1} \left(1 + \frac{\lambda_0}{m} \log \frac{2}{\delta}\right) \cdot \mathbb{E}[X] = \frac{4e}{e-1} \left(\frac{m}{\lambda_0} + \log \frac{2}{\delta}\right) = Y_0.$$

This completes the proof of (6.8).

To derive (6.7) from (6.8), we proceed as follows. Consider the event

$$\mathcal{E} = \{ \lambda_{\min}(Z_B) \leq R^2 \text{ and } \lambda_{\min}(W_B) \geq m \}.$$

Letting $\alpha = 1 - R^2/m$ and rewriting R^2 as $(1 - \alpha)m$, we use a concentration inequality from [33, Theorem 1.1] (following [19, Lemma A.3]) to guarantee that

$$\Pr[\mathcal{E} \mid \mathcal{E}_B, h] \leq d \left(e^\alpha (1 - \alpha)^{1-\alpha} \right)^{-m/R^2}.$$

Then, using the fact that $x^x \geq e^{-1/e}$ for all $x > 0$, we have

$$\begin{aligned} \Pr[\mathcal{E} \mid \mathcal{E}_B, h] &\leq d \left(e^{1-R^2/m-1/e} \right)^{-m/R^2} = d e^{-(m-R^2-m/e)/R^2} \\ &= d \exp \left(-\frac{\left(\frac{e-1}{e}\right)m}{R^2} + 1 \right) \leq \frac{\delta}{2}, \end{aligned}$$

since $m \geq \frac{e}{e-1} R^2 (1 + \log \frac{2d}{\delta})$. Finally, observe that, omitting the conditioning on \mathcal{E}_B, h , we have:

$$\Pr[\lambda_{\min}(Z_B) \leq R^2] \leq \Pr[\mathcal{E}] + \Pr[\lambda_{\min}(W_B) < m] \leq \frac{\delta}{2} + \frac{\delta}{2} = \delta. \quad \square$$

6.2. Reward Simulation with a Diverse Batch History. We consider reward simulation with a batch history, in the sense of Definition 5.1. We show that a sufficiently “diverse” batch history suffices to simulate the reward for any given context vector. Coupled with the results of Section 6.1, it follows that batch history generated by a batch-greedy-style algorithm can simulate rewards as long as the batch size is sufficiently large.

Let us recap the definition of reward simulation (Definition 5.1). Let $\text{Rew}_\theta(\cdot)$ be a randomized function that takes a context x and outputs an independent random sample from $\mathcal{N}(\theta^\top x, 1)$. In other words, this is the realized reward for an action with context vector x .

DEFINITION 6.5. *Consider batch B in the execution of a batch-greedy-style algorithm. Batch history h_B can simulate $\text{Rew}_\theta(\cdot)$ up to radius $R > 0$ if there exists a function $g : \{\text{context vectors}\} \times \{\text{batch histories } h_B\} \rightarrow \mathbb{R}$ such that $g(x, h_B)$ is identically distributed to $\text{Rew}_\theta(x)$ conditional on the batch context matrix, for all θ and all context vectors $x \in \mathbb{R}^d$ with $\|x\|_2 \leq R$.*

Note that we do not require the function g to be efficiently computable. We do not require algorithms to compute g ; a mere existence of such function suffices for our analysis.

The result in this subsection does not rely on the “greedy” property. Instead, it applies to all “batch-style” algorithms, defined as follows: time is divided in batches of

Y consecutive rounds each, and the action at each round t only depends on the history up to the previous batch. The data diversity condition is formalized as $\{\lambda_{\min}(Z_B) \geq R^2\}$; recall that it is a high-probability event, in a precise sense defined in Lemma 6.4. The result is stated as follows:

LEMMA 6.6. *Fix a batch-style algorithm and any batch B in the execution of this algorithm. Assume the batch covariance matrix Z_B satisfies $\lambda_{\min}(Z_B) \geq R^2$. Then batch history h_B can simulate Rew_θ up to radius R .*

Proof. Let us construct a suitable function g for Definition 6.5. Fix a context vector $x \in \mathbb{R}^d$ with $\|x\|_2 \leq R$. Let r_B be the vector of realized rewards in batch B , i.e., $r_B = (r_t : \text{rounds } t \text{ in } B) \in \mathbb{R}^Y$. Define

$$(6.9) \quad g(x, h_B) = w_B^\top r_B + \mathcal{N}(0, 1 - \|w_B\|_2^2), \text{ where } w_B = X_B Z_B^{-1} x \in \mathbb{R}^Y.$$

Recall that the variance of the reward noise is 1. (We can also handle a more general version in which the variance of the reward noise is σ^2 . Then the noise variance in (6.9) should be $\sigma^2(1 - \|w_B\|_2^2)$, with essentially no modifications throughout the rest of the proof.)

Note that w_B is well-defined: indeed, Z_B is invertible since $\lambda_{\min}(Z_B) \geq R^2 > 0$. In the rest of the proof we show that g is as needed for Definition 6.5.

First, we will show that for any $x \in \mathbb{R}^d$ such that $\|x\|_2 \leq R$, the weights $w_B \in \mathbb{R}^t$ as defined above satisfy $X_B^\top w_B = x$ and $\|w_B\|_2 \leq 1$. Then, we'll show that if each $r_\tau \sim \mathcal{N}(\theta^\top x_\tau, 1)$, then $r_B^\top w_B + \mathcal{N}(0, 1 - \|w_B\|_2^2) \sim \mathcal{N}(\theta^\top x, 1)$.

Trivially, we have

$$X_B^\top w_B = X_B^\top X_B (X_B^\top X_B)^{-1} x = x$$

as desired. We must now show that $\|w_B\|_2^2 \leq 1$. Note that

$$\|w_B\|_2^2 = w_B^\top w_B = w_B^\top X_B Z_B^{-1} x = x^\top Z_B^{-1} x = \|x\|_{Z_B^{-1}}^2$$

where $\|v\|_M^2$ simply denotes $v^\top M v$. Thus, it is sufficient to show that $\|x\|_{Z_B^{-1}}^2 \leq 1$. Since $\|x\|_2 \leq R$ and $\lambda_{\min}(Z_B) \geq R^2$, we have by Lemma A.9

$$Z_B \succeq R^2 I \succeq x x^\top.$$

By Lemma A.10, we have

$$I \succeq Z_B^{-1/2} x x^\top Z_B^{-1/2}.$$

Let $z = Z_B^{-1/2} x$, so $I \succeq z z^\top$. Again by Lemma A.9, $\lambda_{\max}(z z^\top) = z^\top z$. This means that

$$1 \geq z^\top z = (Z_B^{-1/2} x)^\top Z_B^{-1/2} x = x^\top Z_B^{-1} x = \|x\|_{Z_B^{-1}}^2 = \|w_B\|_2^2$$

as desired. Finally, observe that

$$r_B^\top w_B = (X_B \theta + \eta)^\top w_B = \theta^\top X_B^\top w_B + \eta^\top w_B = \theta^\top x + \eta^\top w_B$$

where $\eta \sim \mathcal{N}(0, I)$ is the noise vector. Notice that $\eta^\top w_B \sim \mathcal{N}(0, \|w_B\|_2^2)$, and therefore, $\eta^\top w_B + \mathcal{N}(0, 1 - \|w_B\|_2^2) \sim \mathcal{N}(0, 1)$. Putting this all together, we have

$$r_B^\top w_B + \mathcal{N}(0, 1 - \|w_B\|_2^2) \sim \mathcal{N}(\theta^\top x, 1)$$

and therefore D can simulate E for any x up to radius R . \square

6.3. Regret Bounds for BatchBayesGreedy. We apply the tools from Sections 6.1 and 6.2 to derive regret bounds for BatchBayesGreedy. On a high level, we prove that the history collected by BatchBayesGreedy suffices to simulate a “slowed-down” run of any other algorithm ALG_0 . Therefore, when it comes to choosing the next action, BatchBayesGreedy has at least as much information as ALG_0 , so its Bayesian-greedy choice cannot be worse than the choice made by ALG_0 .

Our analysis extends to a more general scenario which is useful for the analysis of BatchFreqGreedy. We formulate and prove our results for this scenario directly. We consider an extended bandit model which separates data collection and reward collection. Each round t proceeds as follows: the algorithm observes available actions and the context vectors for these actions, then it chooses *two* actions, a_t and a'_t , and observes the reward for the former but not the latter. We refer to a'_t as the “prediction” at round t . We will refer to an algorithm in this model as a bandit algorithm (which chooses actions a_t) with “prediction rule” that chooses the predictions a'_t . More specifically, we will be interested in an arbitrary batch-greedy-style algorithm with prediction rule given by BatchBayesGreedy, as per (3.3) on 7. We assume this prediction rule henceforth. We are interested in *prediction regret*: a version of regret (3.1) if actions a_t are replaced with predictions a'_t :

$$(6.10) \quad \text{PReg}(T) = \sum_{t=1}^T \theta^\top x_t^* - \theta^\top x_{a'_t, t}$$

where x_t^* is the context vector of the best action at round t , as in (3.1). More precisely, we are interested in *Bayesian prediction regret*, the expectation of (6.10) over everything: the context vectors, the rewards, the algorithm’s random seed, and the prior over θ .

Thus, the main theorem of this subsection is formulated as follows:

THEOREM 6.7. *Consider perturbed context generation. Let ALG be an arbitrary batch-greedy-style algorithm whose batch size is at least Y_0 from (6.6). Fix any bandit algorithm ALG_0 , and let $R_0^\top(T)$ be regret of this algorithm on a particular problem instance \mathcal{I} . Then on the same instance, ALG has Bayesian prediction regret*

$$(6.11) \quad \mathbb{E} [\text{PReg}^\top(T)] \leq Y \cdot \mathbb{E} [R_0^\top(T/Y)] + \tilde{O}(1/T).$$

Proof sketch. We use a t -round history of ALG to simulate a (t/Y) -round history of ALG_0 . More specifically, we use each batch in the history of ALG to simulate one round of ALG_0 . We prove that the simulated history of ALG_0 has exactly the same distribution as the actual history, for any θ . Since ALG predicts the Bayesian-optimal action given the history (up to the previous batch), this action is at least as good (in expectation over the prior) as the one chosen by ALG_0 after t/Y rounds. \square

As a corollary, we obtain regret bounds for BatchBayesGreedy in Theorems 4.1 and 4.4. We take ALG to be BatchBayesGreedy. For Theorem 4.4(b), we take ALG_0 to be LinUCB. Thus:

COROLLARY 6.8. *In the setting of Theorem 6.7, BatchBayesGreedy has Bayesian regret at most $Y \cdot \mathbb{E} [R_0(T/Y)] + \tilde{O}(1/T)$ on problem instance \mathcal{I} . Further, under the assumptions of Theorem 4.4, BatchBayesGreedy has Bayesian regret at most $\tilde{O}(d^2 K^{2/3} T^{1/3}/\rho^2)$ on all instances.*

We also obtain a similar regret bound on the Bayesian prediction regret of BatchFreqGreedy, which is essential for Section 6.4.

COROLLARY 6.9. *In the setting of Theorem 6.7, BatchFreqGreedy has Bayesian prediction regret (6.11).*

In the remainder of this subsection, we prove Theorem 6.7.

Throughout the proof, we condition on the event that all perturbations are bounded by \hat{R} , more precisely, on the event

$$(6.12) \quad \mathcal{E}_1 = \left\{ \|\varepsilon_{a,t}\|_\infty \leq \hat{R} : \text{for all arms } a \text{ and all rounds } t \right\}.$$

Recall that \mathcal{E}_1 is a high-probability event, by (6.1). We also condition on the event

$$\mathcal{E}_2 = \left\{ \lambda_{\min}(Z_B) \geq R^2 : \text{for each batch } B, \right\}$$

where Z_B is the batch covariance matrix, as usual. Conditioned on \mathcal{E}_1 , this too is a high-probability event (this follows by Lemma 6.4, plugging in δ/T and taking a union bound over all batches).

