

# Anytime Model Selection in Linear Bandits

**Parnian Kassraie**

*ETH Zurich*

PKASSRAIE@ETHZ.CH

**Nicolas Emmenegger**

*ETH Zurich*

NICOLAS.EMMENEGGER@INF.ETHZ.CH

**Andreas Krause**

*ETH Zurich*

KRAUSEA@ETHZ.CH

**Aldo Pacchiano**

*Broad Institute*

APACCHIA@BROADINSTITUTE.ORG

## Abstract

Model selection in the context of bandit optimization is a challenging problem, as it requires balancing exploration and exploitation not only for action selection, but also for model selection. One natural approach is to rely on online learning algorithms that treat different models as experts. Existing methods, however, scale poorly (poly $M$ ) with the number of models  $M$  in terms of their regret. We develop ALEXP, an anytime algorithm, which has an exponentially improved ( $\log M$ ) dependence on  $M$  for its regret. We neither require knowledge of the horizon  $n$ , nor rely on an initial purely exploratory stage. Our approach utilizes a novel time-uniform analysis of the Lasso, by defining a self-normalized martingale sequence based on the empirical process error, establishing a new connection between interactive learning and high-dimensional statistics.

**Keywords:** Model Selection, Linear Bandits, Anytime Martingale Bounds

## 1. Introduction

When solving bandit problems or performing Bayesian optimization, we need to commit to a reward model *a priori*, based on which we estimate the reward function and build a policy for selecting the next action. In practice, there are many ways to model the reward by considering different feature maps or hypothesis spaces. However, is not known a priori which model is going to yield the most sample efficient bandit algorithm, and we can only select the right model as we gather empirical evidence. This leads us to ask, can we perform adaptive model selection, while simultaneously optimizing for reward?

In an idealized setting with no sampling limits, given a model class of size  $M$ , we could initialize  $M$  bandit algorithms (a.k.a agents) in parallel, each using one of the available reward models. Then, as the algorithms run, at every step we can select the most promising agent, according to the cumulative rewards that are obtained so far. Model selection can then be cast into an online optimization problem on a  $M$ -dimensional probability simplex, where the probability of selecting an agent is dependent on its cumulative reward, and the optimizer seeks to find the distribution with the best return in hindsight. This approach is impractical for large model classes, since at every step, it requires drawing  $M$  different samples in parallel from the environment so that the reward for each agent is realized.

In realistic applications of bandit optimization, we can only draw *one* sample at a time, and so we need to design an algorithm which allocates more samples to the agents that are deemed more promising. Prior work (e.g., [Maillard and Munos, 2011](#); [Agarwal et al., 2017](#)) propose to run a single meta algorithm which interacts with the environment by first selecting an agent, and then selecting an action according to the suggestion of that agent. The online model selection problem can still be emulated in this setup, however this time the optimizer receives partial feedback, coming from only one agent. Consequently, many agents need to be queried, and the overall regret scales with  $\text{poly}M$ , restricting this approach to small model classes. Addressing the limited scope of such algorithms, [Agarwal et al. \(2017\)](#) raise an open problem on the feasibility of obtaining a  $\log M$  dependency for the regret.

We show that this rate is achievable, in the particular case of linearly parametrizable rewards. We develop a technique to “hallucinate” the reward for every agent that was not selected, and run the online optimizer with emulated full-information feedback. This allows the optimizer to assess the quality of the agents, without ever having queried them. As a result, our algorithm ALEXP, satisfies a regret of rate  $\mathcal{O}(\max\{\sqrt{n \log^3 M}, n^{3/4} \sqrt{\log M}\})$ , with high probability, simultaneously for all  $n \geq 1$  ([Theorem 1](#)). Our key idea, leading to  $\log M$  dependency, is to employ the Lasso as a low-variance online regression oracle, and estimate the reward for the agents that were not chosen. This trick is made possible through our novel analysis of online Lasso regression ([Theorem 3](#)), where we recognize that the empirical process error is controlled by a sub-Gamma martingale sequence. We establish empirically that ALEXP has favorable exploration–exploitation dynamics, and consistently outperforms prior work across a range of environments ([Appendix F](#)).

## 2. Related Work

Online Model selection (MS) for bandits considers combining a number of agents in a master algorithm, with the goal of performing as well as the best agent ([Maillard and Munos, 2011](#); [Bubeck et al., 2012](#); [Agarwal et al., 2017](#); [Pacchiano et al., 2020](#); [Luo et al., 2022](#)). This literature operates on black-box model classes of size  $M$  and uses variants of Online Mirror Descent (OMD) to sequentially select the agents. Effectively, the optimizer receives partial (a.k.a. bandit) feedback and agents are at risk of starvation, since at every step only the selected agent gets the new data point.

Sparse linear bandits use sparsity-inducing methods, often Lasso ([Tibshirani, 1996](#)), for estimating the reward in presence of many features, as an alternative to model selection. Recent work often consider the contextual case, where at every step only a finite stochastic subset of the action domain is presented to the agent ([Li et al., 2022](#); [Bastani and Bayati, 2020](#); [Kim and Paik, 2019](#); [Oh et al., 2021](#)). Most sparse bandit algorithms either start with a purely exploratory phase ([Kim and Paik, 2019](#); [Li et al., 2022](#); [Hao et al., 2020](#); [Jang et al., 2022](#)), the length of which is set according to the horizon  $n$ . Therefore, such algorithms inherently require the knowledge of  $n$  and can be made anytime only via the doubling trick ([Auer et al., 1995](#)). [Appendix A.1](#) presents an in-depth overview.

ALEXP inherits the best of both worlds: its regret enjoys the  $\log M$  dependency of sparse linear bandits even on compact domains, and it has adaptive probabilistic exploration with anytime guarantees. In contrast to prior literature, we perform model selection with an online optimizer (EXP4), which hallucinates full-information feedback using a low-variance Lasso

estimator. Moreover, our anytime approach lifts the horizon dependence and the exploration requirement of sparse linear bandits. Our analysis relies on the seminal results of Howard et al. (2021), that present time-uniform curved boundaries for sub-Gamma martingale processes.

### 3. Problem Setting

We consider a bandit problem where a learner interacts with the environment in rounds. At step  $t$  the learner selects an action  $\mathbf{x}_t \in \mathcal{X}$ , where  $\mathcal{X} \subset \mathbb{R}^{d_0}$  is a compact domain and observes a noisy reward  $y_t = r(\mathbf{x}_t) + \varepsilon_t$  such that  $\varepsilon_t$  is an i.i.d. zero-mean sub-Gaussian variable with parameter  $\sigma^2$ . We assume the reward function  $r : \mathcal{X} \rightarrow \mathbb{R}$  is linearly parametrizable by some unknown feature map, and that the model class  $\{\phi_j : \mathbb{R}^{d_0} \rightarrow \mathbb{R}^d, j = 1, \dots, M\}$  contains the set of plausible feature maps. We consider the setting where  $M$  can be very large, and while the set  $\{\phi_j\}$  may include misspecified feature maps, it contains at least one feature map that represents the reward function, i.e., there exists  $j^* \in [M]$  such that  $r(\cdot) = \boldsymbol{\theta}_{j^*}^\top \phi_{j^*}(\cdot)$ . We assume the image of  $\phi_j$  spans  $\mathbb{R}^d$ , and no two feature maps are linearly dependent, i.e. for any  $j, j' \in [M]$ , there exists no  $\alpha \in \mathbb{R}$  such that  $\phi_j(\cdot) = \alpha \phi_{j'}(\cdot)$ . This assumption, which is satisfied by design in practice, ensures that the features are not ill-posed and we can explore in all relevant directions. We assume that the concatenated feature map  $\boldsymbol{\phi}(\mathbf{x}) := (\phi_1(\mathbf{x}), \dots, \phi_M(\mathbf{x}))$  is normalized  $\|\boldsymbol{\phi}(\mathbf{x})\| \leq 1$  for all  $\mathbf{x} \in \mathcal{X}$  and that  $\|\boldsymbol{\theta}_{j^*}\| \leq B$ , which implies  $|r(\mathbf{x})| \leq B$  for all  $\mathbf{x} \in \mathcal{X}$ .

We will model this problem in the language of model selection where a meta algorithm aims to optimize the unknown reward function by relying on a number of base learners. In order to interact with the environment the meta algorithm selects an agent that in turn selects an action. In our setting we thinking of each of these  $M$  feature maps as controlled by a base agent running its own algorithm. Base agent  $j$  uses the feature map  $\phi_j$  for modeling the reward. At step  $t$  of the bandit problem, each agent  $j$  is given access to the full history  $H_{t-1} := \{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_{t-1}, y_{t-1})\}$ , and uses it to locally estimate the reward as  $\hat{\boldsymbol{\beta}}_{t-1,j}^\top \phi_j(\cdot)$ , where  $\hat{\boldsymbol{\beta}}_{t-1,j} \in \mathbb{R}^d$  is the estimated coefficients vector. The agent then uses this estimate to develop its action selection policy  $p_{t,j} \in \mathcal{M}(\mathcal{X})$ . Here,  $\mathcal{M}$  denotes the space of probability measures defined on  $\mathcal{X}$ . The condition on existence of  $j^*$  will ensure that there is at least one agent which is using a correct model for the reward, and thereby can solve the bandit problem if executed in isolation. We refer to agent  $j^*$  as the oracle agent.

Our goal is to find a sample-efficient strategy for iterating over the agents, such that their suggested actions maximize the cumulative reward, achieved over any horizon  $n \geq 1$ . This is equivalent to minimizing the cumulative regret  $R(n) = \sum_{t=1}^n r(\mathbf{x}^*) - r(\mathbf{x}_t)$ , where  $\mathbf{x}^*$  is a global maximizer of the reward function. We neither fix  $n$ , nor assume knowledge of it.

### 4. Method

We propose the **Anytime EXPONENTIAL** weighting algorithm based on **Lasso** reward estimates (ALEXP), summarized in Algorithm 1. At step  $t$  we first sample an agent  $j_t$ , and then sample an action  $\mathbf{x}_t$  according to the agent's policy  $p_{t,j_t}$ . Let  $\Delta_M$  denote the  $M$ -dimensional probability simplex. We maintain a probability distribution  $\mathbf{q}_t \in \Delta_M$  over the agents, and update it sequentially as we accumulate evidence on the performance of each agent. Ideally, we would have adjusted  $q_{t,j}$  according to the average return of model  $j$ , that is,  $\mathbb{E}_{\mathbf{x} \sim p_{t,j}} r(\mathbf{x})$ . However, since  $r$  is unknown, we estimate the average reward with some  $\hat{r}_{t,j}$ . We then update

$q_t$  for the next step via,

$$q_{t+1,j} = \frac{\exp(\eta_t \sum_{s=1}^t \hat{r}_{s,j})}{\sum_{i=1}^M \exp(\eta_t \sum_{s=1}^t \hat{r}_{s,i})}$$

for all  $j = 1, \dots, M$ , where  $\eta_t$  is the learning rate, and controls the sensitivity of the updates. This rule allows us to imitate the full-information example that we mentioned above. By utilizing  $\hat{r}_{t,j}$  and hallucinating feedback from all agents, we can reduce the probability of selecting a badly performing agent, without ever having sampled them (c.f. Figure 5). It remains to design the estimator  $\hat{r}_{t,j}$ . We concatenate all feature maps, and, knowing that many features are redundant, use a sparsity inducing estimator over the resulting coefficients vector. Mainly, let  $\boldsymbol{\theta} = (\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_M) \in \mathbb{R}^{M^d}$  be the concatenated coefficients vector. We then solve

$$\hat{\boldsymbol{\theta}}_t = \arg \min_{\boldsymbol{\theta} \in \mathbb{R}^{M^d}} \mathcal{L}(\boldsymbol{\theta}; H_t, \lambda_t) = \arg \min_{\boldsymbol{\theta} \in \mathbb{R}^{M^d}} \frac{1}{t} \|\mathbf{y}_t - \Phi_t \boldsymbol{\theta}\|_2^2 + 2\lambda_t \sum_{j=1}^M \|\boldsymbol{\theta}_j\|_2 \quad (1)$$

where  $\Phi_t = [\boldsymbol{\phi}^\top(\mathbf{x}_s)]_{s \leq t} \in \mathbb{R}^{t \times M^d}$  is the feature matrix,  $\mathbf{y}_t \in \mathbb{R}^t$  is the concatenated reward vector, and  $\lambda_t$  is an adaptive regularization parameter. Problem (1) is the online variant of the group Lasso (Lounici et al., 2011). The second term is the loss is the mixed  $(2, 1)$ -norm of  $\boldsymbol{\theta}$ , which can be seen as the  $\ell_1$ -norm of the vector  $(\|\boldsymbol{\theta}_1\|, \dots, \|\boldsymbol{\theta}_M\|) \in \mathbb{R}^M$ . This norm induces sparsity at the group level, and therefore, the sub-vectors  $\hat{\boldsymbol{\theta}}_{t,j} \in \mathbb{R}^d$  that correspond to redundant feature maps are expected to be  $\mathbf{0}$ , i.e. the null vector. We then estimate the average return of each model by simply taking an expectation  $\hat{r}_{t,j} = \mathbb{E}_{\mathbf{x} \sim p_{t+1,j}}[\hat{\boldsymbol{\theta}}_t^\top \boldsymbol{\phi}(\mathbf{x})]$ . This quantity is the average return of the agent's policy  $p_{t+1,j}$ , according to our Lasso estimator. Appendix A.3 explains why the choice of Lasso is crucial for obtaining a  $\log M$  rate for the regret.

For action selection, with probability  $\gamma_t$ , we sample agent  $j$  with probability  $q_{t,j}$  and draw  $\mathbf{x}_t \sim p_{t,j}$  as per suggestion of the agent. With probability  $1 - \gamma_t$ , we choose the action according to some exploratory distribution  $\pi \in \mathcal{M}(\mathcal{X})$  which aims to sample informative actions. This can be any design where  $\text{supp}(\pi) = \mathcal{X}$ . We mix  $p_{t,j}$  with  $\pi$ , to collect sufficiently diverse data for model selection. We are not restricting the agents' policy, and therefore can not rely on them to explore adequately.

## 5. Main Results

For the regret guarantee, we consider specific choices of base agents and exploratory distribution. Our analysis may be extended to include other policies, since ALEXP can be wrapped around any bandit agent that is described by some  $p_{t,j}$ , and allows for random exploration with any distribution  $\pi$ .

**Base Agents.** We assume that the oracle agent has either a UCB (Abbasi-Yadkori et al., 2011) or a GREEDY (Auer et al., 2002a) policy, and all other agents are free to choose *any arbitrary* policy. Similar treatment can be applied to cases where the oracle uses other (sublinear) policies for solving linear or contextual bandits (e.g., Thompson, 1933; Kaufmann et al., 2012; Agarwal et al., 2014). In either case, agent  $j^*$  calculates a ridge estimate of the coefficients vector based on the history  $H_t$

$$\hat{\boldsymbol{\beta}}_{t,j^*} := \arg \min_{\boldsymbol{\beta} \in \mathbb{R}^d} \|\mathbf{y}_t - \Phi_{t,j^*} \boldsymbol{\beta}\|_2^2 + \tilde{\lambda} \|\boldsymbol{\beta}\|_2^2 = \left( \Phi_{t,j^*}^\top \Phi_{t,j^*} + \tilde{\lambda} \mathbf{I} \right)^{-1} \Phi_{t,j^*}^\top \mathbf{y}_t.$$

Here  $\Phi_{t,j^*} \in \mathbb{R}^{t \times d}$  is the feature matrix, where each row  $s$  is  $\phi_{j^*}^\top(\mathbf{x}_s)$  and  $\tilde{\lambda}$  is the regularization constant. Then at step  $t$ , a GREEDY oracle suggests the action which maximizes the reward estimate  $\hat{\beta}_{t-1,j^*}^\top \phi_{j^*}(\mathbf{x})$ , and a UCB oracle queries  $\arg \max u_{t-1,j^*}(\mathbf{x})$  where  $u_{t-1,j^*}(\cdot)$  is the upper confidence bound that this agent calculates for  $r(\cdot)$ . Proposition 21 shows that the sequence  $(u_{t,j^*})_{t \geq 1}$  is in fact an anytime valid upper bound for  $r$  over the entire domain.