We will prove that **ALG** satisfies

$$(6.13) \quad \mathbb{E} [\text{PReg}^T(T) \mid \mathcal{E}_1, \mathcal{E}_2] \leq Y \cdot \mathbb{E} [R_0^T(\lceil T/Y \rceil) \mid \mathcal{E}_1, \mathcal{E}_2],$$

where the expectation is taken over everything: the context vectors, the rewards, the algorithm's random seed, and the prior over θ . Then we take care of the “failure event” $\overline{\mathcal{E}_1 \cap \mathcal{E}_2}$.

Before we prove (6.13), let us argue about using the history of **ALG** to simulate a (shorter) run of **ALG**₀. Fix round t . We use a t -round history of **ALG** to simulate a $\lfloor t/Y \rfloor$ -round run of **ALG**₀, where Y is the batch size in **ALG**. Stating this formally requires some notation. Let A_t be the set of actions available in round t , and let $\text{con}_t = (x_{a,t} : a \in A_t)$ be the corresponding tuple of contexts. Let **CON** be the set of all possible context tuples, more precisely, the set of all finite subsets of \mathbb{R}^d . Let h_t and h_t^0 denote, resp., the t -round history of **ALG** and **ALG**₀. Let \mathcal{H}_t denote the set of all possible t -round histories. Note that h_t and h_t^0 are random variables which take values on \mathcal{H}_t . We want to use history h_t to simulate history $h_{\lfloor t/Y \rfloor}^0$. Thus, the simulation result is stated as follows:

LEMMA 6.10. *Fix round t and let $\sigma = (\text{con}_1, \dots, \text{con}_{\lfloor t/Y \rfloor})$ be the sequence of context arrivals up to and including round $\lfloor t/Y \rfloor$. Then there exists a “simulation function”*

$$\text{sim} = \text{sim}_t : \mathcal{H}_t \times \text{CON}_{\lfloor t/Y \rfloor} \rightarrow \mathcal{H}_{\lfloor t/Y \rfloor}$$

such that the simulated history $\text{sim}(h_t, \sigma)$ is distributed identically to $h_{\lfloor t/Y \rfloor}^0$, conditional on sequence σ , latent vector θ , and events $\mathcal{E}_1, \mathcal{E}_2$.

Proof. Throughout this proof, condition on events \mathcal{E}_1 and \mathcal{E}_2 . Generically, $\text{sim}(h_t, \sigma)$ outputs a sequence of pairs $\{(x_\tau, r_\tau)\}_{\tau=1}^{\lfloor t/Y \rfloor}$, where x_τ is a context vector and r_τ is a simulated reward for this context vector. We define $\text{sim}(h_t, \sigma)$ by induction on τ with base case $\tau = 0$. Throughout, we maintain a run of algorithm **ALG**₀. For each step $\tau \geq 1$, suppose **ALG**₀ is simulated up to round $\tau - 1$, and the corresponding history is recorded as $((x_1, r_1), \dots, (x_{\tau-1}, r_{\tau-1}))$. Simulate the next round in the execution of **ALG**₀ by presenting it with the action set A_τ and the corresponding context tuple con_τ . Let x_τ be the context vector chosen by **ALG**₀. The corresponding reward r_τ is constructed using the τ -th batch in h_t , denote it with B . By Lemmas 6.4 and 6.6, the batch history h_B can simulate a single reward, in the sense of Definition 6.5. In particular, there exists a function $g(x, h_B)$ with the required

properties (recall that it is explicitly defined in (6.9)). Thus, we define $r_\tau = g(x_\tau, h_B)$, and return r_τ as a reward to ALG_0 . This completes the construction of $\text{sim}(h_t, \sigma)$. The distribution property of $\text{sim}(h_t, \sigma)$ is immediate from the construction. \square

Proof of Equation (6.13). We argue for each batch separately, and then aggregate over all batches in the very end. Fix batch B , and let $t_0 = t_0(B)$ be the last round in this batch. Let $\tau = 1 + t_0/Y$, and consider the context vector x_τ^0 chosen by ALG_0 in round τ . This context vector is a randomized function f of the current context tuple con_τ and the history $h_{\tau-1}^0$:

$$x_\tau^0 = f(\text{con}_\tau; h_{\tau-1}^0).$$

By Lemma 6.10, letting $\sigma = (\text{con}_1, \dots, \text{con}_{\lfloor t/Y \rfloor})$, it holds that

$$(6.14) \quad \mathbb{E}[x_\tau^0 \cdot \theta \mid \sigma, \theta, \mathcal{E}_1, \mathcal{E}_2] = \mathbb{E}[f(\text{con}_\tau; \text{sim}(h_{t_0}, \sigma)) \cdot \theta \mid \sigma, \theta, \mathcal{E}_1, \mathcal{E}_2]$$

Let t be some round in the next batch after B , and let $x'_t = x_{a'_t, t}$, be the context vector predicted by ALG in round t . Recall that x'_t is a Bayesian-greedy choice from the context tuple con_t , based on history h_{t_0} . Observe that the Bayesian-greedy action choice from a given context tuple based on history h_{t_0} cannot be worse, in terms of the Bayesian-expected reward, than any other choice from the same context tuple and based on the same history. Using (6.14), we obtain:

$$(6.15) \quad \mathbb{E}[x'_t \cdot \theta \mid \text{con}_t = \text{con}, \mathcal{E}_1, \mathcal{E}_2] \geq \mathbb{E}[x_\tau^0 \cdot \theta \mid \text{con}_\tau = \text{con}, \mathcal{E}_1, \mathcal{E}_2],$$

for any given context tuple $\text{con} \in \text{CON}$ that has a non-zero arrival probability given $\mathcal{E}_1 \cap \mathcal{E}_2$.

Observe that con_t and con_τ have the same distribution, even conditioned on event $\mathcal{E}_1 \cap \mathcal{E}_2$. (This is because the definitions of \mathcal{E}_1 and \mathcal{E}_2 treat all rounds in the same batch in exactly the same way.) Therefore, we can integrate (6.15) over the context tuples con :

$$(6.16) \quad \mathbb{E}[x'_t \cdot \theta \mid \mathcal{E}_1, \mathcal{E}_2] \geq \mathbb{E}[x_\tau^0 \cdot \theta \mid \mathcal{E}_1, \mathcal{E}_2],$$

Now, let us sum up (6.16) over all rounds t in the next batch after B , denote it $\text{next}(B)$.

$$(6.17) \quad \sum_{t \in \text{next}(B)} \mathbb{E}[x'_t \cdot \theta \mid \mathcal{E}_1, \mathcal{E}_2] \geq Y \cdot \mathbb{E}[x_\tau^0 \cdot \theta \mid \mathcal{E}_1, \mathcal{E}_2].$$

Note that the right-hand side of (6.16) stays the same for all t , hence the factor of Y on the right-hand side of (6.17). This completes our analysis of a single batch B .

We obtain (6.13) by integrating over all batches B . Here it is essential that the expectation $\mathbb{E}[\theta^\top x_t^*]$ does not depend on round t , and therefore the “regret benchmark” $\theta^\top x_t^*$ cancels out from (6.13). In particular, it is essential that the context tuples con_t are identically distributed across rounds. \square

Proof of Theorem 6.7 given Equation (6.13). We must take care of the low-probability failure events $\bar{\mathcal{E}}_1$ and $\bar{\mathcal{E}}_2$. Specifically, we need to upper-bound the expression

$$\mathbb{E}_{\theta \sim P} [\text{PReg}^\mathcal{T}(T) \mid \bar{\mathcal{E}}_1 \cup \bar{\mathcal{E}}_2] \cdot \Pr[\bar{\mathcal{E}}_1 \cup \bar{\mathcal{E}}_2].$$

We know that $\Pr[\bar{\mathcal{E}}_1 \cup \bar{\mathcal{E}}_2] \leq \delta + \delta_R$. Lemma 8.7 with $\ell = \hat{R}$ gives us that the instantaneous regret of every round is at most

$$\begin{aligned} 2 \mathbb{E}_{\theta \sim (P \mid h_{t-1})} \left[\|\theta\|_2 \left(1 + \rho(2 + \sqrt{2 \log K}) + \hat{R} \right) \right] \\ \leq 2 \left[\left(\|\bar{\theta}\|_2 + \sqrt{d \lambda_{\max}(\Sigma)} \right) \left(1 + \rho(2 + \sqrt{2 \log K}) + \hat{R} \right) \right] \end{aligned}$$

by Lemma A.6. Letting $\delta = \delta_R = \frac{1}{T^2}$, we verify that our definition of Y means that Lemma 6.4 indeed holds with probability at least $1 - T^{-2}$. Using (6.13), the Bayesian prediction regret of ALG is

$$\begin{aligned} \mathbb{E}_{\theta \sim \mathcal{P}} [\text{PReg}^T(T)] &\leq Y \mathbb{E}_{\theta \sim \mathcal{P}} [R_0^T(\frac{T}{Y})] \\ &\quad + 2T(\delta + \delta_R) \left[\left(\|\bar{\theta}\|_2 + \sqrt{d\lambda_{\max}(\Sigma)} \right) \left(1 + \rho(2 + \sqrt{2\log K}) + \hat{R} \right) \right] \\ &\leq Y \mathbb{E}_{\theta \sim \mathcal{P}} [R_0^T(\frac{T}{Y})] + \tilde{O}\left(\frac{1}{T}\right). \end{aligned}$$

This completes the proof of Theorem 6.7. \square

6.4. Regret Bounds for BatchFreqGreedy. To analyze BatchFreqGreedy, we show that its Bayesian regret is not too different from its Bayesian prediction regret, and use Corollary 6.9 to bound the latter.

THEOREM 6.11. *Consider perturbed context generation. Suppose prior \mathcal{P} is a multivariate Gaussian distribution with invertible covariance matrix Σ , and the eigenvalues of Σ are at least ρ^{-4}/T . Then BatchFreqGreedy satisfies*

$$|\mathbb{E}[R^T(T) - \text{PReg}^T(T)]| \leq \tilde{O}\left(\frac{\sqrt{d}}{\rho^2}\right) \left(\sqrt{\lambda_{\max}(\Sigma)} + \frac{1}{\sqrt{\lambda_{\min}(\Sigma)}} \right),$$

where Σ is the covariance matrix of the prior and ρ is the perturbation size.

Using Corollary 6.9, we obtain regret bounds for BatchFreqGreedy in Theorem 4.1 and Theorem 4.4.

The remainder of this section is dedicated to proving Theorem 6.11. On a high level, the idea is as follows. As in the proof of Theorem 6.7, we condition on the high-probability event (6.12) that perturbations are bounded. We prove that

$$(6.18) \quad |\mathbb{E}[R^T(T) - \text{PReg}^T(T) \mid \mathcal{E}_1]| \leq \tilde{O}\left(\frac{\sqrt{d}}{\rho^2}\right) \left(\sqrt{\lambda_{\max}(\Sigma)} + \frac{1}{\sqrt{\lambda_{\min}(\Sigma)}} \right).$$

To prove this statement, we fix round t and compare the action a_t taken by BatchFreqGreedy and the predicted action a'_t . We observe that the difference in rewards between these two actions can be upper-bounded in terms of $\theta_t^{\text{bay}} - \theta_t^{\text{fre}}$, the difference in the θ estimates with and without knowledge of the prior. (Recall (3.3) and (3.4) for definitions.) Specifically, we show that

$$(6.19) \quad \mathbb{E}[\theta^\top(x_{a_t, t} - x_{a'_t, t}) \mid \mathcal{E}_1] \leq 2R \mathbb{E}_{\theta \sim \mathcal{P}} [\|\theta_t^{\text{bay}} - \theta_t^{\text{fre}}\|_2].$$

The crux of the proof is to show that the difference $\|\theta_t^{\text{bay}} - \theta_t^{\text{fre}}\|_2$ is small, namely

$$(6.20) \quad \mathbb{E}[\|\theta_t^{\text{bay}} - \theta_t^{\text{fre}}\|_2 \mid \mathcal{E}_1] = \tilde{O}(1/t),$$

ignoring other parameters. Interestingly, the two estimates are much closer to each other than they are to θ , as either estimate is typically $\Omega(1/\sqrt{t})$ away from θ .

Thus, summing over all rounds, we get

$$\mathbb{E}[R^T(T) - \text{PReg}^T(T) \mid \mathcal{E}_1] \leq O(\log T) = \tilde{O}(1).$$

Once we prove that (6.18) holds under event (6.12), the proof of Theorem 6.11 is easily completed as follows. Recall that event (6.12) happens with probability at least $1 - \delta_R$. When this event fails to hold, the total regret is at most

$$2 \left[\left(\|\bar{\theta}\|_2 + \sqrt{d\lambda_{\max}(\Sigma)} \right) \left(1 + \rho(2 + \sqrt{2\log K}) + \hat{R} \right) \right]$$

by Lemma 8.7 (with $\ell = \hat{R}$) and Lemma A.6. Since $\delta_R = T^{-2}$, the contribution of regret when the high-probability bound fails is $\tilde{O}(1/T) \leq \tilde{O}(1)$.

6.4.1. Proof of Eq. (6.18). Let R^t and PReg^t be, resp., instantaneous regret and instantaneous prediction regret at time t . Then

$$(6.21) \quad \mathbb{E}_{\theta \sim \mathcal{P}} [R^{\mathcal{T}}(T) - \text{PReg}^{\mathcal{T}}(T)] = \sum_t \mathbb{E}_{\theta \sim \mathcal{P}} [R^t - \text{PReg}^t].$$

Thus, it suffices to bound the differences in instantaneous regret.

Recall that at time t , the chosen action for BatchFreqGreedy and the predicted action are, resp.,

$$\begin{aligned} a_t &= \arg \max_{a \in A} x_{a,t}^\top \theta_t^{\text{fre}} \\ a'_t &= \arg \max_{a \in A} x_{a,t}^\top \theta_t^{\text{bay}}. \end{aligned}$$

Letting $t_0 - 1 = \lfloor t/Y \rfloor$ be the last round in the previous batch, we can formulate θ_t^{fre} and θ_t^{bay} as

$$\begin{aligned} \theta_t^{\text{fre}} &= (Z_{t_0-1})^{-1} X_{t_0-1}^\top \mathbf{r}_{1:t_0-1} \\ \theta_t^{\text{bay}} &= (Z_{t_0-1} + \Sigma^{-1})^{-1} (X_{t_0-1}^\top \mathbf{r}_{1:t_0-1} + \Sigma^{-1} \bar{\theta}). \end{aligned}$$

Therefore, we have

$$\mathbb{E}_{\theta \sim \mathcal{P} \mid h_{t-1}} [R^t - \text{PReg}^t] = \mathbb{E}_{\theta \sim \mathcal{P} \mid h_{t-1}} [(x_{a'_t,t} - x_{a_t,t})^\top \theta_t^{\text{bay}}] = (x_{a'_t,t} - x_{a_t,t})^\top \theta_t^{\text{bay}},$$

since the mean of the posterior distribution is exactly θ_t^{bay} , and θ_t^{bay} is deterministic given h_{t-1} . Taking expectation over h_{t-1} , we have

$$\mathbb{E}_{\theta \sim \mathcal{P}} [R^t - \text{PReg}^t] = \mathbb{E}_{\theta \sim \mathcal{P}} [(x_{a'_t,t} - x_{a_t,t})^\top \theta_t^{\text{bay}}].$$

For any fixed θ_t^{bay} and θ_t^{fre} , since BatchFreqGreedy chose a_t over a'_t , it must be the case that

$$(6.22) \quad x_{a_t,t}^\top \theta_t^{\text{fre}} \geq x_{a'_t,t}^\top \theta_t^{\text{fre}}.$$

Therefore,

$$\begin{aligned} (x_{a'_t,t} - x_{a_t,t})^\top \theta_t^{\text{bay}} &= (x_{a'_t,t} - x_{a_t,t})^\top \theta_t^{\text{fre}} + (x_{a'_t,t} - x_{a_t,t})^\top (\theta_t^{\text{bay}} - \theta_t^{\text{fre}}) \\ (\text{By (6.22)}) \quad &\leq (x_{a'_t,t} - x_{a_t,t})^\top (\theta_t^{\text{bay}} - \theta_t^{\text{fre}}) \\ &\leq (\|x_{a'_t,t}\|_2 + \|x_{a_t,t}\|_2) \|\theta_t^{\text{bay}} - \theta_t^{\text{fre}}\|_2 \\ &\leq 2R \|\theta_t^{\text{bay}} - \theta_t^{\text{fre}}\|_2 \end{aligned}$$

Eq. (6.19) follows.

The crux is to prove (6.20): to bound the expected distance between the Frequentist and Bayesian estimates for θ . By expanding their definitions, and denoting $M = (Z_{t_0-1} + \Sigma^{-1})^{-1}$ for succinctness, we have

$$\begin{aligned}\theta_t^{\text{bay}} - \theta_t^{\text{fre}} &= M(X_{t_0-1}^\top \mathbf{r}_{1:t_0-1} + \Sigma^{-1}\bar{\theta}) - Z_{t_0-1}^{-1}X_{t_0-1}^\top \mathbf{r}_{1:t_0-1} \\ &= M[X_{t_0-1}^\top \mathbf{r}_{1:t_0-1} + \Sigma^{-1}\bar{\theta} - (Z_{t_0-1} + \Sigma^{-1})Z_{t_0-1}^{-1}X_{t_0-1}^\top \mathbf{r}_{1:t_0-1}] \\ &= M[X_{t_0-1}^\top \mathbf{r}_{1:t_0-1} + \Sigma^{-1}\bar{\theta} - X_{t_0-1}^\top \mathbf{r}_{1:t_0-1} - \Sigma^{-1}Z_{t_0-1}^{-1}X_{t_0-1}^\top \mathbf{r}_{1:t_0-1}] \\ &= M[\Sigma^{-1}\bar{\theta} - \Sigma^{-1}Z_{t_0-1}^{-1}X_{t_0-1}^\top \mathbf{r}_{1:t_0-1}] \\ &= M\Sigma^{-1}(\bar{\theta} - \theta_t^{\text{fre}}).\end{aligned}$$

Next, note that

$$\begin{aligned}\|M\Sigma^{-1}(\bar{\theta} - \theta_t^{\text{fre}})\|_2 &\leq \|M\|_2 \|\Sigma^{-1}(\bar{\theta} - \theta_t^{\text{fre}})\|_2 \\ &\leq \|(Z_{t_0-1} + \Sigma)^{-1}\|_2 (\|\Sigma^{-1}(\bar{\theta} - \theta)\|_2 + \|\Sigma^{-1}\|_2 \|\theta - \theta_t^{\text{fre}}\|_2).\end{aligned}$$

By Lemma A.11, $\lambda_{\min}(Z_{t_0-1} + \Sigma) \geq \lambda_{\min}(Z_{t_0-1})$. Therefore,

$$\|(Z_{t_0-1} + \Sigma)^{-1}\|_2 = \frac{1}{\lambda_{\min}(Z_{t_0-1} + \Sigma)} \leq \frac{1}{\lambda_{\min}(Z_{t_0-1})},$$

giving us

$$\begin{aligned}\|\theta_t^{\text{bay}} - \theta_t^{\text{fre}}\|_2 &\leq \frac{\|\Sigma^{-1}(\bar{\theta} - \theta)\|_2 + \|\Sigma^{-1}\|_2 \|\theta - \theta_t^{\text{fre}}\|_2}{\lambda_{\min}(Z_{t_0-1})} \\ &\leq \frac{\|\Sigma^{-1/2}\|_2 \|\Sigma^{-1/2}(\bar{\theta} - \theta)\|_2 + \|\Sigma^{-1}\|_2 \|\theta - \theta_t^{\text{fre}}\|_2}{\lambda_{\min}(Z_{t_0-1})} \\ &= \frac{\left(\|\Sigma^{-1/2}(\bar{\theta} - \theta)\|_2 + \frac{1}{\sqrt{\lambda_{\min}(\Sigma)}} \|\theta - \theta_t^{\text{fre}}\|_2\right)}{\sqrt{\lambda_{\min}(\Sigma)} \lambda_{\min}(Z_{t_0-1})}.\end{aligned}$$

Next, recall that for

$$t_0 - 1 \geq t_{\min}(\delta) := 160 \frac{R^2}{\rho^2} \log \frac{2d}{\delta} \cdot \log T$$

the following bounds hold, each with probability at least $1 - \delta$:

$$\text{(Lemma 6.1)} \quad \frac{1}{\lambda_{\min}(Z_{t_0-1})} \leq \frac{32 \log T}{\rho^2(t_0 - 1)}$$

$$\text{(Lemma 6.3)} \quad \|\theta - \theta_t^{\text{fre}}\|_2 \leq \frac{\sqrt{2dR(t_0 - 1) \log(d/\delta)}}{\lambda_{\min}(Z_{t_0-1})}$$

Therefore, fixing $t_0 \geq 1 + t_{\min}(\delta/2)$, with probability at least $1 - \delta$ we have

$$(6.23) \quad \|\theta_t^{\text{bay}} - \theta_t^{\text{fre}}\|_2 \leq \frac{32 \log T}{\rho^2(t_0 - 1) \sqrt{\lambda_{\min}(\Sigma)}} \left(\|\Sigma^{-1/2}(\bar{\theta} - \theta)\|_2 + \Phi \sqrt{d} \right),$$

where for succinctness we denote

$$\Phi := \frac{64 \sqrt{R \log(2d/\delta)} \cdot \log T}{\rho^2 \sqrt{(t_0 - 1) \lambda_{\min}(\Sigma)}}.$$

Note that the high-probability events we need are deterministic given h_{t_0-1} , and therefore are independent of the perturbations at time t . This means that Lemma 8.7 applies, with $\ell = 0$: conditioned on any h_{t_0-1} , the expected regret for round t is upper-bounded by $2\|\theta\|_2(1 + \rho(1 + \sqrt{2\log K}))$. In particular, this holds for any h_{t_0-1} not satisfying the high probability events from Lemmas 6.1 and 6.3. Therefore, for all $t \geq t_{\min}(\delta)$,

$$\begin{aligned}
& \mathbb{E}_{\theta \sim \mathcal{P}} \left[\|\theta_t^{\text{bay}} - \theta_t^{\text{fre}}\|_2 \right] \\
& \leq \mathbb{E}_{\theta \sim \mathcal{P}} \left[(1 - \delta) \frac{32 \log T}{\rho^2(t_0 - 1) \sqrt{\lambda_{\min}(\Sigma)}} \left(\|\Sigma^{-1/2}(\bar{\theta} - \theta)\|_2 + \Phi \sqrt{d} \right) \right. \\
& \quad \left. + \delta \cdot 2\|\theta\|_2(1 + \rho(2 + \sqrt{2\log K})) \right] \\
& \leq \frac{32 \log T}{\rho^2(t_0 - 1) \sqrt{\lambda_{\min}(\Sigma)}} \left(\mathbb{E}_{\theta \sim \mathcal{P}} \left[\|\Sigma^{-1/2}(\bar{\theta} - \theta)\|_2 \right] + \Phi \sqrt{d} \right) \\
& \quad + \delta \cdot 2(\|\bar{\theta}\|_2 + \mathbb{E}_{\theta \sim \mathcal{P}} [\|\bar{\theta} - \theta\|_2])(1 + \rho(2 + \sqrt{2\log K})).
\end{aligned}$$

Because $\theta \sim \mathcal{N}(\bar{\theta}, \Sigma)$, we have $\Sigma^{-1/2}(\bar{\theta} - \theta) \sim \mathcal{N}(0, I)$. By Lemma A.6,

$$\mathbb{E}_{\theta \sim \mathcal{P}} \left[\|\Sigma^{-1/2}(\bar{\theta} - \theta)\|_2 \right] \leq \sqrt{d} \quad \text{and} \quad \mathbb{E}_{\theta \sim \mathcal{P}} [\|\bar{\theta} - \theta\|_2] \leq \sqrt{d\lambda_{\max}(\Sigma)}.$$

This means

$$\begin{aligned}
\mathbb{E}_{\theta \sim \mathcal{P}} \left[\|\theta_t^{\text{bay}} - \theta_t^{\text{fre}}\|_2 \right] & \leq \frac{32\sqrt{d} \log T}{\rho^2(t_0 - 1) \sqrt{\lambda_{\min}(\Sigma)}} (1 + \Phi) \\
& \quad + \delta \cdot 2(\|\bar{\theta}\|_2 + \sqrt{d\lambda_{\max}(\Sigma)})(1 + \rho(2 + \sqrt{2\log K})).
\end{aligned}$$

Since $t_0 = \Omega(t)$, for sufficiently small δ , this proves (6.20).