**Exploratory policy.** Performance of ALEXP depends on the quality of the samples that  $\pi$  suggests. The eigenvalues of the covariance matrix  $\Sigma(\pi, \phi) := \mathbb{E}_{\mathbf{x} \sim \pi} \phi(\mathbf{x}) \phi^\top(\mathbf{x})$  reflect how diverse the data is, and thus are a good indicator for data quality. van de Geer and Bühlmann (2009) present a survey on the notions of diversity defined based on these eigenvalues. Let  $\lambda_{\min}(A)$  denote the minimum eigenvalue of a matrix  $A$ . Similar to Hao et al. (2020), we assume that  $\pi$  is the maximizer of the problem below and present our regret bounds in terms of

$$C_{\min} = C_{\min}(\mathcal{X}, \phi) := \max_{\pi \in \mathcal{M}(\mathcal{X})} \lambda_{\min}(\Sigma(\pi, \phi)), \quad (2)$$

which is greater than zero under the conditions specified in our problem setting. Prior works in the sparse bandit literature all require a similar or stronger assumption of this kind, and Table 1 gives an overview. Alternatively, one can work with an arbitrary  $\pi$ , e.g.,  $\text{Unif}(\mathcal{X})$ , as long as  $\lambda_{\min}(\Sigma(\pi, \phi))$  is bounded away from zero. Appendix C.1 reviews some configurations of  $(\phi, \mathcal{X}, \pi)$  that lead to a non-zero minimum eigenvalue, and Corollary 12 bounds the regret of ALEXP with uniform exploration.

For this choice of agents and exploratory distribution, Theorem 1 presents an informal regret bound. Here, we have used the  $\mathcal{O}$  notation, and only included the fastest growing terms. The inequality is made exact in Theorem 14, up to constant multiplicative factors.

**Theorem 1** (Cumulative Regret of ALEXP, Informal). *Let  $\delta \in (0, 1]$  and set  $\pi$  to be the maximizer of (2). Choose learning rate  $\eta_t = \mathcal{O}(C_{\min} t^{-1/2} / C(M, \delta, d))$ , exploration probability  $\gamma_t = \mathcal{O}(t^{-1/4})$  and Lasso regularization parameter  $\lambda_t = \mathcal{O}(C_{\min} t^{-1/2} C(M, \delta, d))$ , where*

$$C(M, \delta, d) = \mathcal{O} \left( \sqrt{1 + \log(M/\delta) + (\log \log d)_+ + \sqrt{d(\log(M/\delta) + (\log \log d)_+)}} \right).$$

*Then ALEXP satisfies the cumulative regret*

$$R(n) = \mathcal{O} \left( n^{3/4} B + C(M, \delta, d) \left( n^{3/4} + C_{\min}^{-1} \sqrt{n} \log M \right) + C_{\min}^{-1/2} n^{5/8} \sqrt{d \log(n) + \log(1/\delta)} \right)$$

*simultaneously for all  $n \geq 1$ , with probability greater than  $1 - \delta$ .*

In this bound, the first term is the regret incurred at exploratory steps (when  $\alpha_t = 1$ ), the second term is due to the estimation error of Lasso (i.e.,  $\|\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_t\|$ ), and the third term is the regret of the exponential weights sub-algorithm. The fourth term, is the regret bound for the oracle agent  $j^*$ , when run within the ALEXP framework. It does not depend on the agent's policy (greedy or optimistic), and is worse than the minimax optimal rate of  $\sqrt{nd \log n}$ . This is because the oracle is suggesting actions based on the history  $H_t$ , which includes uninformative action-reward pairs queried by other, potentially misspecified, agents. In Corollary 12, we provide a regret bound independent of  $C_{\min}$ , for orthogonal feature maps, and show that the same  $\mathcal{O}(\max\{\sqrt{n \log^3 M}, n^{3/4} \sqrt{\log M}\})$  rate may be achieved even with the simple choice  $\pi = \text{Unif}(\mathcal{X})$ . We discuss some implications of this theorem in Appendix A.3, and provide empirical evaluations in Appendix F.

## 6. Time-Uniform Lasso

We develop novel confidence sequences for Lasso with history-dependent data, which are uniformly valid over an unbounded time horizon. We use it for bounding bias and variance terms in the model selection regret decomposition, however, this result may be of independent interest in applications of Lasso for online learning or sequential decision making. The width of the Lasso confidence intervals depends on the quality of feature matrix  $\Phi_t$ , often quantified by the restricted eigenvalue property (Bickel et al., 2009; van de Geer and Bühlmann, 2009):

**Definition 2.** For the feature matrix  $\Phi_t \in \mathbb{R}^{t \times dM}$  we define  $\kappa(\Phi_t, s)$  for any  $1 \leq s \leq M$  as

$$\kappa(\Phi_t, s) := \inf_{(J, \mathbf{b})} \frac{1}{\sqrt{t}} \frac{\|\Phi_t \mathbf{b}\|_2}{\sqrt{\sum_{j \in J} \|\mathbf{b}_j\|_2^2}}$$

$$\text{s.t. } \mathbf{b} \in \mathbb{R}^d \setminus \{0\}, \sum_{j \notin J} \|\mathbf{b}_j\|_2 \leq 3 \sum_{j \in J} \|\mathbf{b}_j\|_2, J \subset \{1, \dots, M\}, |J| \leq s.$$

Our analysis is in terms of this quantity, and Lemma 8 explains the connection between  $\kappa(\Phi_t, s)$  and  $C_{\min}$  as defined in (2), particularly that  $\kappa(\Phi_t, 2)$  is also positive.

**Theorem 3** (Anytime Lasso Confidence Sequences). *Consider the data model  $y_t = \boldsymbol{\theta}^\top \boldsymbol{\phi}(\mathbf{x}_t) + \varepsilon_t$  for all  $t \geq 1$ , where  $\varepsilon_t$  is i.i.d. zero-mean sub-Gaussian noise, and  $\mathbf{x}_t$  is  $\mathcal{F}_t$ -measurable, where  $\mathcal{F}_t := (\mathbf{x}_1, \dots, \mathbf{x}_t, \varepsilon_1, \dots, \varepsilon_{t-1})$ . Then the solution of (1) guarantees*

$$\mathbb{P}\left(\forall t \geq 1: \|\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_t\|_2 \leq \frac{4\sqrt{10}\lambda_t}{\kappa^2(\Phi_t, 2)}\right) \geq 1 - \delta$$

if the regularization parameter satisfies

$$\lambda_t \geq \frac{2\sigma}{\sqrt{t}} \left(1 + \frac{12}{\sqrt{2}} (\log(2M/\delta) + (\log \log d)_+) + \frac{5}{\sqrt{2}} \sqrt{d (\log(2M/\delta) + (\log \log d)_+)}\right)^{1/2}.$$

Our confidence bound enjoys the same rate as Lasso with offline i.i.d. data, up to  $\log \log d$  factors. We prove this theorem by constructing a self-normalized martingale sequence based on the  $\ell_2$ -norm of the empirical process error  $(\Phi_t^\top \boldsymbol{\varepsilon}_t)$ . We then apply the “stitched” time-uniform boundary of Howard et al. (2021). Appendix B elaborates on this technique. Previous work on sparse linear bandits also include analysis of Lasso in an online setting, when  $\mathbf{x}_t$  is  $\mathcal{F}_t$  measurable. Cella and Pontil (2021) imitate offline analysis and then apply a union bound, which multiplies the width by  $\log n$  and requires knowledge of the horizon. Bastani and Bayati (2020) also rely on knowledge of  $n$  and employ a scalar-valued Bernstein inequality on  $\|\Phi_t^\top \boldsymbol{\varepsilon}_t\|_\infty$ , which inflates the width of the confidence sets by a factor of  $\sqrt{d}$ . We work on the  $\ell_2$ -norm, and use a curved boundary for the resulting sub-Gamma martingale sequence, which is known to be uniformly tighter than a Bernstein bound, especially for small  $t$  (Howard et al., 2021).

## 7. Conclusion

We proposed ALEXP, an algorithm for simultaneous online model selection and bandit optimization. As a first, our approach leads to anytime valid guarantees for model selection and bandit regret, and does not rely on a priori determined exploration schedule. Further, we showed how the Lasso can be used together with the exponential weights algorithm to construct a low-variance online learner. This new connection between high-dimensional statistics and online learning opens up avenues for future research on online high-dimensional learning.

## References

- Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. *Advances in Neural Information Processing Systems*, 2011.
- Yasin Abbasi-Yadkori, David Pal, and Csaba Szepesvari. Online-to-confidence-set conversions and application to sparse stochastic bandits. In *Artificial Intelligence and Statistics*. PMLR, 2012.
- Jacob Abernethy and Alexander Rakhlin. Beating the adaptive bandit with high probability. In *2009 Information Theory and Applications Workshop*. IEEE, 2009.
- Jacob Abernethy, Elad Hazan, and Alexander Rakhlin. An efficient algorithm for bandit linear optimization. In *Conference on Learning Theory*, 2008.
- Alekh Agarwal, Daniel Hsu, Satyen Kale, John Langford, Lihong Li, and Robert Schapire. Taming the monster: A fast and simple algorithm for contextual bandits. In *International Conference on Machine Learning*. PMLR, 2014.
- Alekh Agarwal, Haipeng Luo, Behnam Neyshabur, and Robert E Schapire. Corraling a band of bandit algorithms. In *Conference on Learning Theory*. PMLR, 2017.
- Peter Auer, Nicolo Cesa-Bianchi, Yoav Freund, and Robert E Schapire. Gambling in a rigged casino: The adversarial multi-armed bandit problem. In *Foundations of Computer Science*. IEEE, 1995.
- Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 2002a.
- Peter Auer, Nicolo Cesa-Bianchi, Yoav Freund, and Robert E Schapire. The nonstochastic multiarmed bandit problem. *SIAM journal on computing*, 2002b.
- Debangshu Banerjee, Avishek Ghosh, Sayak Ray Chowdhury, and Aditya Gopalan. Exploration in linear bandits with rich action sets and its implications for inference. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 2023.
- Peter L Bartlett, Varsha Dani, Thomas P Hayes, Sham M Kakade, Alexander Rakhlin, and Ambuj Tewari. High-probability regret bounds for bandit online linear optimization. In *Conference on Learning Theory*, 2008.
- Hamsa Bastani and Mohsen Bayati. Online decision making with high-dimensional covariates. *Operations Research*, 2020.
- Alina Beygelzimer, John Langford, Lihong Li, Lev Reyzin, and Robert Schapire. Contextual bandit algorithms with supervised learning guarantees. In *International Conference on Artificial Intelligence and Statistics*, 2011.
- Peter J. Bickel, Ya'acov Ritov, and Alexandre B. Tsybakov. Simultaneous analysis of Lasso and Dantzig selector. *The Annals of Statistics*, 2009.

- Sébastien Bubeck, Nicolo Cesa-Bianchi, et al. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends® in Machine Learning*, 2012.
- Sébastien Bubeck, Yin Tat Lee, and Ronen Eldan. Kernel-based methods for bandit convex optimization. In *ACM SIGACT Symposium on Theory of Computing*, pages 72–85, 2017.
- Peter Bühlmann and Sara Van De Geer. *Statistics for high-dimensional data: methods, theory and applications*. Springer Science & Business Media, 2011.
- Alexandra Carpentier and Rémi Munos. Bandit theory meets compressed sensing for high dimensional stochastic linear bandit. In *Artificial Intelligence and Statistics*. PMLR, 2012.
- Leonardo Cella and Massimiliano Pontil. Multi-task and meta-learning with sparse linear bandits. In *Conference on Uncertainty in Artificial Intelligence*, 2021.
- Dylan J Foster, Akshay Krishnamurthy, and Haipeng Luo. Model selection for contextual bandits. *Advances in Neural Information Processing Systems*, 32, 2019.
- Sébastien Gerchinovitz. Sparsity regret bounds for individual sequences in online linear regression. In *Conference on Learning Theory*, 2011.
- Apostolos Giannopoulos. Notes on isotropic convex bodies, October 2003.
- Botao Hao, Tor Lattimore, and Mengdi Wang. High-dimensional sparse linear bandits. *Advances in Neural Information Processing Systems*, 2020.
- Steven R. Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Time-uniform chernoff bounds via nonnegative supermartingales. *Probability Surveys*, 2020.
- Steven R Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Time-uniform, nonparametric, nonasymptotic confidence sequences. *The Annals of Statistics*, 2021.
- Kyoungseok Jang, Chicheng Zhang, and Kwang-Sung Jun. Popart: Efficient sparse regression and experimental design for optimal sparse linear bandits. In *Advances in Neural Information Processing Systems*, 2022.
- Adel Javanmard and Andrea Montanari. Confidence intervals and hypothesis testing for high-dimensional regression. *The Journal of Machine Learning Research*, 2014.
- Parnian Kassaie, Jonas Rothfuss, and Andreas Krause. Meta-learning hypothesis spaces for sequential decision-making. In *International Conference on Machine Learning*. PMLR, 2022.
- Emilie Kaufmann, Olivier Cappé, and Aurélien Garivier. On bayesian upper confidence bounds for bandit problems. In *Artificial intelligence and statistics*. PMLR, 2012.
- Gi-Soo Kim and Myunghee Cho Paik. Doubly-robust lasso bandit. *Advances in Neural Information Processing Systems*, 32, 2019.
- Tomáš Kocák, Gergely Neu, Michal Valko, and Rémi Munos. Efficient learning by implicit exploration in bandit problems with side observations. *Advances in Neural Information Processing Systems*, 2014.



- Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.
- Chung-Wei Lee, Haipeng Luo, Chen-Yu Wei, and Mengxiao Zhang. Bias no more: high-probability data-dependent regret bounds for adversarial bandits and mdps. *Advances in Neural Information Processing Systems*, 2020.
- Wenjie Li, Adarsh Barik, and Jean Honorio. A simple unified framework for high dimensional bandit problems. In *International Conference on Machine Learning*. PMLR, 2022.
- Karim Lounici, Massimiliano Pontil, Sara Van De Geer, and Alexandre B Tsybakov. Oracle inequalities and optimal inference under group sparsity. *The annals of statistics*, 2011.
- Haipeng Luo, Mengxiao Zhang, Peng Zhao, and Zhi-Hua Zhou. Corraling a larger band of bandits: A case study on switching regret for linear bandits. In *Conference on Learning Theory*. PMLR, 2022.
- Odalric-Ambrym Maillard and Rémi Munos. Adaptive bandits: Towards the best history-dependent strategy. In *International Conference on Artificial Intelligence and Statistics*. JMLR Workshop and Conference Proceedings, 2011.
- Mathurin Massias, Alexandre Gramfort, and Joseph Salmon. Celer: a fast solver for the lasso with dual extrapolation. In *International Conference on Machine Learning*, 2018.
- H. Brendan McMahan and Matthew Streeter. Tighter bounds for multi-armed bandits with expert advice. In *Proceedings of the 22nd Annual Conference on Learning Theory*, 2009.
- Gergely Neu. Explore no more: Improved high-probability regret bounds for non-stochastic bandits. *Advances in Neural Information Processing Systems*, 2015.
- Min-hwan Oh, Garud Iyengar, and Assaf Zeevi. Sparsity-agnostic lasso bandit. In *International Conference on Machine Learning*. PMLR, 2021.
- Aldo Pacchiano, My Phan, Yasin Abbasi Yadkori, Anup Rao, Julian Zimmert, Tor Lattimore, and Csaba Szepesvari. Model selection in contextual stochastic bandit problems. *Advances in Neural Information Processing Systems*, 2020.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017.
- Ali Rahimi, Benjamin Recht, et al. Random features for large-scale kernel machines. In *Advances in Neural Information Processing Systems*, 2007.
- Felix Schur, Parnian Kassraie, Jonas Rothfuss, and Andreas Krause. Lifelong bandit optimization: No prior and no regret. In *Conference on Uncertainty in Artificial Intelligence*, 2023.
- William R Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 1933.

Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 1996.

Sara A. van de Geer and Peter Bühlmann. On the conditions used to prove oracle results for the Lasso. *Electronic Journal of Statistics*, 2009.

Roman Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press, 2018.

Martin J Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*. Cambridge university press, 2019.

Julian Zimmert and Tor Lattimore. Return of the bias: Almost minimax optimal high probability bounds for adversarial linear bandits. In *Conference on Learning Theory*. PMLR, 2022.

## Appendix A. Details of the Main Text

### A.1 Extended Literature Review

The sparse linear bandit literature considers linear reward functions of the form  $\boldsymbol{\theta}^\top \mathbf{x}$ , where  $\mathbf{x} \in \mathbb{R}^p$ , however a sub-vector of size  $d$  is sufficient to span the reward function. This can be formulated as model selection among  $M = \binom{p}{d}$  different linear parametrizations, where each  $\phi_j$  is a  $d$ -dimensional feature map. We present the bounds in terms of  $d$  and  $M$  for coherence with the rest of the text, assuming that  $M = \mathcal{O}(p)$ , which is the case when  $d \ll p$ .