We need to do a careful computation to complete the proof of Eq. (6.18). We know from (6.19) that

$$\mathbb{E}_{\theta \sim \mathcal{P}} [R^{\mathcal{T}}(T) - \text{PReg}^{\mathcal{T}}(T)] \leq \sum_{t=1}^T 2R \mathbb{E}_{\theta \sim \mathcal{P}} [\|\theta_t^{\text{bay}} - \theta_t^{\text{fre}}\|_2].$$

Choosing $\delta = T^{-2}$, we find that

$$\sum_{t=t_{\min}(T^{-2})}^T \delta \cdot 2(\|\bar{\theta}\|_2 + \sqrt{d\lambda_{\max}(\Sigma)})(1 + \rho(2 + \sqrt{2\log K})) = \tilde{O}(1),$$

so this term vanishes. Furthermore,

$$\begin{aligned}
& \sum_{t=t_{\min}(T^{-2})}^T 2R \frac{32\sqrt{d} \log T}{\rho^2(t_0 - 1) \sqrt{\lambda_{\min}(\Sigma)}} (1 + \Phi) \\
& = \tilde{O} \left(\frac{R\sqrt{d}}{\rho^2 \sqrt{\lambda_{\min}(\Sigma)}} \right)
\end{aligned}$$

as long as $\rho^2 \sqrt{\lambda_{\min}(\Sigma)} \geq T^{-1/2}$, since $t_0 \geq t - Y$, and $\sum_{t=1}^T 1/t = O(\log T)$. Using the fact that $R = \tilde{O}(1)$ (since by assumption $\rho \leq d^{-1/2}$), this is simply

$$\tilde{O}\left(\frac{\sqrt{d}}{\rho^2 \sqrt{\lambda_{\min}(\Sigma)}}\right).$$

Finally, we note that on the first $t_{\min}(T^{-2}) = \tilde{O}(1/\rho^2)$ rounds, the regret bound from Lemma 8.7 with $\ell = 0$ applies, so the total regret difference is at most

$$\begin{aligned} & \mathbb{E}_{\theta \sim \mathcal{P}} [R^T(T) - \text{PReg}^T(T)] \\ & \leq \sum_{t=1}^{t_{\min}(T^{-2})} \mathbb{E}_{\theta \sim \mathcal{P}} [R^t - \text{PReg}^t] + \sum_{t=t_{\min}(T^{-2})}^T 2R \mathbb{E}_{\theta \sim \mathcal{P}} [\|\theta_t^{\text{bay}} - \theta_t^{\text{fre}}\|_2], \\ & \leq t_{\min}(T^{-2}) \cdot 2(\|\bar{\theta}\|_2 + \sqrt{d\lambda_{\max}(\Sigma)})(1 + \rho(2 + \sqrt{2\log K})) + \tilde{O}\left(\frac{\sqrt{d}}{\rho^2 \sqrt{\lambda_{\min}(\Sigma)}}\right) \\ & = \tilde{O}\left(\frac{\sqrt{d\lambda_{\max}(\Sigma)}}{\rho^2}\right) + \tilde{O}\left(\frac{\sqrt{d}}{\rho^2 \sqrt{\lambda_{\min}(\Sigma)}}\right), \end{aligned}$$

which implies Eq. (6.18).

7. Lower Bound: Proof of Theorem 4.3. Here, we show that for fully-action-correlated perturbations, i.e., when every arm at a given timestep is perturbed by the same perturbation, no algorithm can achieve regret less than \sqrt{T} .

Consider the following problem instance. There are $d = 2$ dimensions and $K = 2$ arms, with $\mu_{1,t} = [1 \ 0]^\top$ and $\mu_{2,t} = [0 \ 1]^\top$ at each round. (For intuition, one can think of them as, resp., the *horizontal* arm and the *vertical* arm.) There are two possible hidden vectors: $\theta_1 = [1 + \delta \ 1]^\top$ and $\theta_2 = [1 \ 1 + \delta]^\top$, occurring with probability $\frac{1}{2}$ each. Here δ is a parameter which we specify later in the analysis. Thus, arm 1 is preferable in expectation for $\theta = \theta_1$ and arm 2 is preferable for $\theta = \theta_2$. We will show that even under perturbations, we need $\Omega(1/\delta^2)$ samples to distinguish between them, meaning we get \sqrt{T} regret for $\delta \sim 1/\sqrt{T}$.

By definition of fully-action-correlated perturbation, at any round t , both $\mu_{1,t}$ and $\mu_{2,t}$ have the same perturbation $\varepsilon_t \sim N(0, \rho^2 I)$ added to them, where ρ is a perturbation size and I is the 2-dimensional identity matrix. Given $\varepsilon_t = \varepsilon$, the arms' expected rewards under θ_1 and θ_2 are, resp.:

$$\begin{aligned} \mathbb{E}[\theta_1^\top(\mu_{1,t} + \varepsilon)] &= 1 + \delta + \theta_1^\top \varepsilon & \mathbb{E}[\theta_2^\top(\mu_{1,t} + \varepsilon)] &= 1 + \theta_2^\top \varepsilon & (\text{arm } 1), \\ \mathbb{E}[\theta_1^\top(\mu_{2,t} + \varepsilon)] &= 1 + \theta_1^\top \varepsilon & \mathbb{E}[\theta_2^\top(\mu_{2,t} + \varepsilon)] &= 1 + \delta + \theta_2^\top \varepsilon & (\text{arm } 2). \end{aligned}$$

We analyze this problem instance using a standard KL-divergence technique, e.g., see [23, Chapter 2]. Compared to the standard analysis, we need to handle contexts. To this end, we fix the realized sequence of perturbation vectors $\varepsilon_1, \dots, \varepsilon_T$, and condition on the high-probability event that perturbations are not too large:

$$\mathcal{E}_{\text{nice}} = \{ \|\varepsilon_t\|_2^2 \leq \Psi \text{ for all rounds } t \}, \text{ where } \Psi := 2 + 8\sqrt{\log T} + 8\log T.$$

Since each $\|\varepsilon_t\|_2^2$ follows a χ^2 -distribution with 2 degrees of freedom, a standard tail bound (e.g., [23, Lemma 1]) implies that

$$\Pr[\mathcal{E}_{\text{nice}}] \geq 1 - \sum_{t=1}^T \Pr[\|\varepsilon_t\|_2^2 > \Psi] \geq 1 - 1/T.$$

Clearly, it suffices to prove a regret bound for such sequence $\varepsilon_1, \dots, \varepsilon_T$.

The rest of the analysis consists of two parts: a generic K-divergence argument leading to (7.1), and an application of (7.1) to an execution of a given algorithm. We set $\delta = \frac{1}{8\sqrt{T\Psi}}$.

A generic KL-divergence argument. Given perturbation $\varepsilon_t = \varepsilon$, let $\mathcal{D}_{\theta, \varepsilon}^{(i)}$ be the probability distribution of rewards under hidden vector θ when choosing arm i for $\theta \in \{\theta_1, \theta_2\}$ and $i \in \{1, 2\}$. The KL-divergence between Gaussians with variance 1 and means ξ_1, ξ_2 is $(\xi_1 - \xi_2)^2/2$. Since rewards are assumed to be Gaussian with variance 1, the KL-divergence between the reward distributions of the two arms is

$$\begin{aligned} \text{KL} \left(\mathcal{D}_{\theta_1, \varepsilon}^{(i)}, \mathcal{D}_{\theta_2, \varepsilon}^{(i)} \right) &\leq \left(\delta + |(\theta_1 - \theta_2)^\top \varepsilon| \right)^2 / 2 \\ &\leq 2 \max \left(\delta^2, ((\theta_1 - \theta_2)^\top \varepsilon)^2 \right) \\ &\leq 2 \max \left(\delta^2, \|\theta_1 - \theta_2\|_2^2 \|\varepsilon\|_2^2 \right) \\ &\leq 2 \max \left(\delta^2, 2\delta^2 \|\varepsilon\|_2^2 \right) \leq 4\delta^2 \Psi \quad (\text{under event } \mathcal{E}_{\text{nice}}). \end{aligned}$$

Let $\mathcal{D}_{\theta, \varepsilon} = \mathcal{D}_{\theta, \varepsilon}^{(1)} \times \mathcal{D}_{\theta, \varepsilon}^{(2)}$ be the joint distribution of rewards from both arms, fixing the perturbation. By the chain rule,

$$\text{KL}(\mathcal{D}_{\theta_1, \varepsilon}, \mathcal{D}_{\theta_2, \varepsilon}) \leq 8\delta^2 \Psi.$$

Fix a realized sequence of perturbations $\varepsilon = \{\varepsilon_t\}_{t \in [T]}$ which satisfies $\mathcal{E}_{\text{nice}}$. Let $\Omega = (\mathbb{R} \times \mathbb{R})^T$ be the event space of rewards from the two arms, with events of the form $(r_{1,t}, r_{2,t})_{t \in [T]}$. Note that θ_1 and θ_2 each impose distributions over Ω , call them $p = \prod_{t \in [T]} p_t$ and $q = \prod_{t \in [T]} q_t$, respectively. By a standard application of Pinsker's inequality (e.g., see Lemma 2.5 in [30]), for any event $A \subset \Omega$ it holds that

$$\begin{aligned} 2(p(A) - q(A))^2 &\leq \text{KL}(p, q) = \sum_{t \in [T]} \text{KL}(p_t, q_t) \\ &\leq T \cdot \text{KL}(\mathcal{D}_{\theta_1, \varepsilon}, \mathcal{D}_{\theta_2, \varepsilon}) \\ &\leq 8\delta^2 T \Psi \\ (7.1) \quad |p(A) - q(A)| &\leq 2\delta\sqrt{T\Psi}. \end{aligned}$$

Using (7.1) to bound regret. Consider any deterministic algorithm **ALG** for linear contextual bandits. Let A be the event that **ALG** chooses arm 1 in at least $T/2$ rounds. Since the algorithm is deterministic, A can be interpreted as an event in Ω . Note that if A occurs when $\theta = \theta_2$ or if $\neg A$ occurs when $\theta = \theta_1$, then **ALG** incurs expected regret $\Omega(\delta T) = \tilde{\Omega}(\sqrt{T})$, as desired.

We consider two cases, depending on whether $p(A) \geq 1/2$.

Case 1: $p(A) \geq 1/2$. Then, $q(A) \geq 1/2 - 2\delta\sqrt{T\Psi} = 1/4$, and expected regret is

$$\begin{aligned} \mathbb{E}[R(T)] &\geq \Pr[\theta = \theta_2] \cdot \mathbb{E}[R(T) \mid \theta = \theta_2] \\ &\geq 1/2 \cdot \Pr[A \mid \theta = \theta_2] \cdot \mathbb{E}[R(T) \mid \theta = \theta_2, A] \\ &\geq 1/2 \cdot q(A) \cdot \frac{\delta T}{2} \\ &\geq \delta T / 16. \end{aligned}$$

Case 2: $p(A) < 1/2$. Then, expected regret is

$$\begin{aligned}\mathbb{E}[R(T)] &\geq \Pr[\theta = \theta_1] \cdot \mathbb{E}[R(T) \mid \theta = \theta_1] \\ &\geq 1/2 \cdot \Pr[\neg A \mid \theta = \theta_1] \mathbb{E}[R(T) \mid \theta = \theta_1, \neg A] \\ &\geq 1/2 \cdot (1 - p(A)) \cdot \frac{\delta T}{2} \\ &\geq \delta T/8.\end{aligned}$$

Thus, $\mathbb{E}[R(T)] \geq \tilde{\Omega}(\sqrt{T})$. This extends to randomized algorithms by taking expectations over the algorithm's random seed.

8. LinUCB with Perturbed Contexts. We prove Theorem 4.4(a), a Bayesian regret bound for the LinUCB algorithm under perturbed context generation. For this section, we focus on action-independent perturbation with perturbation size ρ , and posit a multivariate Gaussian prior $\mathcal{P} = \mathcal{N}(\bar{\theta}, \Sigma)$, with mean vector $\bar{\theta} \in \mathbb{R}^d$ and invertible covariate matrix $\Sigma \in \mathbb{R}^{d \times d}$.

8.1. Preliminaries: LinUCB algorithm. LinUCB is a well-known algorithm for linear contextual bandits, which implements the paradigm of ‘optimism under uncertainty’. The idea is to evaluate each action “optimistically”—assuming the best-case scenario for this action—and then choose an action with the best optimistic evaluation. For the basic setting of multi-armed bandits, one chooses an action with the highest upper confidence bound (henceforth, UCB) on its mean reward. The UCB is computed as the sample average of the reward for this action plus a term which captures the amount of uncertainty. (This is a seminal algorithm called UCB1 [7].)

Going back linear contextual bandits, the high-level idea is to compute a confidence region $\Theta_t \subset \mathbb{R}^d$ in each round t such that $\theta \in \Theta_t$ with high probability, and choose an action a which maximizes the optimistic reward estimate $\sup_{\theta \in \Theta_t} x_{a,t}^\top \theta$. Concretely, one uses regression to form an empirical estimate $\hat{\theta}_t$ for θ . Concentration techniques lead to high-probability bounds of the form $|x^\top(\theta - \hat{\theta}_t)| \leq f(t)\sqrt{x^\top Z_t^{-1}x}$, where the *interval width function* $f(t)$ may depend on hyperparameters and features of the instance. LinUCB simply chooses an action

$$(8.1) \quad a_t^{\text{LinUCB}} := \arg \max_a x_{a,t}^\top \hat{\theta}_t + f(t)\sqrt{x_{a,t}^\top Z_t^{-1}x_{a,t}}.$$

We focus on a version from [1], with

$$(8.2) \quad f(t) = S + \sqrt{dc_0 \log(T + tTL^2)},$$

here L and S are known upper bounds on $\|x_{a,t}\|_2$ and $\|\theta\|_2$, respectively, and c_0 is a parameter. For any $c_0 \geq 1$, one obtains regret $\tilde{O}(dS\sqrt{c_0 T})$, with only a polylog dependence on TL/d [1].

8.2. Our result. Recall that ρ denotes perturbation size, and $\bar{\theta} = \mathbb{E}[\theta]$ is the prior mean of the latent vector θ . The parameters from (8.2) are set as follows:

$$(8.3) \quad \begin{aligned}L &\geq 1 + \rho\sqrt{2d \log(2T^3 Kd)}, \\ S &\geq \|\bar{\theta}\|_2 + \sqrt{3d \log T} \quad (\text{and } S < T) \\ c_0 &= 1.\end{aligned}$$

REMARK 8.1. Ideally we would like to set L, S according to (8.3) with equalities. We consider a more permissive version with inequalities so as to not require the exact knowledge of ρ and $\|\bar{\theta}\|_2$. While the original result in [1] requires $\|x_{a,t}\|_2 \leq L$ and $\|\theta\|_2 \leq S$, in our setting this only happens with high probability.

We prove the following theorem (which implies Theorem 4.4(a)):

THEOREM 8.2. Assume perturbed context generation, with action-independent perturbation. Further, suppose that the maximal eigenvalue of the covariance matrix Σ of the prior \mathcal{P} is at most 1, and the mean vector satisfies $\|\bar{\theta}\|_2 \geq 1 + \sqrt{3 \log T}$. The version of LinUCB with interval width function (8.2) and parameters given by (8.3) has Bayesian regret at most

$$(8.4) \quad T^{1/3} \left(d^2 S (K^2 / \rho)^{1/3} \right) \cdot \text{polylog}(TKLd).$$

REMARK 8.3. The theorem also holds if the assumption on $\|\bar{\theta}\|_2$ is replaced with $d \geq \frac{\log T}{\log \log T}$. The only change in the analysis is that in the concluding steps (Section 8.5), we use Lemma 8.6(b) instead of Lemma 8.6(a).

8.3. Key steps of the analysis. On a high level, our analysis proceeds as follows. We massage algorithm’s regret so as to elucidate the dependence on the number of rounds with small “gap” between the best and second-best action, call it N . This step does not rely on perturbed context generation, and makes use of the analysis from [1]. The crux is that we derive a much stronger upper-bound on $\mathbb{E}[N]$ under perturbed context generation. The analysis relies on some non-trivial technicalities on bounding the deviations from the “high-probability” behavior, which are gathered in Section 8.4.

We reuse the analysis in [1] via the following lemma.⁹ To state this lemma, define the instantaneous regret at time t as $R_t = \theta^\top x_t^* - \theta^\top x_{a_t,t}$, and let

$$\beta_T = \left(\sqrt{d \log(T(1 + TL^2))} + S \right)^2.$$

LEMMA 8.4 ([1]). Consider a problem instance with reward noise $\mathcal{N}(0, 1)$ and a specific realization of latent vector θ and contexts $x_{a,t}$. Consider LinUCB with parameters L, S, c_0 that satisfy $\|x_{a,t}\|_2 \leq L$, $\|\theta\|_2 \leq S$, and $c_0 = 1$. Then

(a) with probability at least $1 - \frac{1}{T}$ (over the randomness in the rewards),

$$\sum_{t=1}^T R_t^2 \leq 16\beta_T \log(\det(Z_t + I)),$$

where Z_t is the “empirical covariance matrix” at time t :

$$Z_t = \sum_{\tau=1}^t x_\tau x_\tau^\top \in \mathbb{R}^{d \times d}.$$

(b) $\det(Z_t + I) \leq (1 + tL^2/d)^d$.

The following lemma captures the essence of the proof of Theorem 8.2. From here on, we assume perturbed context generation without further notice. In particular, reward noise is $\mathcal{N}(0, 1)$.

⁹Lemma 8.4(a) is implicit in the proof of Theorem 3 from [1], and Lemma 8.4(b) is asserted by [1, Lemma 10].

LEMMA 8.5. Suppose parameter L is set as in (8.3). Consider a problem instance with a specific realization of θ such that $\|\theta\|_2 \leq S$. Then,

$$\mathbb{E}[\text{Regret}(T)] \leq \|\theta\|_2^{-1/3} \left(\frac{1}{2\sqrt{\pi}} + 16\beta_T d \log(1 + TL^2/d) \right) \left(\frac{TK^2}{\rho} \right)^{1/3} + \tilde{O}(1).$$

Proof. We will prove that for any $\gamma > 0$,

$$(8.5) \quad \mathbb{E}[\text{Regret}(T)] \leq T \cdot \frac{\gamma^2 K^2}{2\rho\|\theta\|_2\sqrt{\pi}} + \frac{1}{\gamma} 16\beta_T d \log(1 + TL^2/d) + \tilde{O}(1).$$

The Lemma easily follows by setting $\gamma = (TK^2/(\rho\|\theta\|_2))^{-1/3}$.

Fix some $\gamma > 0$. We distinguish between rounds t with $R_t < \gamma$ and those with $R_t \geq \gamma$:

$$(8.6) \quad \text{Regret}(T) = \sum_{t=1}^T R_t \leq \sum_{t \in \mathcal{T}_\gamma} R_t + \sum_{t=1}^T \frac{R_t^2}{\gamma} \leq \gamma |\mathcal{T}_\gamma| + \frac{1}{\gamma} \sum_{t=1}^T R_t^2,$$

where $\mathcal{T}_\gamma = \{t : R_t \in (0, \gamma)\}$.

We use Lemma 8.4 to upper-bound the second summand in (8.6). To this end, we condition on the event that every component of every perturbation $\varepsilon_{a,t}$ has absolute value at most $\sqrt{2 \log 2T^3 K d}$; denote this event by U . This implies $\|x_{a,t}\|_2 \leq L$ for all actions a and all rounds t . By Lemma A.4, U is a high-probability event: $\Pr[U] \geq 1 - \frac{1}{T^2}$. Now we are ready to apply Lemma 8.4:

$$(8.7) \quad \mathbb{E} \left[\sum_{t=1}^T R_t^2 \mid U \right] \leq 16 d \beta_T \log(1 + tL^2/d).$$

To plug this into (8.6), we need to account for the low-probability event \bar{U} . We need to be careful because R_t could, with low probability, be arbitrarily large. By Lemma 8.7 with $\ell = 0$,

$$\begin{aligned} \mathbb{E}[R_t \mid \bar{U}] &\leq 2 \left[\|\theta\|_2 \left(1 + \rho(1 + \sqrt{2 \log K}) + \sqrt{2 \log(2T^3 K d)} \right) \right] \\ \mathbb{E}[\text{Regret}(T) \mid \bar{U}] \Pr[\bar{U}] &= \sum_{t=1}^T \mathbb{E}[R_t \mid \bar{U}] / T^2 < \tilde{O}(1). \\ \mathbb{E}[\text{Regret}(T) \mid U] \Pr[U] &\leq \gamma \mathbb{E}[|\mathcal{T}_\gamma|] + \frac{1}{\gamma} \mathbb{E} \left[\sum_{t=1}^T R_t^2 \mid U \right] \quad (\text{by (8.6)}) \end{aligned}$$

Putting this together and using (8.7), we obtain:

$$(8.8) \quad \mathbb{E}[\text{Regret}(T)] \leq \gamma \mathbb{E}[|\mathcal{T}_\gamma|] + \frac{16}{\gamma} d \beta_T \log(1 + TL^2/d) + \tilde{O}(1).$$

To obtain (8.5), we analyze the first summand in (8.8). Let Δ_t be the “gap” at time t : the difference in expected rewards between the best and second-best actions at time t (where “best” and “second-best” is according to expected rewards). Here, we’re taking expectations *after* the perturbations are applied, so the only randomness comes from the noisy rewards. Consider the set of rounds with small gap, $\mathcal{G}_\gamma := \{t : \Delta_t < \gamma\}$. Notice that $r_t \in (0, \gamma)$ implies $\Delta_t < \gamma$, so $|\mathcal{T}_\gamma| \leq |\mathcal{G}_\gamma|$.

In what follows we prove an upper bound on $\mathbb{E}[|\mathcal{G}_\gamma|]$. This is the step where perturbed context generation is truly used. For any two arms a_1 and a_2 , the gap between their expected rewards is

$$\theta^\top(x_{a_1,t} - x_{a_2,t}) = \theta^\top(\mu_{a_1,t} - \mu_{a_2,t}) + \theta^\top(\varepsilon_{a_1,t} - \varepsilon_{a_2,t}).$$

Therefore, the probability that the gap between those arms is smaller than γ is

$$\begin{aligned} \Pr[|\theta^\top(\mu_{a_1,t} - \mu_{a_2,t}) + \theta^\top(\varepsilon_{a_1,t} - \varepsilon_{a_2,t})| \leq \gamma] \\ = \Pr[-\gamma - \theta^\top(\mu_{a_1,t} - \mu_{a_2,t}) \leq \theta^\top(\varepsilon_{a_1,t} - \varepsilon_{a_2,t}) \leq \gamma - \theta^\top(\mu_{a_1,t} - \mu_{a_2,t})] \end{aligned}$$

Since $\theta^\top \varepsilon_{a_1,t}$ and $\theta^\top \varepsilon_{a_2,t}$ are both distributed as $\mathcal{N}(0, \rho^2 \|\theta\|_2^2)$, their difference is $\mathcal{N}(0, 2\rho^2 \|\theta\|_2^2)$. The maximum value that the Gaussian measure takes is $\frac{1}{2\rho \|\theta\|_2 \sqrt{\pi}}$, and the measure in any interval of width 2γ is therefore at most $\frac{\gamma}{\rho \|\theta\|_2 \sqrt{\pi}}$. This gives us the bound

$$\Pr[|\theta^\top(\mu_{a_1,t} - \mu_{a_2,t}) + \theta^\top(\varepsilon_{a_1,t} - \varepsilon_{a_2,t})| \leq \gamma] \leq \frac{\gamma}{\rho \|\theta\|_2 \sqrt{\pi}}.$$

Union-bounding over all $\binom{K}{2}$ pairs of actions, we have

$$\begin{aligned} \Pr[\Delta_t \leq \gamma] &\leq \Pr\left[\bigcup_{a_1, a_2 \in [K]} |\theta^\top(x_{a_1,t} - x_{a_2,t})| \leq \gamma\right] \leq \frac{K^2}{2} \frac{\gamma}{\rho \|\theta\|_2 \sqrt{\pi}}. \\ \mathbb{E}[|\mathcal{G}_\gamma|] &= \sum_{t=1}^T \Pr[\Delta_t \leq \gamma] \leq T \cdot \frac{K^2}{2} \frac{\gamma}{\rho \|\theta\|_2 \sqrt{\pi}}. \end{aligned}$$

Plugging this into (8.8) (recalling that $|\mathcal{T}_\gamma| \leq |\mathcal{G}_\gamma|$) completes the proof. \square

8.4. Bounding the Deviations. We make use of two results that bound deviations from the “high-probability” behavior, one on $\|\theta\|_2$ and another on instantaneous regret. First, we prove high-probability upper and lower bounds on $\|\theta\|_2$ under the conditions in Theorem 8.2. Essentially, these bounds allow us to use Lemma 8.5.