Table 1 compares recent work on sparse linear bandits based on a number of important factors. In this table, the ETC algorithms follow the general format of exploring, performing parameter estimation once at  $t = n_0$ , and then repeatedly suggesting the same action which maximizes  $\hat{\boldsymbol{\theta}}_{n_0}^\top \phi(\mathbf{x})$ . Explore-then- $(\epsilon)$ Greedy takes a similar approach, however it does not settle on  $\boldsymbol{\theta}_{n_0}$ , rather it continues to update the parameter estimate and select  $\mathbf{x}_t = \arg \max \hat{\boldsymbol{\theta}}_t^\top \phi(\mathbf{x})$ . The UCB algorithms iteratively update upper confidence bound, and choose actions which maximize them. The regret bounds in Table 1 are simplified to the terms with largest rate of growth, *the reader should check the corresponding papers for rigorous results*. Some of the mentioned bounds depend on problem-dependent parameters (e.g.  $c_K$ ), which may not be treated as absolute constants and have complicated forms. To indicate such parameters we use  $\tau$  in Table 1, following the notation of Hao et al. (2020). Note that  $\tau$  varies across the rows of the table, and is just an indicator for existence of other terms.

Abbasi-Yadkori et al. (2012) use the SEQSEW online regression oracle (Gerchinovitz, 2011) for estimating the parameter vector, together with a UCB policy. The regression oracle is an exponential weights algorithm, which runs on the squared error loss. This subroutine, and thereby the algorithm proposed by Abbasi-Yadkori et al. (2012) are not computationally efficient, and this is believed to be unavoidable. This work considers the data-rich regime and shows  $R(n) = \mathcal{O}(\sqrt{dMn})$ , matching the lower bound of Theorem 24.3 in Lattimore and Szepesvári (2020).

Carpentier and Munos (2012) assume that the action set is a Euclidean ball, and that the noise is directly added to the parameter vector, i.e.  $y_t = \mathbf{x}^\top (\boldsymbol{\theta} + \boldsymbol{\varepsilon}_t)$ . Roughly put, linear bandits with parameter noise are “easier” to solve than stochastic linear bandits with reward noise, since the noise is scaled proportionally to the features  $x_i$  and does less “damage” (Chapter 29.3 Lattimore and Szepesvári, 2020). In this setting, Carpentier and Munos (2012) present a  $\mathcal{O}(d\sqrt{n})$  regret bound.

Recent work considers contextual linear bandits, where at every step  $\mathcal{A}_t$ , a stochastic finite subset of size  $K$  from  $\mathcal{A}$ , is presented to the agent. It is commonly assumed that members of  $\mathcal{A}_t$  are i.i.d., and the sampling distribution is diverse and time-independent. The diversity assumption is often in the form of a restricted eigenvalue condition (Definition 2) on the covariance of the context distribution (e.g. in, Kim and Paik, 2019; Bastani and Bayati, 2020). Li et al. (2022) require a stronger condition which directly assumes that  $\lambda_{\min}(\Phi_t)$  the minimum eigenvalue of the empirical covariance matrix is lower bounded. This is generally not true, but may hold with high probability. Hao et al. (2020) assume that the action set spans  $\mathbb{R}^{dM}$ . We believe that this assumption is the weakest in the literature, and conjecture that it is necessary for model selection. If not met, the agent can not explore in all relevant directions, and may not identify the relevant features. Our diversity assumption is similar

to Hao et al. (2020), adapted to our problem setting. Mainly, we consider reward functions which are linearly parametrizable, i.e.  $\boldsymbol{\theta}^\top \boldsymbol{\phi}(\mathbf{x})$ , as oppose to linear rewards, i.e.  $\boldsymbol{\theta}^\top \mathbf{x}$ .

A key distinguishing factor between ALEXP and existing work on sparse linear bandit is that ALEXP is horizon-independent and does not rely on a forced exploration schedule. As shown on Table 1, majority of prior work relies either on an initial exploration stage, the length of which is determined according to  $n$  (e.g., Carpentier and Munos, 2012; Kim and Paik, 2019; Li et al., 2022; Hao et al., 2020; Jang et al., 2022), or on a hand crafted schedule, which is again designed for a specific horizon (Bastani and Bayati, 2020). Oh et al. (2021), which analyzes  $K$ -armed contextual bandits, does not require explicit exploration, and instead imposes restrictive assumptions on the diversity of context distribution, e.g. relaxed symmetry and balanced covariance. Regardless, the regret bounds hold in expectation, and are not time-uniform.

Table 1: Overview of recent work on high-dimensional Bandits. Parameter  $\tau$  shows existence of other problem-dependent terms which are not constants, and varies across different rows. The regret bounds are simplified and are *not* rigorous.

	$ \mathcal{A}_t $	data-poor	adap. exp.	any-time	action selection policy	MS algo	context or action assumpt.	Regret
Abbasi-Yadkori et al.	$\infty$	✗	✓	✓	UCB	EXP4 on Sqrd error	$\mathcal{A}$ is compact	$\sqrt{dMn}$ , w.h.p.
Foster et al.	$K$	✓	✓	✗	UCB	EXP4 on Sqrd error	$\lambda_{\min}(\Sigma) \geq c_\lambda$	$(Mn)^{3/4}K^{1/4} + \sqrt{KdMn}$ w.h.p
Carpentier and Munos	$\infty$	✓	✗	✗	UCB	Hard Thresh.	$\mathcal{A}$ is a ball param. noise	$d\sqrt{n}$ , w.h.p.
Bastani and Bayati	$K$	✗	✗	✗	Explore then Greedy	Lasso	$\kappa(\Sigma) > c_K$	$\tau K d^2 (\log n + \log M)^2$ w.h.p.
Kim and Paik	$K$	✗	✗	✗	Explore then $\epsilon$ -Greedy	Lasso	$\kappa(\Sigma) > c_K$	$\tau d\sqrt{n} \log(Mn)$ , w.h.p.
Oh et al.	$K$	✓	✓	✗	Greedy	Lasso	$\kappa(\Sigma) > c_\kappa$ + other assumns.	$\tau d\sqrt{n} \log(Mn)$ in expectation
Li et al.	$K$	✓	✗	✗	ETC	Lasso	$\lambda_{\min}(\hat{\Sigma}) > c_\lambda$	$\tau(n^2d)^{1/3} \sqrt{\log Mn}$ in expectation
Hao et al.	$\infty$	✓	✗	✗	ETC	Lasso	$\mathcal{A}$ spans $\mathbb{R}^{dM}$ + is compact	$(ndC_{\min}^{-1})^{2/3} (\log M)^{1/3}$ w.h.p.
Jang et al.	$\infty$	✓	✗	✗	ETC	Hard Thresh.	$\mathcal{A} \subset [-1, 1]^{Md}$ + spans $\mathbb{R}^{Md}$	$(nd)^{2/3} (C_{\min}^{-1} \log M)^{1/3}$ w.h.p.
ALEXP (Ours)	$\infty$	✓	✓	✓	Greedy or UCB	EXP4 on reward est.	$\text{Im}(\boldsymbol{\phi}_j)$ spans $\mathbb{R}^d$ $\mathcal{A}$ is compact	$\sqrt{n} \log M (n^{1/4} + C_{\min}^{-1} \log M)$ w.h.p

## A.2 Proof Sketch of Main Theorem

The regret is caused by two sources: selecting a sub-optimal agent, and an agent selecting a sub-optimal action. Accordingly, for any  $j \in 1, \dots, M$ , we decompose the regret as

$$R(n) = \sum_{t=1}^n r(\mathbf{x}^*) - r(\mathbf{x}_t) = \left( \sum_{t=1}^n r(\mathbf{x}^*) - r_{t,j} \right) + \left( \sum_{t=1}^n r_{t,j} - r(\mathbf{x}_t) \right). \quad (3)$$

The first term shows the cumulative regret of agent  $j$ , when run within ALEXP. The second term evaluates the received reward against the cumulative average reward of a model  $j$ . We bound each term separately.

**Virtual Regret.** The first term  $\tilde{R}_j(n) := \sum_{t=1}^n r(\mathbf{x}^*) - r_{t,j}$  compares the suggestion of agent  $j$ , against the optimal action. We call it the virtual regret since the sequence  $(\mathbf{x}_{t,j})_{t \geq 1}$  of the actions suggested by model  $j$  are not necessarily selected by the meta algorithm, unless  $j_t = j$ . This regret is merely a technical tool, and not actually realized when running ALEXP. The virtual regret of the oracle agent may still be bounded using standard analyses of linear bandit, e.g., [Abbasi-Yadkori et al. \(2011\)](#), however we need to slightly adapt it, to take into account a subtle difference: The confidence sequence of model  $j^*$  is constructed according to the true sequence of actions  $(\mathbf{x}_t)_{t \geq 1}$ , while its virtual regret is calculated based on the sequence  $(\mathbf{x}_{t,j^*})_{t \geq 1}$ , which the model suggests. The two sequences only match at the steps when model  $j^*$  is selected. Adapting the analysis of [Abbasi-Yadkori et al. \(2011\)](#) to this subtlety, we obtain in [Lemma 15](#) that with probability greater than  $1 - \delta$ , simultaneously for all  $n \geq 1$

$$\tilde{R}_{j^*}(n) = \mathcal{O}\left(n^{5/8} C_{\min}^{-1/2} \sqrt{d \log(n) + \log(1/\delta)}\right).$$

**Model Selection Regret.** The second term in [\(3\)](#) is the model selection regret,  $R(n, j) := \sum_{t=1}^n r_{t,j} - r(\mathbf{x}_t)$ , which evaluates the chosen action by ALEXP against the suggestion of the  $j$ -th agent. Our analysis relies on a careful decomposition of  $R(n, j)$ ,

$$R(n, j) = \sum_{t=1}^n \left[ \underbrace{r_{t,j} - \hat{r}_{t,j}}_{\text{(I)}} + \underbrace{\hat{r}_{t,j} - \sum_{i=1}^M q_{t,i} \hat{r}_{t,i}}_{\text{(II)}} + \underbrace{\sum_{i=1}^M q_{t,i} (\hat{r}_{t,i} - r_{t,i})}_{\text{(III)}} + \underbrace{\sum_{i=1}^M q_{t,i} r_{t,i} - r(\mathbf{x}_t)}_{\text{(IV)}} \right]$$

We bound away each term in a modular manner, until we are left with the regret of the standard exponential weights algorithm. The terms (I) and (III) are controlled by the bias of the Lasso estimator, and are  $\mathcal{O}(n^{3/4} C(M, \delta, d))$  ([Lemma 19](#)). The last term (IV) is zero in expectation, and reflects the deviation of  $r(\mathbf{x}_t)$  from its mean. We observe that it is a bounded Martingale difference sequence growing with  $\mathcal{O}(\sqrt{n})$  ([Lemma 18](#)). Term (II) is the regret of our online optimizer, which depends on the variance of the Lasso estimator. We bound this term with  $\mathcal{O}(\sqrt{n} C(M, \delta, d) \log M)$ , by first conducting a standard anytime analysis of exponential weights ([Lemma 17](#)), and then incorporating the anytime Lasso variance bound ([Lemma 20](#)). We highlight that neither of the above steps require assumptions about the base agents. Combining these steps, [Lemma 16](#) establishes the formal bound on the model selection regret.

### A.3 Extended Discussion of Results

In light of [Theorem 1](#) and [Theorem 3](#), we discuss some properties of ALEXP.

**Sparse EXP4.** Our approach presents a new connection between online learning and high-dimensional statistics. The rule for updating the probabilities in ALEXP is inspired by the exponential weighting for Exploration and Exploitation with Experts (EXP4) algorithm, which was proposed by [Auer et al. \(2002b\)](#) and has been extensively studied in the adversarial bandit and learning with expert advice literature (e.g., [McMahan and Streeter, 2009](#); [Beygelzimer et al., 2011](#)). EXP4 classically uses importance-weighted (IW) or ordinary least squares

(LS) estimators to estimate the reward  $r_{t,j}$ , both of which are unbiased but high-variance choices (Bubeck et al., 2012). In particular, in our linearly parametrized setting, the variance of IW and LS scales with  $Md$ , which will lead to a  $\text{poly}(M)$  regret. However, it is known that introducing bias can be useful if it reduces the variance (Zimmert and Lattimore, 2022). For instance, EXP3-IX (Kocák et al., 2014) and EXP4-IX (Neu, 2015) construct a biased IW estimator. Equivalently, others have crafted regularizers for the reward of the online optimizer, seeking to improve the bias-variance balance (e.g., Abernethy et al., 2008; Bartlett et al., 2008; Abernethy and Rakhlin, 2009; Bubeck et al., 2017; Lee et al., 2020; Zimmert and Lattimore, 2022). A key technical observation in this work is that our online Lasso estimator leads EXP4 to achieve sublinear regret which depends logarithmically on  $M$ . This is due to the fact that while the estimator itself is  $Md$ -dimensional, its bias squared and variance scale with  $\sqrt{d \log M}$ . To the best of our knowledge, this work is first to instantiate the EXP4 algorithm with a sparse low-variance estimator.

**Adaptive and Anytime.** To estimate the reward, prior work on sparse bandits commonly emulate the Lasso analysis on offline data or on a martingale sequence with a known length (Hao et al., 2020; Bastani and Bayati, 2020). These works require a long enough sequence of exploratory samples, and knowledge of the horizon to plan this sequence. ALEXP removes both of these constraints, and presents a fully adaptive algorithm. Crucially, we employ the elegant martingale bounds of Howard et al. (2020) to present the first time-uniform analysis of the Lasso with history-dependent data (Theorem 3). This way we can use all the data points and explore only adaptively with a probability which vanishes at a  $\mathcal{O}(t^{-1/4})$  rate. Our anytime confidence bound for Lasso, together with the horizon-independent analysis of the exponential weights algorithm, also allows ALEXP to be stopped at any time with valid guarantees.

**Rate Optimality.** For  $M \gg n$ , we obtain a  $\mathcal{O}(\sqrt{n \log^3 M})$  regret, which matches the rate conjectured by Agarwal et al. (2017). However, if  $M$  is comparable to  $n$  or smaller, our regret scales with  $\mathcal{O}(n^{3/4} \sqrt{\log M})$ , and while it is still sublinear and scales logarithmically with  $M$ , the dependency on  $n$  is sub-optimal. This may be due to the conservative nature of our model selection analysis, during which we do not make assumptions about the dynamics of the base agents. Therefore, to ensure sufficiently diverse data for successful model selection, we need to occasionally choose exploratory actions with a vanishing probability of  $\gamma_t$ . We conjecture that this is avoidable, if we make more assumptions about the agents, e.g., that a sufficient number of agents can achieve sublinear regret if executed in isolation. Banerjee et al. (2023) show that the data collected by sublinear algorithms organically satisfies a minimum eigenvalue lowerbound, which may also be sufficient for model selection. We leave this as an open question for future work.

## Appendix B. Time Uniform Lasso Analysis

We start by showing that the sum of squared sub-gaussian variables is a sub-Gamma process (c.f. Definition 22).

**Lemma 4** (Empirical Process is sub-Gamma). *For  $t \geq 1$ , suppose  $\xi_t$  are a sequence conditionally standard sub-Gaussians adapted to the filtration  $\mathcal{F}_t = \sigma(\xi_1, \dots, \xi_t)$ , Let  $v_t \in \mathbb{R}$ , and  $Z_t := \xi_t^2 - 1$ . Define the processes  $S_t := \sum_{i=1}^t Z_i v_i$  and  $V_t := 4 \sum_{i=1}^t v_i^2$ . Then  $(S_t)_{t=0}^\infty$  is sub-Gamma with variance process  $(V_t)_{t=0}^\infty$  and scale parameter  $c = 4 \max_{i \geq 1} v_i$ .*

*Proof of Lemma 4.* By definition (c.f. Definition 1, Howard et al., 2021),  $S_t$  is sub-Gamma if for each  $\lambda \in [0, 1/c)$ , there exists a supermartingale  $(M_t(\lambda))_{t=0}^\infty$  w.r.t.  $\mathcal{F}_t$ , such that  $\mathbb{E} M_0 = 1$  and for all  $t \geq 1$ :

$$\exp \left\{ \lambda S_t - \frac{\lambda^2}{2(1-c\lambda)} V_t \right\} \leq M_t(\lambda) \quad a.s.$$

We show the above holds in equality by proving that the left hand side itself, is a supermartingale w.r.t.  $\mathcal{F}_t$ . We define,  $M_t(\lambda) := \exp\{\lambda S_t - \lambda^2 V_t / 2(1-c\lambda)\}$ , therefore,

$$\begin{aligned} \mathbb{E}[M_t | \mathcal{F}_{t-1}] &\leq \mathbb{E} \left[ \exp \left( \lambda S_{t-1} - \frac{\lambda^2}{2(1-c\lambda)} V_{t-1} + \lambda Z_t v_t - \frac{2\lambda^2 v_t^2}{(1-c\lambda)} \right) | \mathcal{F}_{t-1} \right] \\ &= \mathbb{E}[M_{t-1} | \mathcal{F}_{t-1}] \mathbb{E} [\exp(\lambda Z_t v_t) | \mathcal{F}_{t-1}] \exp \left( -\frac{2\lambda^2 v_t^2}{1-c\lambda} \right) \\ &= M_{t-1} \mathbb{E} [\exp(\lambda Z_t v_t) | \mathcal{F}_{t-1}] \exp \left( -\frac{2\lambda^2 v_t^2}{1-c\lambda} \right). \end{aligned}$$

Note that  $Z_t$  is  $\mathcal{F}_{t-1}$ -measurable, conditionally centered and conditionally sub-exponential with parameters  $(\nu, \alpha) = (2, 4)$  (c.f. Vershynin (2018, Lemma 2.7.6) and Wainwright (2019, Example 2.8)). Therefore, for  $\lambda < 1/c$ ,

$$\mathbb{E} [\exp(\lambda v_t Z_t) | \mathcal{F}_{t-1}] \leq \exp(2\lambda^2 v_t^2) \leq \exp \left( \frac{2\lambda^2 v_t^2}{1-c\lambda} \right),$$

where the last inequality holds due to the fact that  $0 \leq 1-c\lambda < 1$ . Therefore,

$$\mathbb{E}[M_t | \mathcal{F}_{t-1}] \leq M_{t-1} \exp \left( \frac{2\lambda^2 v_t^2}{1-c\lambda} \right) \exp \left( -\frac{2\lambda^2 v_t^2}{1-c\lambda} \right) = M_{t-1}.$$

for  $\lambda \in [0, 1/c)$ , concluding the proof.  $\square$

We now construct a self-normalizing martingale sequence based on  $\ell_2$ -norm of the empirical process error term, and recognize that it is a sub-gamma process. We then employ our curved Bernstein bound Lemma 25 to control the boundary. This step will allow us to “ignore” the empirical process error term later in the lasso analysis.