LEMMA 8.6. *Assume the latent vector θ comes from a multivariate Gaussian, $\theta \sim \mathcal{N}(\bar{\theta}, \Sigma)$, here the covariate matrix Σ satisfies $\lambda_{\max}(\Sigma) \leq 1$.*

(a) *If $\|\bar{\theta}\|_2 \geq 1 + \sqrt{3 \log T}$, then for sufficiently large T , with probability at least $1 - \frac{2}{T}$, it holds that*

$$(8.9) \quad \frac{1}{2 \log T} \leq \|\theta\|_2 \leq \|\bar{\theta}\|_2 + \sqrt{3d \log T}.$$

(b) *Same conclusion if $d \geq \frac{\log T}{\log \log T}$.*

Proof. We consider two cases, based on whether $d \geq \log T / \log \log T$. We need both cases to prove part (a), and we obtain part (b) as an interesting by-product. We repeatedly use Lemma A.7, a concentration inequality for χ^2 random variables, to show concentration on the Gaussian norm.

Case 1: $d \geq \log T / \log \log T$.

Since the Gaussian measure is decreasing in distance from 0, the $\Pr[\|\theta\|_2 \leq c] \leq \Pr[\|\bar{\theta} - \theta\|_2 \leq c]$ for any c . In other words, the norm of a Gaussian is most likely to be small when its mean is 0. Let $X = \Sigma^{-1/2}(\bar{\theta} - \theta)$. Note that X has distribution $\mathcal{N}(0, I)$, and therefore $\|X\|_2^2$ has χ^2 distribution with d degrees of freedom. We can

bound this as

$$\begin{aligned}
\Pr \left[\|\bar{\theta} - \theta\|_2 \leq \frac{1}{2 \log T} \right] &= \Pr \left[\|\Sigma^{-1/2} X\|_2 \leq \frac{1}{2 \log T} \right] \\
&\leq \Pr \left[\sqrt{\lambda_{\max}(\Sigma)} \|X\|_2 \leq \frac{1}{2 \log T} \right] \\
&\leq \Pr \left[\|X\|_2 \leq \frac{1}{2 \log T} \right] \\
&= \Pr \left[\|X\|_2^2 \leq \frac{1}{4(\log T)^2} \right] \\
&\leq \left(\frac{1}{4d(\log T)^2} e^{1-1/((4 \log T)^2 d)} \right)^{d/2} \\
&\leq \left(\frac{\log \log T}{(\log T)^3} \right)^{\log T / (2 \log \log T)} \\
&= \frac{T^{\log \log \log T / (2 \log \log T)}}{T^{3/2}} \\
&\leq T^{-1}
\end{aligned}$$

(By Lemma A.7)

($d \geq \log T / \log \log T$)

Similarly, we can show

$$\begin{aligned}
\Pr \left[\|\bar{\theta} - \theta\|_2 \geq \sqrt{d \log T} \right] &= \Pr \left[\|\Sigma^{-1/2} X\|_2 \geq \sqrt{d \log T} \right] \\
&\leq \Pr \left[\sqrt{\lambda_{\max}(\Sigma)} \|X\|_2 \geq \sqrt{d \log T} \right] \\
&\leq \Pr \left[\|X\|_2 \geq \sqrt{d \log T} \right] \\
&= \Pr \left[\|X\|_2^2 \geq d \log T \right] \\
&\leq (\log T e^{1-\log T})^{d/2} \\
&\leq (\exp(1 + \log \log T - \log T))^{\log T / (2 \log \log T)} \\
&= T^{(1+\log \log T - \log T) / (2 \log \log T)} \\
&\leq T^{-1}
\end{aligned}$$

(By Lemma A.7)

($d \geq \log T / \log \log T$)

for $\log T > 1 + 3 \log \log T$. By the triangle inequality,

$$\|\bar{\theta}\|_2 - \|\bar{\theta} - \theta\|_2 \leq \|\theta\|_2 \leq \|\bar{\theta}\|_2 + \|\bar{\theta} - \theta\|_2.$$

Thus, in this case, $\frac{1}{2 \log T} \leq \|\theta\|_2 \leq \|\bar{\theta}\|_2 + \sqrt{d \log T}$ with probability at least $1 - 2T^{-1}$.

Case 2: $\|\bar{\theta}\|_2 \geq 1 + \sqrt{3 \log T}$ and $d < \log T / \log \log T$.

For this part of the proof, we just need that $d < \log T$, which it is by assumption. Using the triangle inequality, if $\|\bar{\theta}\|_2$ is large, it suffices to show that $\|\bar{\theta} - \theta\|_2$ is small

with high probability. Again, let $X = \Sigma^{-1/2}(\bar{\theta} - \theta)$. Then,

$$\begin{aligned}
\Pr \left[\|\bar{\theta} - \theta\|_2 \geq \sqrt{3 \log T} \right] &= \Pr \left[\|\Sigma^{1/2} X\|_2 \geq \sqrt{3 \log T} \right] \\
&\geq \Pr \left[\sqrt{\lambda_{\max}(\Sigma)} \|X\|_2 \geq \sqrt{3 \log T} \right] \\
&= \Pr \left[\|X\|_2 \geq \frac{\sqrt{3 \log T}}{\sqrt{\lambda_{\max}(\Sigma)}} \right] \\
&\geq \Pr \left[\|X\|_2 \geq \sqrt{3 \log T} \right] \\
&= \Pr \left[\|X\|_2^2 \geq 3 \log T \right]
\end{aligned}$$

By Lemma A.7,

$$\begin{aligned}
\Pr \left[\|X\|_2^2 \geq 3 \log T \right] &\leq \left(\frac{3 \log T}{d} e^{1 - \frac{3 \log T}{d}} \right)^{d/2} \\
&= \left(T^{-3/d} e \frac{3 \log T}{d} \right)^{d/2} \\
&= T^{-1} \left(T^{-1/d} e \frac{3 \log T}{d} \right)^{d/2} \\
&\leq T^{-1}
\end{aligned}$$

(for sufficiently large T)

Because $\|\bar{\theta}\|_2 \geq 1 + \sqrt{3 \log T}$, $1 \leq \|\theta\|_2 \leq \|\bar{\theta}\|_2 + \sqrt{3 \log T}$ with probability at least $1 - T^{-1}$. \square

Next, we show how to upper-bound expected instantaneous regret in the worst case.¹⁰

LEMMA 8.7. *Fix round t and parameter $\ell > 0$. For any θ , conditioned on any history h_{t-1} and the event that $\|\varepsilon_{a,t}\|_\infty \geq \ell$ for each arm a , the expected instantaneous regret of any algorithm at round t is at most*

$$2 \|\theta\|_2 \left(1 + \rho(2 + \sqrt{2 \log K}) + \ell \right).$$

Proof. The expected regret at round t is upper-bounded by the reward difference between the best arm x_t^* and the worst arm x_t^\dagger , which is

$$\theta^\top (x_t^* - x_t^\dagger).$$

Note that $x_t^* = \mu_t^* + \varepsilon_t^*$ and $x_t^\dagger = \mu_t^\dagger + \varepsilon_t^\dagger$. Then, this is

$$\begin{aligned}
\theta^\top (x_t^* - x_t^\dagger) &= \theta^\top (\mu_t^* - \mu_t^\dagger) + \theta^\top (\varepsilon_t^* - \varepsilon_t^\dagger) \\
&\leq 2 \|\theta\|_2 + \theta^\top (\varepsilon_t^* - \varepsilon_t^\dagger)
\end{aligned}$$

since $\|\mu_{a,t}\|_2 \leq 1$. Next, note that

$$\theta^\top \varepsilon_t^* \leq \max_a \theta^\top \varepsilon_{a,t}$$

¹⁰We state and prove this result in a slightly more general version which we use to support Section 4. For the sake of this section, a special case of $\ell = 0$ suffices.

and

$$\theta^\top \varepsilon_t^\dagger \geq \min_a \theta^\top \varepsilon_{a,t}.$$

Since $\varepsilon_{a,t}$ has symmetry about the origin conditioned on the event that at least one component of one of the perturbations has absolute value at least ℓ , i.e. v and $-v$ have equal likelihood, $\max_a \theta^\top \varepsilon_{a,t}$ and $-\min_a \theta^\top \varepsilon_{a,t}$ are identically distributed. Let $E_{\ell,t}$ be the event that at least one of the components of one of the perturbations has absolute value at least ℓ . This means for any choice $\mu_{a,t}$ for all a ,

$$\mathbb{E} \left[\theta^\top (x_t^* - x_t^\dagger) \mid E_{\ell,t} \right] \leq 2\|\theta\|_2 + 2 \mathbb{E} \left[\max_a \theta^\top \varepsilon_{a,t} \mid E_{\ell,t} \right]$$

where the expectation is taken over the perturbations at time t .

Without loss of generality, let $(\varepsilon_{a',t})_j$ be the component such that $|(\varepsilon_{a',t})_j| \geq \ell$. Then, all other components have distribution $\mathcal{N}(0, \rho^2)$. Then,

$$\begin{aligned} & \mathbb{E} \left[\max_a \theta^\top \varepsilon_{a,t} \mid E_{\ell,t} \right] \\ &= \mathbb{E} \left[\max_a \theta^\top \varepsilon_{a,t} \mid |(\varepsilon_{a',t})_j| \geq \ell \right] \\ &= \mathbb{E} \left[\max(\theta^\top \varepsilon_{a',t}, \max_{a \neq a'} \theta^\top \varepsilon_{a,t}) \mid |(\varepsilon_{a',t})_j| \geq \ell \right] \\ &\leq \mathbb{E} \left[\max \left(|\theta_j(\varepsilon_{a',t})_j| + \sum_{i \neq j} \theta_i(\varepsilon_{a',t})_i, \max_{a \neq a'} \theta^\top \varepsilon_{a,t} \right) \mid |(\varepsilon_{a',t})_j| \geq \ell \right] \end{aligned}$$

Let $(\tilde{\varepsilon}_{a,t})_i = 0$ if $a = a'$ and $i = j$, and $(\varepsilon_{a,t})_i$ otherwise. In other words, we simply zero out the component $(\varepsilon_{a',t})_j$. Then, this is

$$\begin{aligned} & \mathbb{E} \left[\max \left(|\theta_j(\varepsilon_{a',t})_j| + \theta^\top \tilde{\varepsilon}_{a',t}, \max_{a \neq a'} \theta^\top \tilde{\varepsilon}_{a,t} \right) \mid |(\varepsilon_{a',t})_j| \geq \ell \right] \\ &\leq \mathbb{E} \left[\max_a (|\theta_j(\varepsilon_{a',t})_j| + \theta^\top \tilde{\varepsilon}_{a,t}) \mid |(\varepsilon_{a',t})_j| \geq \ell \right] \\ &= \mathbb{E} \left[|\theta_j(\varepsilon_{a',t})_j| + \max_a (\theta^\top \tilde{\varepsilon}_{a,t}) \mid |(\varepsilon_{a',t})_j| \geq \ell \right] \\ &= \mathbb{E} [|\theta_j(\varepsilon_{a',t})_j| \mid |(\varepsilon_{a',t})_j| \geq \ell] + \mathbb{E} \left[\max_a (\theta^\top \tilde{\varepsilon}_{a,t}) \right] \\ &\leq \mathbb{E} [|\theta_j(\varepsilon_{a',t})_j| \mid |(\varepsilon_{a',t})_j| \geq \ell] + \rho \|\theta\|_2 \sqrt{2 \log K} \end{aligned}$$

because by Lemma A.5,

$$\mathbb{E} \left[\max_a \theta^\top \tilde{\varepsilon}_{a,t} \right] \leq \mathbb{E} \left[\max_a \theta^\top \varepsilon_{a,t} \right] \leq \rho \|\theta\|_2 \sqrt{2 \log K}$$

Next, note that by symmetry and since $\theta_j \leq \|\theta\|_2$,

$$\mathbb{E} [|\theta_j(\varepsilon_{a',t})_j| \mid |(\varepsilon_{a',t})_j| \geq \ell] \leq \|\theta\|_2 \mathbb{E} [(\varepsilon_{a',t})_j \mid (\varepsilon_{a',t})_j \geq \ell].$$

By Lemma A.1,

$$\mathbb{E} [(\varepsilon_{a',t})_j \mid (\varepsilon_{a',t})_j \geq \ell] \leq \max(2\rho, \ell + \rho) \leq 2\rho + \ell$$

Putting this all together, the expected instantaneous regret is bounded by

$$2 \left(\|\theta\|_2 \left(1 + \rho(2 + \sqrt{2 \log K}) + \ell \right) \right),$$

proving the lemma. \square

8.5. Finishing the Proof of Theorem 8.2. We focus on the “nice event” that (8.9) holds, denote it \mathcal{E} for brevity. In particular, note that it implies $\|\theta\|_2 \leq S$. Lemma 8.5 guarantees that expected regret under this event, $\mathbb{E}[\text{Regret}(T) \mid \mathcal{E}]$, is upper-bounded by the expression (8.4) in the theorem statement.

In what follows we use Lemma 8.6(a) and Lemma 8.7 guarantee that if \mathcal{E} fails, then the corresponding contribution to expected regret is small. Indeed, Lemma 8.7 with $\ell = 0$ implies that

$$\mathbb{E}[R_t \mid \bar{\mathcal{E}}] \leq BT \|\theta\|_2 \quad \text{for each round } t,$$

where $B = 1 + \rho(2 + \sqrt{2 \log K})$ is the “blow-up factor”. Since (8.9) fails with probability at most $\frac{2}{T}$ by Lemma 8.6(a), we have

$$\begin{aligned} \mathbb{E}[\text{Regret}(T) \mid \bar{\mathcal{E}}] &\Pr[\bar{\mathcal{E}}] \leq \frac{2B}{T} \mathbb{E}[\|\theta\|_2 \mid \bar{\mathcal{E}}] \\ &\leq \frac{2B}{T} \mathbb{E}\left[\|\theta\|_2 \mid \|\theta\|_2 \geq \frac{1}{2 \log T}\right] \\ &\leq O\left(\frac{B}{T}\right) (\|\bar{\theta}\|_2 + d \log T) \\ &\leq O(1). \end{aligned}$$

The antecedent inequality follows by Lemma A.2 with $\alpha = \frac{1}{2 \log T}$, using the assumption that $\lambda_{\max}(\Sigma) \leq 1$. The theorem follows.

Acknowledgments. We thank Dylan Foster, Jon Kleinberg, and Aaron Roth for helpful discussions about these topics.

REFERENCES

- [1] Y. ABBASI-YADKORI, D. PÁL, AND C. SZEPESVÁRI, *Improved algorithms for linear stochastic bandits*, in 25th Advances in Neural Information Processing Systems (NIPS), 2011, pp. 2312–2320.
- [2] D. ACEMOGLU, A. MAKHDOUNI, A. MALEKIAN, AND A. OZDAGLAR, *Learning From Reviews: The Selection Effect and the Speed of Learning*, 2021. Conditionally accepted in *Econometrica*. Working paper available since 2017.
- [3] A. AGARWAL, S. BIRD, M. COZOWICZ, M. DUDIK, L. HOANG, J. LANGFORD, L. LI, D. MELAMED, G. OSHRI, S. SEN, AND A. SLIVKINS, *Multiworld testing: A system for experimentation, learning, and decision-making*, 2016. A white paper, available at <https://github.com/Microsoft/mwt-ds/raw/master/images/MWT-WhitePaper.pdf>.
- [4] A. AGARWAL, S. BIRD, M. COZOWICZ, L. HOANG, J. LANGFORD, S. LEE, J. LI, D. MELAMED, G. OSHRI, O. RIBAS, S. SEN, AND A. SLIVKINS, *Making contextual decisions with low technical debt*, 2017. Technical report at arxiv.org/abs/1606.03966.
- [5] A. AGARWAL, D. HSU, S. KALE, J. LANGFORD, L. LI, AND R. SCHAPIRE, *Taming the monster: A fast and simple algorithm for contextual bandits*, in 31st Intl. Conf. on Machine Learning (ICML), 2014.
- [6] P. AUER, *Using confidence bounds for exploitation-exploration trade-offs*, J. of Machine Learning Research (JMLR), 3 (2002), pp. 397–422. Preliminary version in *41st IEEE FOCS*, 2000.
- [7] P. AUER, N. CESA-BIANCHI, AND P. FISCHER, *Finite-time analysis of the multiarmed bandit problem.*, Machine Learning, 47 (2002), pp. 235–256.
- [8] H. BASTANI, M. BAYATI, AND K. KHOSRAVI, *Mostly exploration-free algorithms for contextual bandits*, Management Science, 67 (2021), pp. 1329–1349. Working paper available on arxiv.org since 2017.
- [9] A. BIETTI, A. AGARWAL, AND J. LANGFORD, *A contextual bandit bake-off*, CoRR arXiv:1802.04064, (2018).
- [10] S. BUBECK AND N. CESA-BIANCHI, *Regret Analysis of Stochastic and Nonstochastic Multi-armed Bandit Problems*, Foundations and Trends in Machine Learning, 5 (2012), pp. 1–122. Published with *Now Publishers* (Boston, MA, USA). Also available at <https://arxiv.org/abs/1204.5721>.

- [11] V. CHANDRASEKARAN, B. RECHT, P. A. PARRILO, AND A. S. WILLSKY, *The convex geometry of linear inverse problems*, Foundations of Computational Mathematics, 12 (2012), pp. 805–849.
- [12] W. CHU, L. LI, L. REYZIN, AND R. E. SCHAPIRE, *Contextual Bandits with Linear Payoff Functions*, in 14th Intl. Conf. on Artificial Intelligence and Statistics (AISTATS), 2011.
- [13] J. D. COOK, *Upper and lower bounds for the normal distribution function*, 2009.
- [14] V. DANI, T. P. HAYES, AND S. KAKADE, *Stochastic Linear Optimization under Bandit Feedback*, in 21th Conf. on Learning Theory (COLT), 2008, pp. 355–366.
- [15] S. DASGUPTA AND A. GUPTA, *An elementary proof of a theorem of Johnson and Lindenstrauss*, Random Structures & Algorithms, 22 (2003), pp. 60–65.
- [16] M. DUDÍK, D. ERHAN, J. LANGFORD, AND L. LI, *Sample-efficient nonstationary policy evaluation for contextual bandits*, in 28th Conf. on Uncertainty in Artificial Intelligence (UAI), 2012, pp. 247–254.
- [17] N. IMMORLICA, J. MAO, A. SLIVKINS, AND S. WU, *Incentivizing exploration with selective data disclosure*, in ACM Conf. on Economics and Computation (ACM-EC), 2020. Working paper available at <https://arxiv.org/abs/1811.06026>.
- [18] S. JANSON, *Tail bounds for sums of geometric and exponential variables*, Statistics Probability Letters, 135 (2018), pp. 1–6.
- [19] S. KANNAN, J. MORGENSTERN, A. ROTH, B. WAGGONER, AND Z. S. WU, *A smoothed analysis of the greedy algorithm for the linear contextual bandit problem*, in Advances in Neural Information Processing Systems (NIPS), 2018.
- [20] A. KRISHNAMURTHY, A. AGARWAL, AND M. DUDÍK, *Contextual semibandits via supervised learning oracles*, in 29th Advances in Neural Information Processing Systems (NIPS), 2016.
- [21] J. LANGFORD AND T. ZHANG, *The Epoch-Greedy Algorithm for Contextual Multi-armed Bandits*, in 21st Advances in Neural Information Processing Systems (NIPS), 2007.
- [22] T. LATTIMORE AND C. SZEPEŠVÁRI, *Bandit Algorithms*, Cambridge University Press, Cambridge, UK, 2020. Versions available at <https://banditalgs.com/> since 2018.
- [23] B. LAURENT AND P. MASSART, *Adaptive estimation of a quadratic functional by model selection*, Annals of Statistics, (2000), pp. 1302–1338.
- [24] L. LI, W. CHU, J. LANGFORD, AND R. E. SCHAPIRE, *A contextual-bandit approach to personalized news article recommendation*, in 19th Intl. World Wide Web Conf. (WWW), 2010.
- [25] L. LI, W. CHU, J. LANGFORD, AND X. WANG, *Unbiased offline evaluation of contextual-bandit-based news article recommendation algorithms*, in 4th ACM Intl. Conf. on Web Search and Data Mining (WSDM), 2011.
- [26] V. PERCHET, P. RIGOLLET, S. CHASSANG, AND E. SNOWBERG, *Batched bandit problems*, Ann. Statist., 44 (2016), pp. 660–681, <https://doi.org/10.1214/15-AOS1381>, <https://doi.org/10.1214/15-AOS1381>.
- [27] M. RAGHAVAN, A. SLIVKINS, J. W. VAUGHAN, AND Z. S. WU, *The externalities of exploration and how data diversity helps exploitation*, in Conf. on Learning Theory (COLT), 2018, pp. 1724–1738.
- [28] M. RAGHAVAN, A. SLIVKINS, J. W. VAUGHAN, AND Z. S. WU, *The externalities of exploration and how data diversity helps exploitation*, CoRR, abs/1806.00543 (2018), <http://arxiv.org/abs/1806.00543>, <https://arxiv.org/abs/1806.00543>.
- [29] P. RIGOLLET AND A. ZEEVI, *Nonparametric Bandits with Covariates*, in 23rd Conf. on Learning Theory (COLT), 2010, pp. 54–66.
- [30] A. SLIVKINS, *Introduction to multi-armed bandits*, Foundations and Trends® in Machine Learning, 12 (2019), pp. 1–286. Published with Now Publishers (Boston, MA, USA). Also available at <https://arxiv.org/abs/1904.07272>. Latest online revision: June 2021.
- [31] A. SLIVKINS, *Exploration and persuasion*, in Online and Matching-Based Market Design, F. Echenique, N. Immorlica, and V. Vazirani, eds., Cambridge University Press, 2021. To appear. Available at <http://slivkins.com/work/ExplPers.pdf>.
- [32] D. A. SPIELMAN AND S. TENG, *Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time*, J. of the ACM, 51 (2004), pp. 385–463.
- [33] J. A. TROPP, *User-friendly tail bounds for sums of random matrices*, Foundations of Computational Mathematics, 12 (2012), pp. 389–434.

Appendix A. Auxiliary Lemmas. Our proofs use a number of tools that are either known or easily follow from something that is known. We state these tools and provide the proofs for the sake of completeness.

A.1. (Sub)gaussians and Concentration. We rely on several known facts about Gaussian and subgaussian random variables. A random variable X is called σ -subgaussian, for some $\sigma > 0$, if $E[e^{\sigma X^2}] < \infty$. This includes variance- σ^2 Gaussian random variables as a special case.