**Lemma 5** (Empirical Process is dominated by regularization.). *Let*

$$A_j = \left\{ \forall t \geq 1 : \left\| (\Phi_t^\top \varepsilon_t)_j \right\|_2 / t \leq \lambda_t / 2 \right\}.$$

*Then, for any  $0 \leq \delta < 1$ , the event  $A = \bigcap_{j=1}^M A_j$  happens with probability  $1 - \delta$ , if for all  $t \geq 1$ ,*

$$\lambda_t \geq \frac{2\sigma}{\sqrt{t}} \sqrt{1 + \frac{5}{\sqrt{2}} \sqrt{d(\log(2M/\delta) + (\log \log d)_+)}} + \frac{12}{\sqrt{2}} (\log(2M/\delta) + (\log \log d)_+).$$

*Proof of Lemma 5.* This proof includes a treatment of the empirical process similar to Lemma B.1 in Kassaie et al. (2022), but adapts it to our time-uniform setting. Since  $\varepsilon_i$  are zero-mean sub-gaussian variables, as driven in Lemma 3.1 (Lounici et al., 2011), it holds that

$$A_j^c = \left\{ \exists t : \frac{1}{t^2} \varepsilon_t^\top \Phi_{t,j} \Phi_{t,j}^\top \varepsilon_t \geq \frac{\lambda^2}{4} \right\} = \left\{ \exists t : \frac{\sum_{i=1}^t v_i (\xi_i^2 - 1)}{\sqrt{2} \|\mathbf{v}_t\|} \geq \alpha_{t,j} \right\}$$

where  $\xi_i$  are sub-gaussian variables with variance proxy 1, scalar  $v_i$  denotes the  $i$ -th eigenvalue of  $\Phi_{t,j}\Phi_{t,j}^\top/t$  with the concatenated vector  $\mathbf{v}_t = (v_1, \dots, v_t)$ , and

$$\alpha_{t,j} = \frac{t^2\lambda^2/(4\sigma^2) - \text{tr}(\Phi_{t,j}^\top\Phi_{t,j})}{\sqrt{2}\left\|\Phi_{t,j}^\top\Phi_{t,j}\right\|_{\text{Fr}}}.$$

We can apply Lemma 25 to control the probability of event  $A_j^c$  by tuning  $\lambda$ . Mainly, for  $A_j^c$  to happen with probability less than  $\delta/M$ , Lemma 25 states that the following must hold for all  $t$ ,

$$\sqrt{2}\|\mathbf{v}_t\|_2\alpha_{t,j} \geq \frac{5}{2}\sqrt{\max\left\{4\|\mathbf{v}_t\|_2^2, 1\right\}\omega_{\delta/M}(\|\mathbf{v}_t\|_2) + 12\omega_{\delta/M}(\|\mathbf{v}_t\|_2)\max_{t \geq 1} v_t} \quad (4)$$

Recall that w.l.o.g. feature maps are bounded everywhere  $\|\phi_j(\cdot)\|_2 \leq 1$ , and  $\text{rank}(\Phi_j) \leq d$  which allows for the following matrix inequalities,

$$\begin{aligned} \text{tr}(\Phi_{t,j}^\top\Phi_{t,j}) &= \text{tr}(\Phi_{t,j}\Phi_{t,j}^\top) = \sum_{i=1}^t \phi_j^\top(x_i)\phi_j(x_i) \leq t \\ \left\|\Phi_{t,j}\Phi_{t,j}^\top\right\| &\leq \text{tr}(\Phi_{t,j}\Phi_{t,j}^\top) \leq t \\ \left\|\Phi_{t,j}^\top\Phi_{t,j}\right\| &\leq \left\|\Phi_{t,j}\Phi_{t,j}^\top\right\|_{\text{Fr}} \leq \sqrt{d}\left\|\Phi_{t,j}\Phi_{t,j}^\top\right\| \leq t\sqrt{d} \end{aligned}$$

Therefore,

$$\|\mathbf{v}_t\| = \left\|\Phi_{t,j}\Phi_{t,j}^\top\right\|_{\text{Fr}}/t \leq \sqrt{d}, \quad \max_{t \geq 1} v_t = \max_{t \geq 1} \left\|\Phi_{t,j}\Phi_{t,j}^\top\right\|/t \leq 1.$$

For Equation (4) to hold, it suffices that for all  $t \geq 1$ ,

$$\lambda \geq \frac{2\sigma}{\sqrt{t}}\sqrt{1 + \frac{5}{2\sqrt{2}}\sqrt{4d(\log(2M/\delta) + (\log \log d)_+) + \frac{12}{\sqrt{2}}(\log(2M/\delta) + (\log \log d)_+)}}.$$

Therefore, if  $\lambda_t$  are chosen to satisfy the above inequality, each  $A_j^c$  happens with probability less than  $\delta/M$ . Then by applying union bound,  $\cup_{j=1}^M A_j^c$  happens with probability less than  $\delta$ .  $\square$

**Proof of Theorem 3.** The theorem statement requires that the regularization parameter  $\lambda_t$  is chosen such that condition of Lemma 5 is met, and therefore event  $A$  happens with probability  $1 - \delta$ . Throughout this proof, which adapts the analysis of Theorem 3.1. Lounici et al. (2011) to the time-uniform setting, we condition on  $A$  happening, and later incorporate the probability.

**Step 1.** Let  $\hat{\boldsymbol{\theta}}_t$  be a minimizer of  $\mathcal{L}$  and  $\boldsymbol{\theta}$  be the true coefficients vector, then  $\mathcal{L}(\hat{\boldsymbol{\theta}}_t; H_t, \lambda_t) \leq \mathcal{L}(\boldsymbol{\theta}; H_t, \lambda_t)$ . Writing out the loss and re-ordering the inequality we obtain,

$$\frac{1}{t}\left\|\Phi_t(\hat{\boldsymbol{\theta}}_t - \boldsymbol{\theta})\right\|_2^2 \leq \frac{2}{t}\boldsymbol{\varepsilon}_t^T\Phi_t(\hat{\boldsymbol{\theta}}_t - \boldsymbol{\theta}) + 2\lambda_t \sum_{j=1}^M \left(\|\boldsymbol{\theta}_j\|_2 - \left\|\hat{\boldsymbol{\theta}}_{t,j}\right\|_2\right).$$

which is often referred to as the Basic inequality (Bühlmann and Van De Geer, 2011). By Cauchy-Schwarz and conditioned on event  $A$ ,

$$\boldsymbol{\varepsilon}_t^T\Phi_t(\hat{\boldsymbol{\theta}}_t - \boldsymbol{\theta}) \leq \sum_{j=1}^M \left\|(\Phi_t^T\boldsymbol{\varepsilon}_t)_j\right\|_2 \left\|\hat{\boldsymbol{\theta}}_{t,j} - \boldsymbol{\theta}_j\right\|_2 \leq \frac{t\lambda}{2} \sum_{j=1}^M \left\|\hat{\boldsymbol{\theta}}_{t,j} - \boldsymbol{\theta}_j\right\|_2$$



then adding  $\lambda_t \sum_{j=1}^M \|\hat{\boldsymbol{\theta}}_{t,j} - \boldsymbol{\theta}_j\|_2$  to both sides, applying the triangle inequality, and recalling from Section 3 that  $\boldsymbol{\theta}_j = 0$  for  $j \neq j^*$  gives

$$\begin{aligned} \frac{1}{t} \|\Phi_t(\hat{\boldsymbol{\theta}}_t - \boldsymbol{\theta})\|_2^2 + \lambda_t \sum_{j=1}^M \|\hat{\boldsymbol{\theta}}_{t,j} - \boldsymbol{\theta}_j\|_2 &\leq 2\lambda_t \sum_{j=1}^M \|\hat{\boldsymbol{\theta}}_{t,j} - \boldsymbol{\theta}_j\|_2 + 2\lambda_t \sum_{j=1}^M \left( \|\boldsymbol{\theta}_j\|_2 - \|\hat{\boldsymbol{\theta}}_{t,j}\|_2 \right) \\ &\leq 4\lambda_t \|\hat{\boldsymbol{\theta}}_{t,j^*} - \boldsymbol{\theta}_{j^*}\|_2. \end{aligned}$$

Since each term on the left hand side is positive, then each is also individually smaller than the right hand side, and we obtain,

$$\frac{1}{t} \|\Phi_t(\hat{\boldsymbol{\theta}}_t - \boldsymbol{\theta})\|_2^2 \leq 4\lambda_t \|\hat{\boldsymbol{\theta}}_{t,j^*} - \boldsymbol{\theta}_{j^*}\|_2 \quad (5)$$

$$\sum_{\substack{j=1 \\ j \neq j^*}}^M \|\hat{\boldsymbol{\theta}}_{t,j} - \boldsymbol{\theta}_j\|_2 \leq 3 \|\hat{\boldsymbol{\theta}}_{t,j^*} - \boldsymbol{\theta}_{j^*}\|_2 \quad (6)$$

**Step 2.** Consider a sequence  $(c_1, \dots, c_k, \dots)$ , where  $c_1 \geq \dots \geq c_k \dots$ , then

$$c_k \leq \frac{1}{k} (kc_k + \sum_{i>k} c_i) \leq \sum_{i \geq 1} \frac{c_i}{k}. \quad (7)$$

Define  $J_1 = \{j^*\}$  and  $J_2 = \{j^*, j'\}$  where

$$j' = \arg \max_{\substack{j \in [M] \\ j \neq j^*}} \|\hat{\boldsymbol{\theta}}_{t,j} - \boldsymbol{\theta}_j\|_2.$$

For any  $J \subset [M]$  the complementing set is denoted as  $J^c = [M] \setminus J$ . For simplicity let  $c_j = \|\hat{\boldsymbol{\theta}}_{t,j} - \boldsymbol{\theta}_j\|_2$ , and let  $\pi(k)$  denote the index of the  $k$ -th largest element of  $\{c_j : j \in J_1^c\}$ . By definition of  $J_2^c$  we have,

$$\begin{aligned} \sum_{j \in J_2^c} \|\hat{\boldsymbol{\theta}}_{t,j} - \boldsymbol{\theta}_j\|_2^2 &= \sum_{\substack{k>1 \\ \pi(k) \in J_1^c}} c_k^2 \stackrel{(7)}{\leq} \sum_{\substack{k>1 \\ \pi(k) \in J_1^c}} \frac{(\sum_{i \in J_1^c} c_i)^2}{k^2} \\ &\leq \left( \sum_{i \in J_1^c} c_i \right)^2 \sum_{\substack{k>1 \\ \pi(k) \in J_1^c}} \frac{1}{k^2} \stackrel{(6)}{\leq} 9c_{j^*}^2 \\ &\leq 9(c_{j^*}^2 + c_{j'}^2) = 9 \sum_{j \in J_2} \|\hat{\boldsymbol{\theta}}_{t,j} - \boldsymbol{\theta}_j\|_2^2, \end{aligned}$$

which, in turn, gives

$$\|\hat{\boldsymbol{\theta}}_t - \boldsymbol{\theta}\|_2 = \sqrt{\sum_{j=1}^M \|\hat{\boldsymbol{\theta}}_{t,j} - \boldsymbol{\theta}_j\|_2^2} \leq \sqrt{10 \sum_{j \in J_2} \|\hat{\boldsymbol{\theta}}_{t,j} - \boldsymbol{\theta}_j\|_2^2}. \quad (8)$$

**Step 3.** On the other hand, due to (6), and by definition of  $J_2$  it also holds that

$$\sum_{j \in J_2^c} \left\| \hat{\boldsymbol{\theta}}_{t,j} - \boldsymbol{\theta}_j \right\|_2 \leq 3 \sum_{j \in J_2} \left\| \hat{\boldsymbol{\theta}}_{t,j} - \boldsymbol{\theta}_j \right\|_2.$$

From the theorem assumptions and Definition 2, we know that there exists  $0 < \kappa(\Phi_t, 2)$ , therefore by Definition 2, the feature matrix  $\Phi_t$  satisfies,

$$\begin{aligned} \sum_{j \in J_2} \left\| \hat{\boldsymbol{\theta}}_{t,j} - \boldsymbol{\theta}_j \right\|_2^2 &\leq \frac{1}{t\kappa^2(\Phi_t, 2)} \left\| \Phi_t(\hat{\boldsymbol{\theta}}_t - \boldsymbol{\theta}) \right\|_2^2 \\ &\stackrel{(5)}{\leq} \frac{1}{\kappa^2(\Phi_t, 2)} 4\lambda_t \left\| \hat{\boldsymbol{\theta}}_{t,j^*} - \boldsymbol{\theta}_{j^*} \right\|_2 \leq \frac{1}{\kappa^2(\Phi_t, 2)} 4\lambda_t \sqrt{\sum_{j \in J_2} \left\| \hat{\boldsymbol{\theta}}_{t,j} - \boldsymbol{\theta}_j \right\|_2^2}. \end{aligned}$$

From here, by applying (8) we get,

$$\left\| \hat{\boldsymbol{\theta}}_t - \boldsymbol{\theta} \right\|_2 \leq \frac{4\sqrt{10}\lambda_t}{\kappa^2(\Phi_t, 2)}.$$

If  $\lambda_t$  are chosen according to Lemma 5, event  $A$  and, in turn, the inequality above hold with probability greater than  $1 - \delta$ .  $\square$

## Appendix C. Results on Exploration

In this section we present lower-bounds on the eigenvalues of the covariance matrix  $\Phi_t \Phi_t^\top$ , as it is later used in our regret analysis. In particular, we show that the feature matrix  $\Phi_t$  satisfies the restricted eigenvalue condition (Definition 2) required for valid Lasso confidence set (Theorem 3), and calculate a lower bound on  $\kappa(\Phi_t, 2)$ . The lower bound is later used by Lemma 19 and Lemma 20 to develop the model selection regret. We show this bound in three steps.

Equivalent to Definition 2, we write  $\kappa(\Phi_t, s)$  as  $\kappa(\Phi_t, s) = \inf_{b \in \Xi_s} \|\Phi_t \mathbf{b}\|_2 / \sqrt{t}$  where

$$\Xi_s := \left\{ \mathbf{b} \in \mathbb{R}^d \setminus \{0\} \mid \sum_{j \notin J} \|\mathbf{b}_j\|_2 \leq 3 \sum_{j \in J} \|\mathbf{b}_j\|_2, \sqrt{\sum_{j \in J} \|\mathbf{b}_j\|_2^2} \leq 1 \text{ s.t. } J \subset \{1, \dots, M\}, |J| \leq s \right\}. \quad (9)$$

Using the set  $\Xi_s$  we define

$$\tilde{\kappa}(\Sigma, s) := \min_{b \in \Xi_s} b^\top \Sigma b. \quad (10)$$

and use in this section for simplicity.

**Step I.** Consider the exploratory steps at which  $\alpha_t = 1$ . Let  $\Phi_{\pi,t}$  be a sub-matrix of  $\Phi_t$  where only rows from exploratory steps are included. Note that  $\Phi_{\pi,t} \in \mathbb{R}^{t' \times dM}$  is a random matrix, where  $t'$  the number of rows are also random. We show that  $\kappa^2(\Phi_t, s)$  is lower bounded by  $\kappa^2(\Phi_{\pi,t}, s)$ .

**Lemma 6.** *Suppose  $\Phi_{\pi,t}$  has  $t'$  rows. Then,*

$$\kappa^2(\Phi_t, s) \geq \frac{t'}{t} \kappa^2(\Phi_{\pi,t}, s)$$

While the number of rows of  $\Phi_{\pi,t}$  is a random variable, we condition on the event that  $\Phi_{\pi,t}$  has  $t'$  rows, and investigate the distribution of its restricted eigenvalues.

**Step II.** The restricted eigenvalues of the *exploratory* sub-matrix are well bounded away from zero.

**Lemma 7.** *Let  $\pi$  be the solution to (2), and  $s \in \mathbb{N}$ . Suppose  $\Phi_{\pi,t}$  has  $t'$  rows. Then for all  $\delta > 0$ ,*

$$\mathbb{P}\left(\forall t' : \kappa^2(\Phi_{\pi,t}, s) \geq \tilde{\kappa}(\Sigma, s) - \frac{80s}{\sqrt{t'}} \sqrt{(\log \log 4t')_+ + \log(2Md/\delta)}\right) \geq 1 - \delta$$

where  $\tilde{\kappa}(\Sigma, s)$  is defined in (10).

**Step III.** Remains to combine the two lemmas above and incorporate a high probability bound on  $t'$ , showing that it is close to  $\sum_{s=1}^t \gamma_s$ .

**Lemma 8.** *Let  $\pi$  be the solution to (2). There exists absolute constants  $C_1, C_2$  which satisfy,*

$$\mathbb{P}\left(\forall t \geq 1 : \kappa^2(\Phi_t, 2) \geq C_1 \tilde{\kappa}(\Sigma, 2)t^{-1/4} - C_2 t^{-5/8} \sqrt{\log(Md/\delta) + (\log \log t)_+}\right) \geq 1 - \delta$$

if  $\gamma_t = \mathcal{O}(t^{-1/4})$ , as prescribed by Theorem 1. It further holds that,

$$\mathbb{P}\left(\forall t \geq 1 : \kappa^2(\Phi_t, 2) \geq C_1 C_{\min} t^{-1/4} - C_2 t^{-5/8} \sqrt{\log(Md/\delta) + (\log \log t)_+}\right) \geq 1 - \delta$$

The regret analysis of Hao et al. (2020) also relies on connecting  $\kappa(\Phi_t, s)$  to  $C_{\min}$ , and for this, they use Theorem 2.4 of Javanmard and Montanari (2014). This theorem states that there exists a problem-dependent constant  $C_1$  for which  $\kappa^2(\Phi_t, s) \geq C_1 C_{\min}$  with high probability, if  $t \geq n_0$  and roughly  $n_0 = \mathcal{O}(\sqrt[3]{n^2 \log M})$ . We highlight that Lemma 8, presents a lower-bound which holds for all  $t \geq 1$ , however this comes at the cost of getting a looser lower bound than the result of Javanmard and Montanari (2014) for the larger time steps  $t$ . In fact, due to the sub-optimal dependency of Lemma 8 on  $t$ , we later obtain sub-optimal dependency on the horizon for the case when  $n \gg M$ . It is unclear to us if this rate can be improved without assuming knowledge of  $n$ , or that  $n \geq n_0$ .

For the last lemma in this section we show that the empirical sub-matrices  $\Phi_{t,j}$  are also bounded away from zero. This will be required later for proving Lemma 15.

**Lemma 9** (Base Model  $\lambda_{\min}$  Bound). *Assume  $\pi$  is the maximizer of Equation (2). Then, with probability greater than  $1 - \delta$ , simultaneously for all  $j = 1, \dots, M$  and  $t \geq 1$ ,*

$$\lambda_{\min}(\Phi_{t,j}^\top \Phi_{t,j}) \geq C_1 C_{\min} t^{3/4} - C_2 t^{3/8} \sqrt{\log(Md/\delta) + (\log \log t)_+}$$

if  $\gamma_t = \mathcal{O}(t^{-1/4})$ , as prescribed by Theorem 1.

## C.1 ALEXP with Uniform Exploration

We present our main regret bound (Theorem 1) in terms of  $C_{\min}$ , which only depends on properties of the feature maps and the action domain. We give a lower-bound on  $C_{\min}$  for a toy scenario which corresponds to the problem of linear feature selection over convex action sets.

**Proposition 10** (Invertible Features). *Suppose  $\phi(\mathbf{x}) := A\mathbf{x} : \mathbb{R}^d \rightarrow \mathbb{R}^d$  is an invertible linear map, and  $\mathcal{X} \in \mathbb{R}^d$  is a convex body. Then,*

$$C_{\min} \geq \frac{\lambda_{\min}(A)}{\lambda_{\max}^2(T)} > 0$$

where  $T$  is the transformation which maps  $\mathcal{X}$  to an isotropic body.

We expect this proposition to be applicable to any action set which span  $\mathbb{R}^d$ , and to non-linear invertible maps. The lower-bound of Proposition 10 is achieved by simply exploring via  $\pi = \text{Unif}(\mathcal{X})$ . Inspired by Schur et al. (2023, Lemma E.13), we show that even for non-convex action domains and orthogonal feature maps, the uniform exploration yields a constant lower-bound on restricted eigenvalues.

**Proposition 11** (Orthonormal Features). *Suppose  $\phi_j : \mathcal{X} \rightarrow \mathbb{R}$  are chosen from an orthogonal basis of  $L^2(\mathcal{X})$ , and satisfy  $\|\phi_i\|_{L^2_{\mu}(\mathcal{X})}/\text{Vol}(\mathcal{X}) \geq 1$ . Then there exist absolute constants  $C_1$  and  $C_2$  for which the exploration distribution  $\pi = \text{Unif}(\mathcal{X})$  satisfies*

$$\mathbb{P}\left(\forall t \geq 1 : \kappa^2(\Phi_t, 2) \geq C_1 t^{-1/4} - C_2 t^{-5/8} \sqrt{\log(Md/\delta) + (\log \log t)_+}\right) \geq 1 - \delta.$$

The  $d = 1$  condition is met without loss of generality, by splitting the higher dimensional feature maps and introducing more base features, which will increase  $M$ . Moreover, the orthonormality condition is met by orthogonalizing and re-scaling the feature maps. Basis functions such as Legendre polynomials and Fourier features (Rahimi et al., 2007) satisfy these conditions.

By invoking Proposition 11, instead of Lemma 8 in the proof of Theorem 1, we obtain the regret of ALEXP with uniform exploration.

**Corollary 12** (ALEXP with Uniform Exploration). *Let  $\delta \in (0, 1]$ . Suppose  $\phi_j : \mathcal{X} \rightarrow \mathbb{R}$  are chosen from an orthogonal basis of  $L^2(\mathcal{X})$ , and satisfy  $\|\phi_i\|_{L^2_{\mu}(\mathcal{X})}/\text{Vol}(\mathcal{X}) \geq 1$ . Assume the oracle agent employs a UCB or a Greedy policy, as laid out in Section 5. Choose  $\eta_t = \mathcal{O}(1/\sqrt{t}C(M, \delta, d))$  and  $\gamma_t = \mathcal{O}(t^{-1/4})$  and  $\lambda_t = \mathcal{O}(C(M, \delta, d)/\sqrt{t})$ , then ALEXP with uniform exploration  $\pi = \text{Unif}(\mathcal{X})$  attains the regret*

$$R(n) = \mathcal{O}\left(Bn^{3/4} + \sqrt{n}C(M, \delta, d) \log M + B^2\sqrt{n} + B\sqrt{n((\log \log nB^2)_+ + \log(1/\delta))} + (n^{3/4} + \log n)C(M, \delta, d) + n^{5/8}\sqrt{d \log n + \log(1/\delta) + B^2}\right)$$

with probability greater than  $1 - \delta$ , simultaneously for all  $n \geq 1$ . Here,

$$C(M, \delta, d) = \mathcal{O}\left(\sqrt{1 + \sqrt{d(\log(M/\delta) + (\log \log d)_+)}} + (\log(M/\delta) + (\log \log d)_+)\right).$$

## C.2 Proof of Results on Exploration

As an intermediate step, we consider the restricted eigenvalue property of the empirical covariance matrix. Given  $t'$  samples, the empirical estimate of  $\Sigma$  is

$$\hat{\Sigma}_{t'} := \frac{1}{t'} \sum_{s=1}^{t'} \phi(\mathbf{x}_s) \phi^\top(\mathbf{x}_s) \quad (11)$$

where  $\mathbf{x}_s$  are sampled according to  $\pi$ . We show that every entry of  $\hat{\Sigma}_{t'}$  is close to the corresponding entry in  $\Sigma$ , and later use it in the proofs of eigenvalue lemmas.

**Lemma 13** (Anytime Bound for The Entries of Empirical Covariance Matrix). *Let  $\hat{\Sigma}_{t'}$  be the empirical covariance matrix corresponding to  $\Sigma(\pi, \phi)$  given  $t'$  samples. Then,*

$$\mathbb{P}\left(\exists t' : d_\infty(\Sigma, \hat{\Sigma}_{t'}) \geq \frac{5}{\sqrt{t'}} \sqrt{((\log \log 4t')_+ + \log(2Md/\delta))}\right) \leq \delta$$

where  $d_\infty(A, B) := \max_{i,j} |A_{i,j} - B_{i,j}|$ .

We are now ready to present the proofs to Lemmas in Appendix C

*Proof of Lemma 6.* Let  $\Psi^{(t)} = \phi(\mathbf{x}_t)\phi^\top(\mathbf{x}_t) \in \mathbb{R}^{dM \times dM}$  for all  $t = 1, \dots, n$ . Note that  $\Psi^{(t)}$  is positive semi-definite by construction. We have,

$$\|\Phi_t b\|_2^2 = \sum_{s=1}^t b^\top \Psi^{(s)} b = \sum_{s \in T_\pi} b^\top \Psi^{(s)} b + \sum_{s \notin T_\pi} b^\top \Psi^{(s)} b$$

where the set  $T_\pi$  contains the indices of the steps at which the action is selected according to  $\pi$ . We use  $\alpha_t$  to indicate whether  $s \in T_\pi$  or not. Therefore, ,

$$\begin{aligned} \kappa^2(\Phi_t, s) &= \frac{1}{t} \min_{b \in \Xi_s} \|\Phi_t b\|_2^2 \\ &= \min_{b \in \Xi_s} \frac{1}{t} \sum_{s \in T_\pi} b^\top \Psi^{(s)} b + \frac{1}{t} \sum_{s \notin T_\pi} b^\top \Psi^{(s)} b \\ &\geq \min_{b \in \Xi_s} \frac{1}{t} \sum_{s \in T_\pi} b^\top \Psi^{(s)} b \end{aligned}$$

where the last inequality holds due to  $\Psi^{(s)}$  being PSD. Then we have,

$$\kappa^2(\Phi_t, s) \geq \min_{b \in \Xi_s} b^\top \left( \frac{1}{t} \sum_{s \in T_\pi} \Psi^{(s)} \right) b = \frac{|T_\pi|}{t} \kappa^2(\Phi_{\pi, t}, s)$$

□

*Proof of Lemma 7.* We observe that the restricted minimum eigenvalue of  $\hat{\Sigma}_{t'}$  is the  $\kappa^2$  constant of  $\Phi_{\pi, t'}$ :

$$\kappa^2(\Phi_{\pi, t}, s) = \min_{b \in \Xi_s} b^\top \hat{\Sigma}_{t'} b = \tilde{\kappa}(\hat{\Sigma}_{t'}, s).$$

Inspired by Lemma 10.1 in [van de Geer and Bühlmann \(2009\)](#), we show that element-wise closeness of matrices  $\Sigma$  and  $\hat{\Sigma}_{t'}$  (c.f. Lemma 13) implies closeness in  $\tilde{\kappa}$ :

$$\left| \kappa^2(\Phi_{\pi, t}, s) - \tilde{\kappa}(\Sigma, s) \right| = \left| \tilde{\kappa}(\hat{\Sigma}_{t'}, s) - \tilde{\kappa}(\Sigma, s) \right| = \left| \tilde{\kappa}(\hat{\Sigma}_{t'} - \Sigma, s) \right| \leq \min_{b \in \Xi_s} d_\infty(\Sigma, \hat{\Sigma}_{t'}) \|b\|_1^2$$

where the last line holds due to Hölder's. Moreover, since  $b \in \Xi_s$ , for any  $J \subset [dM]$  where  $|J| \leq s$  it additionally holds that  $\|b_J\|_2 \leq 1$  and

$$\|b\|_1 \leq (1 + 3) \|b_J\|_1 \leq 4\sqrt{s} \|b_J\|_2 \leq 4\sqrt{s}$$

which gives,

$$\kappa^2(\Phi_{\pi,t}, s) \geq \tilde{\kappa}(\Sigma, s) - 16sd_\infty(\Sigma, \hat{\Sigma}_{t'}).$$

Therefore by Lemma 13,

$$\kappa^2(\Phi_{\pi,t}, s) \geq \tilde{\kappa}(\Sigma, s) - \frac{80s}{\sqrt{t'}} \sqrt{((\log \log 4t')_+ + \log(2Md/\delta))} \quad (12)$$

with probability greater than  $1 - \delta$ , simultaneously for all  $t' \geq 1$ .  $\square$

*Proof of Lemma 8.* In Lemma 6 we showed that

$$\kappa^2(\Phi_t, s) \geq \frac{t'}{t} \kappa^2(\Phi_{\pi,t}, s)$$

where  $t'$  indicates the number of rows in the exploratory sub-matrix of  $\Phi_t$ . Recall that  $t' = \sum_{s=1}^t \alpha_s$  where  $\alpha_s$  are i.i.d Bernoulli random variables with success probability of  $\gamma_s$ . Due to Lemma 24,

$$\mathbb{P} \left( \forall t \geq 1 : \left| t' - \left( \sum_{s=1}^t \gamma_s \right) \right| \leq \Delta_t \right) \geq 1 - \delta/2$$

where

$$\Delta_t := \frac{5}{2} \sqrt{\frac{((\log \log t)_+ + \log(8/\delta))}{t}}.$$

For the rest of the proof, to keep the calculations simple, we ignore the values of the absolute constants, and use the notation  $g(t) = o(f(t))$  if there exists an absolute constant  $c$  where  $g(t) \leq c|f(t)|$  for all  $t \geq 1$ . Let  $\Gamma_t = \sum_{s=1}^t \gamma_s$ , then since  $\gamma_s = \mathcal{O}(s^{-1/4})$  there exists  $C$  such that  $\Gamma_t = Ct^{3/4}$ .