LEMMA A.1. *If $X \sim \mathcal{N}(0, \sigma^2)$, then for any $t \geq 0$,*

$$\mathbb{E}[X \mid X \geq t] \leq \begin{cases} 2\sigma & t \leq \sigma \\ t + \frac{\sigma^2}{t} & t > \sigma \end{cases}$$

Proof. We begin with

$$(A.1) \quad \mathbb{E}[X \mid X \geq t] = \frac{\frac{1}{\sigma\sqrt{2\pi}} \int_t^\infty x \exp(x^2/(2\sigma^2)) dx}{\Pr[X \geq t]}.$$

X can be represented as $X = \sigma Y$, where Y is a standard normal random variable. Using a tail bound for the latter (from [13]),

$$\Pr[X \geq t] = \Pr\left[Y \geq \frac{t}{\sigma}\right] \geq \frac{1}{\sqrt{2\pi}} \frac{t/\sigma}{(t/\sigma)^2 + 1} \exp\left(-\frac{t^2}{2\sigma^2}\right).$$

The numerator in (A.1) is

$$\begin{aligned} \frac{1}{\sigma\sqrt{2\pi}} \int_t^\infty x \exp(x^2/(2\sigma^2)) dx &= -\frac{1}{\sigma\sqrt{2\pi}} \cdot \sigma^2 e^{-x^2/(2\sigma^2)} \Big|_t^\infty \cdot e^{-t^2/(2\sigma^2)} \\ &= \frac{\sigma}{\sqrt{2\pi}} \exp\left(-\frac{t^2}{2\sigma^2}\right). \end{aligned}$$

Combining, we have

$$\mathbb{E}[X \mid X \geq t] \leq \frac{\frac{\sigma}{\sqrt{2\pi}} \exp\left(-\frac{t^2}{2\sigma^2}\right)}{\frac{1}{\sqrt{2\pi}} \frac{t/\sigma}{(t/\sigma)^2 + 1} \exp\left(-\frac{t^2}{2\sigma^2}\right)} = \frac{\sigma^2((t/\sigma)^2 + 1)}{t} = t + \frac{\sigma^2}{t}$$

For $t \leq \sigma$, $\mathbb{E}[X \mid X \geq t] \leq \mathbb{E}[X \mid X \geq \sigma] \leq 2\sigma$ by the above bound. \square

LEMMA A.2. *Suppose $X \sim \mathcal{N}(0, \Sigma)$ is a Gaussian random vector with covariance matrix Σ . Then*

$$\mathbb{E}[\|X\|_2 \mid \|X\|_2 > \alpha] \leq d \left(\alpha + \frac{\lambda_{\max}(\Sigma)}{\alpha} \right) \quad \text{for any } \alpha \geq 0.$$

Proof. Assume without loss of generality that Σ is diagonal, since the norm is rotationally invariant. Observe that $\|X\|_2 \mid \forall i X_i > \alpha$ stochastically dominates $\|X\|_2 \mid \|X\|_2 > \alpha$. (Geometrically, the latter conditioning shifts the probability mass away from the origin.) Therefore,

$$\begin{aligned} \mathbb{E}[\|X\|_2 \mid \|X\|_2 > \alpha] &\leq \mathbb{E}[\|X\|_2 \mid \forall i X_i > \alpha] \\ &= \mathbb{E}\left[\sum_{i=1}^d X_i \mid \forall i X_i > \alpha\right] \leq \sum_{i=1}^d \left(t + \frac{\lambda_i(\Sigma)}{\alpha}\right) \end{aligned}$$

by Lemma A.1, where $\lambda_i(\Sigma) \leq \lambda_{\max}(\Sigma)$ is the i th eigenvalue of Σ . \square

FACT A.3. If X is a σ -subgaussian random variable, then

$$\Pr[|X - \mathbb{E}[X]| > t] \leq 2e^{-t^2/(2\sigma^2)}.$$

LEMMA A.4. If X_1, \dots, X_n are independent σ -subgaussian random variables, then

$$\Pr \left[\max_i |X_i - \mathbb{E}[X_i]| > \sigma \sqrt{2 \log \frac{2n}{\delta}} \right] \leq \delta.$$

Proof. For any X_i , we know from Fact A.3 that

$$\Pr \left[|X_i - \mathbb{E}[X_i]| > \sigma \sqrt{2 \log \frac{2n}{\delta}} \right] \leq 2 \exp \left(-\frac{2\sigma^2 \log \frac{2n}{\delta}}{2\sigma^2} \right) = 2 \exp \left(-\log \frac{2n}{\delta} \right) = \frac{\delta}{n}.$$

A union bound completes the proof. \square

LEMMA A.5. If X_1, \dots, X_K are independent zero-mean σ -subgaussian random variables, then

$$\mathbb{E} [\max_i X_i] \leq \sigma \sqrt{2 \log K}.$$

Proof. Let $X = \max_i X_i$. Since each X_i is σ -subgaussian, it follows that

$$\mathbb{E} [e^{\lambda X_i}] \leq \exp \left(\frac{\lambda^2 \sigma^2}{2} \right).$$

Using Jensen's inequality, we have

$$\exp(\lambda \mathbb{E}[X]) \leq \mathbb{E}[\exp(\lambda X)] = \mathbb{E} \left[\max_i \exp(\lambda X_i) \right] \leq \sum_i \mathbb{E}[\exp(\lambda X_i)] \leq K e^{\lambda^2 \sigma^2 / 2}.$$

Rearranging, we have

$$\mathbb{E}[X] \leq \frac{\log K}{\lambda} + \frac{\lambda \sigma^2}{2}.$$

Setting $\lambda = \frac{\sqrt{2 \log K}}{\sigma}$, we have $\mathbb{E}[X] \leq \sigma \sqrt{2 \log K}$ as needed \square

LEMMA A.6. If $\theta \sim \mathcal{N}(\bar{\theta}, \Sigma)$ where $\bar{\theta} \in \mathbb{R}^d$ and $\Sigma \in \mathbb{R}^{d \times d}$, then $\mathbb{E} [\|\theta - \bar{\theta}\|_2] \leq \sqrt{d \lambda_{\max}(\Sigma)}$.

Proof. From [11], the expected norm of a standard normal d -dimensional Gaussian is at most \sqrt{d} . Using the fact that $\Sigma^{-1/2}(\theta - \bar{\theta}) \sim \mathcal{N}(0, I)$, we have

$$\begin{aligned} \mathbb{E} [\|\theta - \bar{\theta}\|_2] &= \mathbb{E} [\|\Sigma^{1/2} \Sigma^{-1/2}(\theta - \bar{\theta})\|_2] \\ &\leq \|\Sigma^{1/2}\|_2 \mathbb{E} [\|\Sigma^{-1/2}(\theta - \bar{\theta})\|_2] \leq \sqrt{d \lambda_{\max}(\Sigma)}. \end{aligned} \quad \square$$

LEMMA A.7 (Lemma 2.2 in [15]). If $X \sim \chi^2(d)$, i.e., $X = \sum_{i=1}^d X_i^2$, where X_1, \dots, X_d are independent standard Normal random variables, then

$$\begin{aligned} \Pr[X \leq \beta d] &\leq (\beta e^{1-\beta})^{d/2} && \text{for any } \beta \in (0, 1), \\ \Pr[X \geq \beta d] &\leq (\beta e^{1-\beta})^{d/2} && \text{for any } \beta > 1. \end{aligned}$$

LEMMA A.8 (Hoeffding bound). *If $\bar{X} = \frac{1}{n} \sum_{i=1}^n X_i$, where the X_i 's are independent σ -subgaussian random variables with zero mean, then*

$$\begin{aligned} \max(\Pr[\bar{X} \geq t], \Pr[\bar{X} \leq -t]) &\leq \exp\left(-\frac{nt^2}{2\sigma^2}\right) \quad \text{for all } t > 0, \\ \max\left(\Pr\left[\bar{X} \leq -\sigma\sqrt{\frac{2}{n}\log\frac{1}{\delta}}\right], \Pr\left[\bar{X} \geq \sigma\sqrt{\frac{2}{n}\log\frac{1}{\delta}}\right]\right) &\leq \delta \quad \text{for all } \delta > 0. \end{aligned}$$

A.2. Linear Algebra. We use several facts from linear algebra. In what follows, recall that $\lambda_{\min}(M)$ and $\lambda_{\max}(M)$ denote the minimal and the maximal eigenvalues of matrix M , resp. For two matrices A, B , let us write $B \succeq A$ to mean that $B - A$ is positive semidefinite.

LEMMA A.9. $\lambda_{\max}(vv^\top) = \|v\|_2^2$ for any $v \in \mathbb{R}^d$.

Proof. vv^\top has rank one, so it has one eigenvector with nonzero eigenvalue. v is an eigenvector since $(vv^\top)v = (v^\top v)v$, and it has eigenvalue $v^\top v = \|v\|_2^2$. This is the only nonzero eigenvalue, so $\lambda_{\max}(vv^\top) = \|v\|_2^2$. \square

LEMMA A.10. *For symmetric matrices A, B with B invertible,*

$$B \succeq A \iff I \succeq B^{-1/2}AB^{-1/2}$$

Proof.

$$\begin{aligned} (\forall x) \quad B \succeq A &\iff x^\top Bx \geq x^\top Ax \\ (\forall x) \quad &\iff x^\top (B - A)x \geq 0 \\ (\forall x) \quad &\iff x^\top B^{1/2}(I - B^{-1/2}AB^{-1/2})B^{1/2}x \geq 0 \\ (\forall x) \quad &\iff x^\top (I - B^{-1/2}AB^{-1/2})x \geq 0 \\ &\iff I \succeq B^{-1/2}AB^{-1/2}. \end{aligned} \quad \square$$

LEMMA A.11. *If $A \succeq 0$ and $B \succeq 0$, then $\lambda_{\min}(A + B) \geq \lambda_{\min}(A)$.*

Proof.

$$\begin{aligned} \lambda_{\min}(A + B) &= \min_{\|x\|_2=1} x^\top (A + B)x \\ &= \min_{\|x\|_2=1} x^\top Ax + x^\top Bx \\ (\text{because } x^\top Bx \geq 0) \quad &\geq \min_{\|x\|_2=1} x^\top Ax \\ &= \lambda_{\min}(A) \end{aligned} \quad \square$$

A.3. Logarithms. We use several variants of standard inequalities about logarithms.

LEMMA A.12. $x \geq \log(ex)$ for all $x > 0$.

Proof. Equivalently, $x - \log(ex) \geq 0$ for $x > 0$. To show this, observe that

1. At $x = 1$, this holds with equality.
2. At $x = 1$, the derivative is

$$\left. \frac{d}{dx} x - \log(ex) \right|_{x=1} = 1 - \frac{1}{x} \Big|_{x=1} = 0.$$

3. The entire function is convex for $x > 0$, since

$$\frac{d^2}{dx^2}x - \log(ex) = \frac{d}{dx}1 - \frac{1}{x} = \frac{1}{x^2} > 0.$$

This proves the lemma. \square

COROLLARY A.13. $x - \log x \geq \frac{e-1}{e}x$.

Proof. Using Lemma A.12 and letting $z = x/e$,

$$x - \log x = \frac{e-1}{e}x + \frac{1}{e}x - \log x = \frac{e-1}{e}x + z - \log(ez) \geq \frac{e-1}{e}x \quad \square$$

LEMMA A.14. $\log\left(\frac{1}{1-x}\right) \leq \frac{7x}{6}$ for any $x \in [0, 1/4]$.

Proof. First, we note that

$$\frac{d}{dx} \log\left(\frac{1}{1-x}\right) = 1 - x(-(1-x)^{-2}) \cdot (-1) = \frac{1}{1-x} = \sum_{i=0}^{\infty} x^i.$$

Integrating both sides, we have

$$\log\left(\frac{1}{1-x}\right) = C + \sum_{i=0}^{\infty} \frac{x^i}{i},$$

for some constant C that does not depend on x . Taking $x = 0$ yields $C = 0$. Therefore,

$$\log\left(\frac{1}{1-x}\right) \leq x + \frac{x^2}{2} \sum_{i=0}^{\infty} x^i = x + \frac{x^2}{2(1-x)} = x \left(1 + \frac{x}{2(1-x)}\right) \leq \frac{7x}{6}. \quad \square$$