Due to Lemma 7, with probability greater than  $1 - \delta/2$  the following holds for all  $t \geq 1$

$$\begin{aligned} \kappa^2(\Phi_t, 2) &\geq \frac{t'}{t} C_{\min} - \frac{160\sqrt{t'}}{t} \sqrt{(\log \log 4t')_+ + \log(4Md/\delta)} \\ &\geq \frac{\Gamma_t - \Delta_t}{t} C_{\min} - 160 \sqrt{\frac{\Gamma_t + \Delta_t}{t^2}} \sqrt{(\log \log (4\Gamma_t + \Delta_t))_+ + \log(4Md/\delta)} \end{aligned}$$

It is straightforward to observe that there exists absolute constants  $\tilde{C}_i$  which satisfy,

$$\begin{aligned} \kappa^2(\Phi_t, 2) &\geq \tilde{C}_1 t^{-1/4} C_{\min} - \frac{5t^{-3/2} C_{\min}}{2} \sqrt{((\log \log t)_+ + \log(8/\delta))} \\ &\quad - \tilde{C}_2 t^{-5/8} \sqrt{\log(Md/\delta) + (\log \log t)_+} - o\left(t^{-5/8} \sqrt{\log(Md/\delta) + (\log \log t)_+}\right) \\ &\geq \tilde{C}_1 t^{-1/4} C_{\min} - \tilde{C}_3 t^{-5/8} \sqrt{\log(Md/\delta) + (\log \log t)_+} \end{aligned}$$

The last inequality holds since  $t^{-3/2} \sqrt{\log \log t} = o(t^{-5/8} \sqrt{\log \log t})$ . The above chain of inequalities imply that there exist absolute constants  $C_1, C_2$  (w.r.t the parameters of the problem), for which

$$\mathbb{P} \left( \forall t \geq 1 : \kappa^2(\Phi_t, 2) \geq C_1 C_{\min} t^{-1/4} - C_2 t^{-5/8} \sqrt{\log(Md/\delta) + (\log \log t)_+} \right) \geq 1 - \delta.$$

$\square$

*Proof of Lemma 9.* Fix  $j \in \{1, \dots, M\}$ , and construct the set

$$\Xi_{1,j} = \left\{ \mathbf{b} \in \mathbb{R}^d \setminus \{0\} \mid \mathbf{b} = (\mathbf{b}_1, \dots, \mathbf{b}_M), \text{ s.t. } \mathbf{b}_j \in \mathbb{R}^d, \|\mathbf{b}_j\|_2 \leq 1 \text{ and } \forall j' \neq j : \mathbf{b}_{j'} = 0 \right\}.$$

Note that  $\Xi_{1,j} \subset \Xi_s$ . Therefore,

$$\inf_{\mathbf{b} \in \Xi_{1,j}} \|\Phi_t \mathbf{b}\|_2 \geq \inf_{\mathbf{b} \in \Xi_s} \|\Phi_t \mathbf{b}\|_2 = \sqrt{t\kappa}(\Phi_t, s).$$

Moreover, by construction of  $\Xi_{1,j}$  we have for all  $\mathbf{b} \in \Xi_{1,j}$  that  $\Phi_t \mathbf{b} = \Phi_{t,j} \mathbf{b}_j$ , therefore,

$$\inf_{\mathbf{b} \in \Xi_{1,j}} \|\Phi_t \mathbf{b}\|_2^2 = \inf_{\substack{\mathbf{b}_j \in \mathbb{R}^d \\ \|\mathbf{b}_j\|_2^2 \leq 1}} \|\Phi_{t,j} \mathbf{b}_j\|_2^2 = \lambda_{\min}(\Phi_{t,j}^\top \Phi_{t,j}).$$

From the above equations we conclude that  $\lambda_{\min}(\Phi_{t,j}^\top \Phi_{t,j}) \geq t\kappa^2(\Phi_t, s)$ , for all  $j = 1, \dots, M$ . Therefore, using Lemma 8 we obtain that there exists  $C_1, C_2$  such that

$$\mathbb{P} \left( \forall t \geq 1, j = 1, \dots, M : \lambda_{\min}(\Phi_{t,j}^\top \Phi_{t,j}) \geq C_1 C_{\min} t^{3/4} - C_2 t^{3/8} \sqrt{\log(Md/\delta) + (\log \log t)_+} \right) \geq 1 - \delta$$

□

*Proof of Proposition 10.* Since  $\mathcal{X}$  is a convex body, then there exists an invertible map  $T$ , such that  $T(\mathcal{X})$  is an isotropic body (e.g. Proposition 1.1.1., [Giannopoulos, 2003](#)). Then by definition,  $\bar{X} \sim \text{Unif}(T(\mathcal{X}))$  is an isotropic distribution and  $\text{Cov}(\bar{X}) = I_d$  (e.g., c.f. Chapter 3.3.5 [Vershynin, 2018](#)). Since  $\phi$  is linear and invertible, it may be written as  $\phi(\mathbf{x}) = A\mathbf{x}$ , where  $A$  is an invertible matrix. Therefore,

$$\Sigma(\pi, \phi) = \text{Cov}(\phi(X)) = A^\top \text{Cov}(X)A = A^\top \text{Cov}(T^{-1}\bar{X})A = A^\top (T^{-1})^2 A.$$

As for the minimum eigenvalue, suppose  $\mathbf{v} \in \mathbb{R}^d$  and  $\|\mathbf{v}\| = 1$ , then

$$C_{\min} \geq \lambda_{\min}(\Sigma(\pi, \phi)) \geq \mathbf{v}^\top A^\top (T^{-1})^2 A \mathbf{v} \geq \|A\mathbf{v}\|_2 \lambda_{\min}(T^{-2}) = \frac{\|A\mathbf{v}\|_2}{\lambda_{\max}^2(T)} \geq \frac{\lambda_{\min}(A)}{\lambda_{\max}^2(T)}.$$

□

*Proof of Proposition 11.* By the assumption of the proposition, for all  $i \in M$

$$[\Sigma(\pi, \phi)]_{i,i} = \mathbb{E}_{\mathbf{x} \sim \pi} \phi_i^2(\mathbf{x}) = \frac{1}{\text{Vol}(\mathcal{X})} \int_{\mathcal{X}} \phi_i^2(\mathbf{x}) d\mu(\mathbf{x}) \geq 1$$

and for all  $i \neq j$ ,

$$[\Sigma(\pi, \phi)]_{i,j} = \mathbb{E}_{\mathbf{x} \sim \pi} \phi_i(\mathbf{x}) \phi_j(\mathbf{x}) = \frac{1}{\text{Vol}(\mathcal{X})} \int_{\mathcal{X}} \phi_i(\mathbf{x}) \phi_j(\mathbf{x}) d\mu(\mathbf{x}) = 0$$

We use  $\Sigma = \Sigma(\pi, \phi)$ . For any  $\mathbf{b} \in \mathbb{R}^{Md}$  where  $\|\mathbf{b}\| \leq 1$ ,

$$\mathbf{b}^\top \Sigma \mathbf{b} = \sum_{i,j \in [M]} \mathbf{b}_j^\top \Sigma_{i,j} \mathbf{b}_i = \sum_{i \in [M]} \mathbf{b}_i^\top \Sigma_{i,i} \mathbf{b}_i + \sum_{i,j \in [M], i \neq j} \mathbf{b}_j^\top \Sigma_{i,j} \mathbf{b}_i$$

$$= \sum_{i \in [M]} \mathbf{b}_i^\top \Sigma_{i,i} \mathbf{b}_i \geq 1 \sum_{i \in [M]} \|\mathbf{b}_i\|_2^2 \geq 1.$$

Which implies,

$$\tilde{\kappa}(\Sigma, s) = \min_{\mathbf{b} \in \Xi_s} \mathbf{b}^\top \Sigma \mathbf{b} \geq 1.$$

By Lemma 8, there exist absolute constants  $C_1$  and  $C_2$  for which,

$$\mathbb{P} \left( \forall t \geq 1 : \kappa^2(\Phi_t, 2) \geq \tilde{\kappa}(\Sigma, 2) C_1 t^{-1/4} - C_2 t^{-5/8} \sqrt{\log(Md/\delta) + (\log \log t)_+} \right) \geq 1 - \delta.$$

□

*Proof of Lemma 13.* We show the element-wise convergence of  $\Sigma$  to  $\hat{\Sigma}_t$  for the  $(i, j)$  entry where  $i, j = 1, \dots, dM$ . Consider the random sequence  $X_s := \Sigma_{i,j} - \phi_i(\mathbf{x}_s) \phi_j(\mathbf{x}_s)$ . We show that  $X_1, \dots, X_n$  satisfies conditions of Lemma 26. We first observe that

$$\mathbb{E}[X_s | X_{1:s-1}] = \mathbb{E}X_s = \Sigma(i, j) - \mathbb{E}_{\mathbf{x} \sim \pi} \phi_i(\mathbf{x}) \phi_j(\mathbf{x}) = 0$$

since by definition  $\Sigma_{i,j} = \mathbb{E}_{\mathbf{x} \sim \pi} \phi_i(\mathbf{x})^\top \phi_j(\mathbf{x})$ . Moreover, we have normalized features  $\|\phi(\cdot)\| \leq 1$ , therefore, each entry  $\phi_i(\cdot) \phi_j(\cdot)$  is also bounded, yielding  $|X_s| \leq 2$ . Then Lemma 26 implies that for all  $\tilde{\delta} > 0$ ,

$$\mathbb{P} \left( \exists t' : \frac{1}{t'} \sum_{s=1}^{t'} X_s \geq \frac{5}{\sqrt{t'}} \sqrt{\left( (\log \log 4t')_+ + \log(2/\tilde{\delta}) \right)} \right) \leq \tilde{\delta}.$$

Setting  $\tilde{\delta} = \delta/(dM)$  and taking a union bound over all indices concludes the proof. □

## Appendix D. Proof of Regret Bound

**Theorem 14** (Anytime Regret, Formal). *Let  $\delta \in (0, 1]$  and  $\pi$  be the maximizer of (2). Assume the oracle agent employs a UCB or a Greedy policy, as laid out in Section 5. Suppose  $\eta_t = \mathcal{O}(C_{\min}/\sqrt{t}C(M, \delta, d))$  and  $\gamma_t = \mathcal{O}(t^{-1/4})$  and  $\lambda_t = \mathcal{O}(C(M, \delta, d)/\sqrt{t})$ , then exists absolute constants  $C_i$  for which ALEXP attains the regret*

$$\begin{aligned} R(n) &\leq C_1 B n^{3/4} + C_2 \sqrt{n} C_{\min}^{-1} C(M, \delta, d) \log M \\ &\quad + C_3 B^2 C_{\min} \sqrt{n} + C_4 B \sqrt{n \left( (\log \log n B^2)_+ + \log(1/\delta) \right)} \\ &\quad + C_5 \left( 1 + C_{\min}^{-1} n^{-3/8} \sqrt{\log(Md/\delta) + (\log \log n)_+} \right) \\ &\quad \times \left[ B n^{1/4} + \left( n^{3/4} + \frac{\log n}{C_{\min}} \right) C(M, \delta, d) + \frac{n^{5/8}}{\sqrt{C_{\min}}} \sqrt{d \log n + \log(1/\delta) + B^2} \right] \end{aligned}$$

with probability greater than  $1 - \delta$ , simultaneously for all  $n \geq 1$ . Here,

$$C(M, \delta, d) = C_6 \sigma \sqrt{1 + \sqrt{d \left( \log(M/\delta) + (\log \log d)_+ \right)} + \left( \log(M/\delta) + (\log \log d)_+ \right)}.$$

Our main regret bound is an immediate corollary of Lemma 15 and Lemma 16, considering the regret decomposition of (3).



**Lemma 15** (Virtual Regret of the Oracle). *Let  $\delta \in (0, 1]$  and  $\tilde{\lambda} > 0$ . Assume the oracle agent employs a UCB or a Greedy policy, as laid out in Section 5. If  $\gamma_t = \mathcal{O}(t^{1/4})$ , there exists an absolute constant  $C_1$  for which with probability greater than  $1 - \delta$ , simultaneously for all  $n \geq 1$ ,*

$$\begin{aligned} \tilde{R}_{j^*}(n) &= \frac{C_1 n^{5/8}}{\sqrt{C_{\min}}} \left( 1 + n^{-3/8} C_{\min}^{-1} \sqrt{\log(Md/\delta) + (\log \log n)_+} \right) \\ &\quad \times \sqrt{\sigma^2 d \log\left(\frac{n}{\tilde{\lambda}d} + 1\right) + 2\sigma^2 \log(1/\delta) + \tilde{\lambda}B^2} \end{aligned}$$

**Lemma 16** (Any-Time Model-Selection Regret, Formal). *Let  $\delta \in (0, 1]$  and  $\pi$  be the maximizer of (2). Suppose  $\eta_t = \mathcal{O}(C_{\min}/\sqrt{t}C(M, \delta, d))$  and  $\gamma_t = \mathcal{O}(t^{-1/4})$  and  $\lambda_t = \mathcal{O}(C(M, \delta, d)/\sqrt{t})$ , then exists absolute constants  $C_i$  for which ALEXP attains*

$$\begin{aligned} R(n, j) &\leq C_1 B n^{3/4} + C_2 \sqrt{n} C_{\min}^{-1} C(M, \delta, d) \log M \\ &\quad + C_3 B^2 C_{\min} \sqrt{n} + C_4 B \sqrt{n ((\log \log n B^2)_+ + \log(1/\delta))} \\ &\quad + C_5 \left( B n^{1/4} + (n^{3/4} + \frac{\log n}{C_{\min}}) C(M, \delta, d) \right) \left( 1 + C_{\min}^{-1} n^{-3/8} \sqrt{\log(Md/\delta) + (\log \log n)_+} \right) \end{aligned}$$

with probability greater than  $1 - \delta$ , simultaneously for all  $n \geq 1$ . Here,

$$C(M, \delta, d) = C_6 \sigma \sqrt{1 + \sqrt{d (\log(M/\delta) + (\log \log d)_+)} + (\log(M/\delta) + (\log \log d)_+)}.$$

## D.1 Proof of Model Selection Regret

Our technique for bounding the model selection regret relies on a classic horizon-independent analysis of the exponential weights algorithm, presented in Lemma 17.

**Lemma 17** (Anytime Exponential Weights Guarantee). *Assume  $\eta_t \hat{r}_{t,j} \leq 1$  for all  $1 \leq j \leq M$  and  $t \geq 1$ . If the sequence  $\eta_1, \dots, \eta_t, \dots$  is non-increasing, then for all  $n \geq 1$ ,*

$$\sum_{t=1}^n \hat{r}_{t,k} - \sum_{t=1}^n \sum_{j=1}^M q_{t,j} \hat{r}_{t,j} \leq \frac{\log M}{\eta_n} + \sum_{t=1}^n \eta_t \sum_{j=1}^M q_{t,j} \hat{r}_{t,j}^2$$

for any arm  $k \in [M]$ .

*Proof of Lemma 17.* Define  $\hat{R}_{t,i} := \sum_{s=1}^t \hat{r}_{s,i}$  to be the expected cumulative reward of agent  $i$  after  $t$  steps. We rewrite for a fixed  $k$

$$\sum_{t=1}^n \hat{r}_{t,k} - \sum_{t=1}^n \sum_{j=1}^M q_{t,j} \hat{r}_{t,j} = \sum_{t=1}^n \hat{r}_{t,k} - \sum_{t=1}^n \mathbb{E}_{j \sim q_t} [\hat{r}_{t,j}].$$

We focus on a single term in the second sum. For any  $t$ , we have

$$\begin{aligned} -\mathbb{E}_{j \sim q_t} [\hat{r}_{t,j}] &= \log(\exp(-\mathbb{E}_{j \sim q_t} [\frac{\eta_t}{\eta_t} \hat{r}_{t,j}])) = \log(\exp(-\mathbb{E}_{j \sim q_t} [\eta_t \hat{r}_{t,j}])^{1/\eta_t}) \\ &= \frac{1}{\eta_t} \log(\exp(-\mathbb{E}_{j \sim q_t} [\eta_t \hat{r}_{t,j}])) \end{aligned}$$

$$= \frac{1}{\eta_t} \log(\mathbb{E}_{i \sim q_t} \exp(-\mathbb{E}_{j \sim q_t} [\eta_t \hat{r}_{t,j}])) \quad (13)$$

The inner expectation is over  $j$ , while the outer one is over  $i$  and therefore has no effect. Moreover,

$$\begin{aligned} \frac{1}{\eta_t} \log \mathbb{E}_{i \sim q_t} \exp(\eta_t \mathbb{E}_{j \sim q_t} [\hat{r}_{t,j}] - \eta_t \hat{r}_{t,i}) &= \frac{1}{\eta_t} \log (\exp(\eta_t \mathbb{E}_{j \sim q_t} [\hat{r}_{t,j}]) \mathbb{E}_{i \sim q_t} \exp(-\eta_t \hat{r}_{t,i})) \\ &= \frac{1}{\eta_t} \log \mathbb{E}_{i \sim q_t} \exp(\eta_t \mathbb{E}_{j \sim q_t} [\hat{r}_{t,j}]) \\ &\quad + \frac{1}{\eta_t} \log \mathbb{E}_{i \sim q_t} \exp(-\eta_t \hat{r}_{t,i}) \end{aligned} \quad (14)$$

where again, the expectation can be reintroduced to get the last line. Combining (13) and (14),

$$-\mathbb{E}_{j \sim q_t} [\hat{r}_{t,j}] = \frac{1}{\eta_t} \log \mathbb{E}_{i \sim q_t} \exp(-\eta_t \mathbb{E}_{j \sim q_t} [\hat{r}_{t,j}] + \eta_t \hat{r}_{t,i}) - \frac{1}{\eta_t} \log \mathbb{E}_{i \sim q_t} \exp(\eta_t \hat{r}_{t,i}) \quad (15)$$

This transformation is at the core of many exponential weight proofs (Bubeck et al., 2012; Lattimore and Szepesvári, 2020). We first bound the first term in (15):

$$\begin{aligned} \log \mathbb{E}_{i \sim q_t} \exp(-\eta_t \mathbb{E}_{j \sim q_t} [\hat{r}_{t,j}] + \eta_t \hat{r}_{t,i}) &= \log \mathbb{E}_{i \sim q_t} \exp(\eta_t \hat{r}_{t,i}) - \eta_t \mathbb{E}_{j \sim q_t} \hat{r}_{t,j} \\ &\stackrel{(I)}{\leq} \mathbb{E}_{i \sim q_t} \exp(\eta_t \hat{r}_{t,i}) - 1 - \eta_t \mathbb{E}_{j \sim q_t} \hat{r}_{t,j} \\ &= \mathbb{E}_{i \sim q_t} [\exp(\eta_t \hat{r}_{t,i}) - 1 - \eta_t \hat{r}_{t,i}] \\ &\stackrel{(II)}{\leq} \mathbb{E}_{i \sim q_t} [\eta_t^2 \hat{r}_{t,i}^2] \end{aligned} \quad (16)$$

where in (I) we use the fact that  $\log(z) \leq z - 1$  and in (II) we use the fact that for  $x \leq 1$ , we have  $\exp(x) \leq 1 + x + x^2$ , and hence  $\exp(x) - 1 - x \leq x^2$ . For the second term in (15), we will mirror the potential argument in Bubeck et al. (2012), but with a slightly different potential function. We expand the definition of  $q_t$ :

$$\begin{aligned} -\frac{1}{\eta_t} \log \mathbb{E}_{i \sim q_t} \exp(\eta_t \hat{r}_{t,i}) &= -\frac{1}{\eta_t} \log \frac{\sum_{i=1}^M \exp(\eta_t \hat{R}_{t,i})}{\sum_{i=1}^M \exp(\eta_t \hat{R}_{t-1,i})} \\ &= -\frac{1}{\eta_t} \log \frac{1}{M} \sum_{i=1}^M \exp(\eta_t \hat{R}_{t,i}) + \frac{1}{\eta_t} \log \frac{1}{M} \sum_{i=1}^M \exp(\eta_t \hat{R}_{t-1,i}) \\ &= J_t(\eta_t) - J_{t-1}(\eta_t), \end{aligned} \quad (17)$$

where we define  $J_t(\eta) = -\frac{1}{\eta} \log \frac{1}{M} \sum_{i=1}^M \exp(\eta \hat{R}_{t,i})$ . We also define  $F_t(\eta) = \frac{1}{\eta} \log \frac{1}{M} \sum_{i=1}^M \exp(-\eta \hat{R}_{t,i})$ . We observe the relation  $J(\eta) = F(-\eta)$ . From this, it follows that for any  $\eta$ , we have  $J'(\eta) = -F'(-\eta) \leq 0$ , by the argument in Bubeck et al. (2012, Theorem 3.1) that shows  $F'(\eta) \geq 0$  for any  $\eta$ .

**Putting together the pieces** Now, we can bound (15) by inputing (16) and (17):

$$-\mathbb{E}_{j \sim q_t} [\hat{r}_{t,j}] \leq \mathbb{E}_{i \sim q_t} [\eta_t \hat{r}_{t,i}^2] + J_t(\eta_t) - J_{t-1}(\eta_t)$$

With this, we rewrite (D.1) as

$$\sum_{t=1}^n \hat{r}_{t,k} - \sum_{t=1}^n \mathbb{E}_{j \sim q_t} [\hat{r}_{t,j}] = \sum_{t=1}^n \hat{r}_{t,k} + \sum_{t=1}^n \mathbb{E}_{i \sim q_t} [\eta_t \hat{r}_{t,i}^2] + \sum_{t=1}^n J_t(\eta_t) - J_{t-1}(\eta_t) \quad (18)$$

**Potential manipulation** We can do an Abel transformation on the sum of potentials in (18), namely obtaining

$$\sum_{t=1}^n J_t(\eta_t) - J_{t-1}(\eta_t) = \sum_{t=1}^{n-1} (J_t(\eta_t) - J_t(\eta_{t+1})) + J(\eta_n)$$

We know  $J'(\eta) \leq 0$  and so  $J$  is decreasing and since  $\eta_{t+1} \leq \eta_t$ , we have  $J(\eta_{t+1}) \geq J(\eta_t)$  or  $(J_t(\eta_t) - J_t(\eta_{t+1})) \leq 0$ , so that for any fixed  $k$

$$\begin{aligned} \sum_{t=1}^n J_t(\eta_t) - J_{t-1}(\eta_t) &\leq J_n(\eta_n) \leq \frac{\log(M)}{\eta_n} - \frac{1}{\eta_n} \log \left( \sum_{i=1}^K \exp(\eta_n \hat{R}_{n,i}) \right) \\ &\stackrel{(*)}{\leq} \frac{\log(M)}{\eta_n} - \frac{1}{\eta_n} \log \left( \exp(\eta_n \hat{R}_{n,k}) \right) \\ &= \frac{\log(M)}{\eta_n} - \sum_{t=1}^n \hat{r}_{t,k} \end{aligned} \quad (19)$$

where  $(*)$  follows because  $\exp$  is positive and  $-\log$  is decreasing (notice that we drop  $M-1$  terms from the sum). Plugging (19) into (18), we obtain

$$\begin{aligned} \sum_{t=1}^n \hat{r}_{t,k} - \sum_{t=1}^n \mathbb{E}_{j \sim q_t} [\hat{r}_{t,j}] &\leq \sum_{t=1}^n \hat{r}_{t,k} + \sum_{t=1}^n \mathbb{E}_{i \sim q_t} [\eta_t \hat{r}_{t,i}^2] + \sum_{t=1}^n J_t(\eta_t) - J_{t-1}(\eta_t) \\ &\leq \sum_{t=1}^n \hat{r}_{t,k} + \sum_{t=1}^n \mathbb{E}_{i \sim q_t} [\eta_t \hat{r}_{t,i}^2] + \frac{\log(M)}{\eta_n} - \sum_{t=1}^n \hat{r}_{t,k} \\ &\leq \sum_{t=1}^n \mathbb{E}_{i \sim q_t} [\eta_t \hat{r}_{t,i}^2] + \frac{\log(M)}{\eta_n} \\ &= \sum_{t=1}^n \eta_t \sum_{j=1}^M q_{t,j} \hat{r}_{t,j}^2 + \frac{\log(M)}{\eta_n}. \end{aligned}$$

□

We expressed in Appendix A.3, that the model selection regret of ALEXP, is closely tied to the bias and variance of the reward estimates  $\hat{r}_{t,j}$ . The following lemma formalizes this claim.

**Lemma 18.** (Anytime Generic regret bound) *If  $\eta_t$  is picked such that  $\eta_t \hat{r}_{t,j} \leq 1$  for all  $1 \leq j \leq M$  and  $1 \leq t$  almost surely, then Algorithm 1 satisfies with probability greater than  $1 - 2\delta/3$ , that simultaneously for all  $n \geq 1$*

$$R(n, i) \leq 2B \sum_{t=1}^n \gamma_t + \frac{\log M}{\eta_n} + \sum_{t=1}^n \eta_t \sum_{j=1}^M q_{t,j} \hat{r}_{t,j}^2 + \sum_{t=1}^n (\omega_{t,i} + \sum_{j=1}^M q_{t,j} \omega_{t,j})$$

$$+ 10B\sqrt{n((\log \log nB^2)_+ + \log(12/\delta))}$$

where  $\omega_{t,i} = |r_{t,i} - \hat{r}_{t,i}|$ .

*Proof of Lemma 18.* At each step  $t$ , if  $\alpha_t = 1$ , then ALEXP accumulates a regret of at most  $2B$ , since  $\|\boldsymbol{\theta}\|_\infty \leq B$  and  $\|\phi(\cdot)\| \leq 1$ . We can decompose the regret as,

$$R(n, i) \leq \sum_{t=1}^n 2B\alpha_t + (r_{t,i} - r_t)(1 - \alpha_t)$$

For the first term, by Lemma 24, we have

$$2B \sum_{t=1}^n \alpha_t \leq 2B \left( \sum_{t=1}^n \gamma_t + \frac{5}{2} \sqrt{n((\log \log n)_+ + \log(4/\delta_1))} \right).$$

simultaneously for all  $n \geq 1$ , with probability  $1 - \delta_1$ . Let  $\hat{r}_t := \sum_{j=1}^M q_{t,j} \hat{r}_{t,j}$ . We may re-write the second term of the regret as follows,

$$\begin{aligned} \sum_{t=1}^n (1 - \alpha_t) (r_{t,i} - r_t) &\leq \sum_{t=1}^n (1 - \alpha_t) \left[ (r_{t,i} - \hat{r}_{t,i}) + (\hat{r}_{t,i} - \hat{r}_t) + (\hat{r}_t - r_t) \right] \\ &\leq \sum_{t=1}^n \omega_{t,i} + (1 - \alpha_t) \left[ (\hat{r}_{t,i} - \hat{r}_t) + (\hat{r}_t - r_t) \right] \end{aligned}$$

We bound the second term on the right hand side, using Lemma 17

$$\sum_{t=1}^n (1 - \alpha_t) (\hat{r}_{t,i} - \hat{r}_t) \leq \sum_{t=1}^n (\hat{r}_{t,i} - \hat{r}_t) \leq \frac{\log M}{\eta_n} + \sum_{t=1}^n \eta_t \sum_{j=1}^M q_{t,j} \hat{r}_{t,j}^2.$$

As for the third term,

$$\begin{aligned} (1 - \alpha_t) \left( \sum_{j=1}^M q_{t,j} \hat{r}_{t,j} - r_t \right) &= (1 - \alpha_t) \left[ \sum_{j=1}^M q_{t,j} (\hat{r}_{t,j} - r_{t,j} + r_{t,j}) - r_t \right] \\ &\leq \sum_{j=1}^M q_{t,j} \omega_{t,j} + (1 - \alpha_t) \left( r_t - \sum_{j=1}^M q_{t,j} r_{t,j} \right). \end{aligned}$$

It remains to bound the deviation term. For all  $t$  that satisfy  $\alpha_t = 0$ , the action/model is selected according to  $q_{t,j}$ , therefore the conditional expectation of  $r_t$  can be written as

$$\mathbb{E}_{t-1} r_t = \sum_{j=1}^M q_{t,j} r_{t,j}$$

The sequence  $X_t := r_t - \mathbb{E}_{t-1} r_t$  is a martingale difference sequence adapted to the history  $H_t$ , since for every  $t \geq 1$ ,

$$\mathbb{E}_{t-1} X_t = \mathbb{E} [r_t - \mathbb{E}_{t-1} r_t | H_{t-1}] = 0.$$

Since  $r_t \leq B$ , then  $X_t \leq 2B$  almost surely, which allows for an application of anytime Azuma-Hoeffding (Lemma 26):

$$\mathbb{P} \left( \exists n : \sum_{t=1}^n \left( r_t - \sum_{j=1}^M q_{t,j} r_{t,j} \right) \geq \frac{5B}{2} \sqrt{n ((\log \log n B^2)_+ + \log(2/\delta_2))} \right) \leq \delta_2$$

which, in turn leads us to,

$$\begin{aligned} \sum_{t=1}^n (1 - \alpha_t) \left( r_t - \sum_{j=1}^M q_{t,j} r_{t,j} \right) &\stackrel{\text{a.s.}}{\leq} \sum_{t=1}^n \left( r_t - \sum_{j=1}^M q_{t,j} r_{t,j} \right) \\ &\stackrel{\text{w.h.p.}}{\leq} \frac{5B}{2} \sqrt{n ((\log \log n B^2)_+ + \log(2/\delta_2))} \end{aligned}$$

simultaneously for all  $n \geq 1$ . We set  $\delta_1 = \delta_2 = \delta/3$ , take a union bound and put the terms together obtaining,

$$\begin{aligned} R(n, i) &\leq 2B \sum_{t=1}^n \gamma_t + \frac{\log M}{\eta_n} + \eta \sum_{t=1}^n \sum_{j=1}^M q_{t,j} \hat{r}_{t,j}^2 + \sum_{t=1}^n (\omega_{t,i} + \sum_{j=1}^M q_{t,j} \omega_{t,j}) \\ &\quad + \frac{5B}{2} \sqrt{n ((\log \log n B^2)_+ + \log(6/\delta))} + 5B \sqrt{n ((\log \log n)_+ + \log(12/\delta))} \end{aligned}$$

We upper bound the sum of last two terms to conclude the proof.  $\square$

The next two lemmas bound the bias and variance terms which appear in Lemma 18.

**Lemma 19** (Anytime Bound on the Bias Term). *If the regularization parameter of Lasso is chosen at every step as*

$$\lambda_t = \frac{2\sigma}{\sqrt{t}} \sqrt{1 + \frac{5}{\sqrt{2}} \sqrt{d (\log(2M/\delta) + (\log \log d)_+)} + \frac{12}{\sqrt{2}} (\log(2M/\delta) + (\log \log d)_+)}$$

and  $\gamma_t = O(t^{-1/4})$ , then with probability greater than  $1 - \delta$ , simultaneously for all  $n \geq 1$ ,

$$\sum_{t=1}^n |\hat{r}_{t,i} - r_{t,i}| \leq n^{3/4} C_{\min}^{-1} C(M, \delta, d) \left( 1 + n^{-3/8} C_{\min}^{-1} \sqrt{\log(Md/\delta) + (\log \log n)_+} \right)$$

where

$$C(M, \delta, d) := C\sigma \sqrt{1 + \sqrt{d (\log(M/\delta) + (\log \log d)_+)} + (\log(M/\delta) + (\log \log d)_+)}$$

and  $C$  is an absolute constant.

*Proof of Lemma 19.* By the definition of the expected reward and its estimate,

$$\sum_{t=1}^n |\hat{r}_{t,i} - r_{t,i}| = \sum_{t=1}^n \left| \int_{\mathcal{X}} (r(\mathbf{x}) - \hat{r}_t(\mathbf{x})) dp_{t+1,i}(\mathbf{x}) \right|$$

$$\begin{aligned}
&\leq \sum_{t=1}^n \int_{\mathcal{X}} |r(\mathbf{x}) - \hat{r}_t(\mathbf{x})| dp_{t+1,i}(\mathbf{x}) \\
&\stackrel{\text{C.S.}}{\leq} \sum_{t=1}^n \int_{\mathcal{X}} \|\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_t\|_2 \|\phi(\mathbf{x})\|_2 dp_{t+1,i}(\mathbf{x}) \\
&\stackrel{\text{bdd. } \phi}{\leq} \sum_{t=1}^n \|\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_t\|_2 \int_{\mathcal{X}} dp_{t+1,i}(\mathbf{x}) = \sum_{t=1}^n \|\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_t\|_2.
\end{aligned}$$

From Theorem 3, with probability greater than  $1 - \delta/2$  simultaneously for all  $n \geq 1$ ,

$$\sum_{t=1}^n |\hat{r}_{t,i} - r_{t,i}| \leq \sum_{t=1}^n \frac{4\sqrt{10}\lambda_t}{\kappa^2(\Phi_t, 2)} = \tilde{C}(M, \delta, d) \sum_{t=1}^n \frac{1}{\kappa^2(\Phi_t, 2)\sqrt{t}}$$

where,

$$\tilde{C}(M, \delta, d) := 8\sigma \sqrt{1 + \frac{5}{\sqrt{2}} \sqrt{d(\log(4M/\delta) + (\log \log d)_+)} + \frac{12}{\sqrt{2}} (\log(4M/\delta) + (\log \log d)_+)}.$$

From Lemma 8, there exist absolute constants  $C_1, C_2$  for which,

$$\mathbb{P}\left(\forall t \geq 1 : \kappa^2(\Phi_t, 2) \geq C_1 C_{\min} t^{-1/4} - C_2 t^{-5/8} \sqrt{\log(Md/\delta) + (\log \log t)_+}\right) \geq 1 - \delta.$$

Using Taylor approximation we observe that,  $\frac{1}{1-x^{-1}} = 1 + x^{-1} + o(x^{-1}) = \mathcal{O}(1 + x^{-1})$ . Therefore, there exists absolute constant  $C_3, C_4$ , for which with probability greater than  $1 - \delta$  for all  $t \geq 1$

$$\begin{aligned}
\sum_{t=1}^n \frac{\tilde{C}(M, \delta, d)}{\kappa^2(\Phi_t, 2)\sqrt{t}} &\leq \sum_{t=1}^n \frac{\tilde{C}(M, \delta, d)}{\sqrt{t}} \frac{1}{C_1 C_{\min} t^{-1/4} - C_2 t^{-5/8} \sqrt{\log(Md/\delta) + (\log \log t)_+}} \\
&\leq \sum_{t=1}^n \frac{\tilde{C}(M, \delta, d)}{\sqrt{t}} \frac{\tilde{C}_3}{C_{\min} t^{-1/4}} \left(1 + \frac{C_2 t^{-5/8} \sqrt{\log(Md/\delta) + (\log \log t)_+}}{C_1 C_{\min} t^{-1/4}}\right) \\
&\leq \sum_{t=1}^n \frac{\tilde{C}_3 \tilde{C}(M, \delta, d) t^{-1/4}}{C_{\min}} \left(1 + \tilde{C}_4 \frac{t^{-3/8}}{C_{\min}} \sqrt{\log(Md/\delta) + (\log \log t)_+}\right) \\
&= \frac{C_3 \tilde{C}(M, \delta, d) n^{3/4}}{C_{\min}} \left(1 + C_4 \frac{n^{-3/8}}{C_{\min}} \sqrt{\log(Md/\delta) + (\log \log n)_+}\right).
\end{aligned}$$

□

**Lemma 20** (Anytime Bound on Variance Term). *Suppose  $\lambda_t$  is chosen according to Lemma 19,  $\gamma_t = \mathcal{O}(t^{-1/4})$  and  $\eta_t = \mathcal{O}(C_{\min} t^{-1/2}/C(M, \delta, d))$ . Then with probability greater than  $1 - \delta$ , the following holds simultaneously for all  $n \geq 1$*

$$\hat{r}_{t,j} \leq \frac{4\sqrt{10}\lambda_t}{\kappa^2(\Phi_t, 2)} + B, \quad \forall j \in [M]$$

$$\begin{aligned} \sum_{t=1}^n \eta_t \sum_{j=1}^M q_{t,j} \hat{r}_{t,j}^2 &\leq C_1 B^2 C_{\min} \sqrt{n} + C_2 B n^{1/4} \left( 1 + \frac{n^{-3/8}}{C_{\min}} \sqrt{\log(Md/\delta) + (\log \log n)_+} \right) \\ &\quad + C(M, \delta, d) \frac{\log n}{C_{\min}} \left( 1 + \frac{n^{-3/8}}{C_{\min}} \sqrt{\log(Md/\delta) + (\log \log n)_+} \right) \end{aligned}$$

where  $C_i$  are absolute constants, and  $C(M, \delta, d)$  is as defined in Lemma 19, up to constant factors.

*Proof of Lemma 20.* We start by upper bounding  $\hat{r}_{t,j}$ . For all  $j$  and  $t$  it holds that:

$$\hat{r}_{t,j} = \int_{\mathcal{X}} \langle \hat{\boldsymbol{\theta}}_t, \boldsymbol{\phi}(\mathbf{x}) \rangle dp_{t+1,j}(\mathbf{x}) \leq \|\hat{\boldsymbol{\theta}}_t\| \int_{\mathcal{X}} \|\boldsymbol{\phi}(\mathbf{x})\| dp_{t+1,j}(\mathbf{x}) \leq \|\hat{\boldsymbol{\theta}}_t\| \leq B + \|(\hat{\boldsymbol{\theta}}_t - \boldsymbol{\theta})\|$$

since  $\|\boldsymbol{\theta}\|_2 \leq B$ . To bound the last term, we only need to invoke Theorem 3, which, in turn, will simultaneously bound  $\hat{r}_{t,j}$  for all  $j = 1, \dots, M$ :

$$\mathbb{P} \left( \forall t \geq 1, \forall j \in [M] : \hat{r}_{t,j} \leq \frac{4\sqrt{10}\lambda_t}{c_{\kappa,t}^2} + B \right) \geq 1 - \delta$$

Which implies for all  $t \geq 1$ ,

$$\sum_{j=1}^M q_{t,j} \hat{r}_{t,j}^2 \leq \left( \frac{4\sqrt{10}\lambda_t}{c_{\kappa,t}^2} + B \right)^2 \sum_{j=1}^M q_{t,j} = \frac{160\lambda_t^2}{c_{\kappa,t}^4} + B^2 + \frac{8B\sqrt{10}\lambda_t}{c_{\kappa,t}^2}.$$

For the last term, similar to the proof of Lemma 19 we have,

$$\sum_{t=1}^n \eta_t \frac{8B\sqrt{10}\lambda_t}{c_{\kappa,t}^2} \leq C_1 B n^{1/4} \left( 1 + \frac{n^{-3/8}}{C_{\min}} \sqrt{\log(Md/\delta) + (\log \log n)_+} \right)$$

for some absolute constant  $C_1$ . We treat the squared term similarly,

$$\begin{aligned} \sum_{t=1}^n \eta_t \frac{160\lambda_t^2}{c_{\kappa,t}^4} &\leq C(M, \delta, d) \sum_{t=1}^n \frac{\bar{C}_3}{t C_{\min}} \left( 1 + \bar{C}_4 \frac{t^{-3/8}}{C_{\min}} \sqrt{\log(Md/\delta) + (\log \log t)_+} \right) \\ &\leq C(M, \delta, d) \frac{C_3 \log n}{C_{\min}} \left( 1 + C_4 \frac{n^{-3/8}}{C_{\min}} \sqrt{\log(Md/\delta) + (\log \log n)_+} \right). \end{aligned}$$

Note that the last inequality is not tight. This term will not be fastest growing term in the regret, so we have little motivation to bound it tightly. Therefore,

$$\begin{aligned} \sum_{t=1}^n \eta_t \sum_{j=1}^M q_{t,j} \hat{r}_{t,j}^2 &\leq C_1 B^2 C_{\min} \sqrt{n} + C_2 B n^{1/4} \left( 1 + \frac{n^{-3/8}}{C_{\min}} \sqrt{\log(Md/\delta) + (\log \log t)_+} \right) \\ &\quad + C(M, \delta, d) \frac{\log n}{C_{\min}} \left( 1 + C_4 \frac{n^{-3/8}}{C_{\min}} \sqrt{\log(Md/\delta) + (\log \log n)_+} \right) \end{aligned}$$

where  $C_i$  are absolute constants. □

**Proof of Lemma 16.** We start by conditioning on the event  $E$  that  $\eta_t$  is picked such that  $\eta_t \hat{r}_{t,j}$  for all  $t \geq 1$  and  $j = 1, \dots, M$ . Then by application of Lemma 18 we get, with probability greater than  $1 - 2\delta/3$ ,

$$\begin{aligned} R(n, i) &\leq 2B \sum_{t=1}^n \gamma_t + \frac{\log M}{\eta_n} + \sum_{t=1}^n \eta_t \sum_{j=1}^M q_{t,j} \hat{r}_{t,j}^2 \\ &\quad + \sum_{t=1}^n (\omega_{t,i} + \sum_{j=1}^M q_{t,j} \omega_{t,j}) \\ &\quad + 10B \sqrt{n ((\log \log n B^2)_+ + \log(12/\delta))} \end{aligned}$$

We invoke Lemma 19 and Lemma 20 with  $\delta \rightarrow \delta/3$  take a union bound, to bound the variance and  $\omega_{t,i}$  terms as well. These lemmas require one application of Theorem 3 to hold simultaneously and no additional union bound is required between them, since the randomness comes only from the confidence interval over  $\hat{\theta}_t$ .

$$\begin{aligned} R(n, i) &\leq C_1 B n^{3/4} + \frac{\log M}{\eta_n} \\ &\quad + C_2 B^2 C_{\min} \sqrt{n} + C_3 B n^{1/4} \left( 1 + \frac{n^{-3/8}}{C_{\min}} \sqrt{\log(Md/\delta) + (\log \log n)_+} \right) \\ &\quad + C(M, \delta, d) \frac{\log n}{C_{\min}} \left( 1 + \frac{n^{-3/8}}{C_{\min}} \sqrt{\log(Md/\delta) + (\log \log n)_+} \right) \\ &\quad + n^{3/4} C(M, \delta, d) \left( 1 + \frac{n^{-3/8}}{C_{\min}} \sqrt{\log(Md/\delta) + (\log \log n)_+} \right) \\ &\quad + 10B \sqrt{n ((\log \log n B^2)_+ + \log(12/\delta))} \end{aligned}$$

with probability greater than  $1 - \delta$ , conditioned on event  $E$ . First, suppose this condition is met with probability greater than  $1 - 2\delta$ , we will later verify this. Now let  $\mathcal{B} = \mathcal{B}(\delta, M, d, n, B, \sigma, C_{\min})$  denote the right-hand-side of the regret inequality above. By the chain rule we may write,

$$\begin{aligned} \mathbb{P}(\text{Reg}(n, i) \leq \mathcal{B}) &\geq \mathbb{P}\left(R(n, i) \leq \mathcal{B} \mid \forall t \in [n], j \in [M] : \eta_t \hat{r}_{t,j} \leq 1\right) \mathbb{P}(\forall t \in [n], j \in [M] : \eta_t \hat{r}_{t,j} \leq 1) \\ &\geq \mathbb{P}\left(R(n, i) \leq \mathcal{B} \mid \forall t \in [n], j \in [M] : \eta_t \hat{r}_{t,j} \leq 1\right) (1 - 2\delta) \\ &\geq (1 - \delta)(1 - 2\delta) \geq 1 - 3\delta. \end{aligned}$$

It remains to verify that event  $E$  is met with probability  $1 - 2\delta$ . Recall that  $\eta_t = \mathcal{O}(C_{\min}/\sqrt{t}C(M, \delta, d))$ , and that from Lemma 8 with probability  $1 - \delta$ ,

$$\frac{C_{\min}}{4\sqrt{t}C(M, \delta, d)} \leq \frac{C_1 C_{\min} t^{1/4} - C_2 t^{-1/8} \sqrt{\log(Md/\delta) + (\log \log t)_+}}{4\sqrt{10}C(M, \delta, d)} \leq \frac{\kappa^2(\Phi_t, 2)}{4\sqrt{10}\lambda_t}$$



Therefore, from Lemma 20, there exists  $C_\eta$  such that  $\eta_t = C_\eta C_{\min}/B\sqrt{t}C(M, \delta, d)$  satisfying,

$$\mathbb{P}(\forall t \geq 1, j \in [M] : \eta_t \hat{r}_{t,j} \leq 1) \geq 1 - 2\delta$$

The proof is then finished by setting  $\delta \leftarrow 3\delta$  (and updating the absolute constants).  $\square$

## D.2 Proof of Virtual Regret

**Proposition 21.** *For any fixed  $\tilde{\lambda} > 0$ , there exists an absolute constant  $C_1$  such that*

$$\mathbb{P}(\forall t \geq 1 : \|\hat{\beta}_{t,j^*} - \theta_{j^*}\|_2 \leq \omega(t, \delta, d)) \geq 1 - \delta.$$

where

$$\omega(t, \delta, d) := C_1 \sqrt{\frac{\sigma^2 d \log\left(\frac{t}{\tilde{\lambda}d} + 1\right) + 2\sigma^2 \log(1/\delta) + \tilde{\lambda}B^2}{\tilde{\lambda} + C_{\min} t^{3/4}}} \left(1 + C_{\min}^{-1} t^{-3/8} \sqrt{\log(Md/\delta) + (\log \log t)_+}\right).$$

Moreover, for  $u_{t,j^*}(\cdot) := \hat{\beta}_{t,j^*}^\top \phi_{j^*}(\cdot) + \omega(t, \delta, d)$ ,

$$\mathbb{P}(\forall t \geq 1, \mathbf{x} \in \mathcal{X} : r(\mathbf{x}) \leq u_{t,j^*}(\mathbf{x})) \geq 1 - \delta.$$

*Proof of Proposition 21.* Define for convenience  $V_t = \Phi_{t,j^*}^\top \Phi_{t,j^*} + \tilde{\lambda}I$ . We first observe that

$$\hat{\beta}_{t,j^*} = V_t^{-1}(\Phi_{t,j^*})^\top \mathbf{y}_t$$

We can apply results from [Abbasi-Yadkori et al. \(2011\)](#) to get an anytime-valid confidence set. Their Theorem 2 asserts that with probability  $1 - \delta$ , for all  $t \geq 1$  we have<sup>1</sup>

$$\|\hat{\beta}_{t,j^*} - \theta_{j^*}\|_{V_t}^2 \leq \beta_t$$

where

$$\beta_t = 2\sigma^2 \log\left(\frac{\det(V_t)^{1/2}}{\det(\tilde{\lambda}I)^{1/2}\delta}\right) + \tilde{\lambda}B^2$$

Clearly,  $V_t \succeq \lambda_{\min}(V_t)I$ , and therefore

$$\|\hat{\beta}_{t,j^*} - \theta_{j^*}\|_2 \leq \sqrt{\frac{\beta_t}{\lambda_{\min}(V_t)}}$$

uniformly over time. Our assumption is that  $\|\phi_j(\mathbf{x})\| \leq 1$ , and hence, denoting by  $\nu_i$  the eigenvalues of  $V_t$ , the geometric-arithmetic mean inequality yields

$$\det(V_t) \leq \prod_{i=1}^d \nu_i \leq \left(\frac{1}{d} \text{trace}(V_t)\right)^d.$$

---

1. Their theorem statement is slightly different, but they prove the stronger version we state below.

Given that

$$\text{trace}(V_t) = \sum_{i=1}^d \sum_{s=1}^t (\phi_{j^*}(x))_i^2 + \tilde{\lambda}d \leq t + \tilde{\lambda}d$$

we can conclude that

$$\beta_t \leq 2\sigma^2 \log \left( \frac{(t/d + \tilde{\lambda})^{d/2}}{\tilde{\lambda}^{d/2}\delta} \right) + \tilde{\lambda}B^2 = d\sigma^2 \log \left( \frac{t}{\tilde{\lambda}d} + 1 \right) + 2\sigma^2 \log(1/\delta) + \tilde{\lambda}B^2$$

We note that

$$\lambda_{\min}(V_t) = \lambda_{\min}(\Phi_{t,j}^\top \Phi_{t,j}) + \tilde{\lambda}.$$

Then due to Lemma 9, there exist absolute constants  $C_1$  and  $C_2$  such that for all  $t \geq 1$ ,

$$\lambda_{\min}(V_t) \geq \tilde{\lambda} + C_1 C_{\min} t^{3/4} - C_2 t^{3/8} \sqrt{\log(Md/\delta) + (\log \log t)_+}$$

therefore, there exists  $C_3$  such that

$$\begin{aligned} \frac{1}{\sqrt{\lambda_{\min}(V_t)}} &\leq \frac{1}{\sqrt{\tilde{\lambda} + C_1 C_{\min} t^{3/4} - C_2 t^{3/8} \sqrt{\log(Md/\delta) + (\log \log t)_+}}} \\ &\leq \frac{C_3}{\sqrt{\tilde{\lambda} + C_1 C_{\min} t^{3/4}}} \left( 1 + \frac{t^{3/8} \sqrt{\log(Md/\delta) + (\log \log t)_+}}{\tilde{\lambda} + C_1 C_{\min} t^{3/4}} \right) \\ &\leq \frac{C_3}{\sqrt{\tilde{\lambda} + C_1 C_{\min} t^{3/4}}} \left( 1 + \frac{t^{3/8} \sqrt{\log(Md/\delta) + (\log \log t)_+}}{\tilde{\lambda} + C_1 C_{\min} t^{3/4}} \right) \end{aligned}$$

with high probability for all  $t \geq 1$ . Setting

$$\omega(t, \delta, d) = C_4 \sqrt{\frac{\sigma^2 d \log \left( \frac{t}{\tilde{\lambda}d} + 1 \right) + 2\sigma^2 \log(1/\delta) + \tilde{\lambda}B^2}{\tilde{\lambda} + C_{\min} t^{3/4}}} \left( 1 + C_{\min}^{-1} t^{-3/8} \sqrt{\log(Md/\delta) + (\log \log t)_+} \right)$$

where  $C_4$  is an absolute constant. The upper confidence bound then simply follows: for any  $\mathbf{x} \in \mathcal{X}$

$$r(\mathbf{x}) - \hat{\beta}_{t,j^*}^\top \phi_{j^*}(\mathbf{x}) = \langle \boldsymbol{\theta}_{j^*} - \hat{\boldsymbol{\beta}}_{t,j^*}, \phi_{j^*}(\mathbf{x}) \rangle \leq \left\| \boldsymbol{\theta}_{j^*} - \hat{\boldsymbol{\beta}}_{t,j^*} \right\|_2 \|\phi_{j^*}(\mathbf{x})\|_2 \leq \omega(t, \delta, d)$$

where the last inequality holds with high probability simultaneously for all  $t \geq 1$ .  $\square$

**Proof of Lemma 15.** Using Proposition 21 and the Cauchy-Schwarz inequality we obtain,

$$\begin{aligned} \tilde{R}_{j^*}(n) &= \sum_{t=1}^n r(\mathbf{x}^*) - r(\tilde{\mathbf{x}}_{t,j}) \\ &= \sum_{t=1}^n r(\mathbf{x}^*) - \hat{r}_t(\mathbf{x}^*) + \hat{r}_t(\mathbf{x}^*) - \hat{r}_t(\tilde{\mathbf{x}}_{t,j}) + \hat{r}_t(\tilde{\mathbf{x}}_{t,j}) - r(\tilde{\mathbf{x}}_{t,j}) \\ &\leq \sum_{t=1}^n \left\| \boldsymbol{\theta}_j - \hat{\boldsymbol{\theta}}_{t,j} \right\|_2 (\|\phi_j(\mathbf{x}^*)\|_2 + \|\phi_j(\tilde{\mathbf{x}}_{t,j})\|_2) + \hat{r}_t(\mathbf{x}^*) - \hat{r}_t(\tilde{\mathbf{x}}_{t,j}) \end{aligned}$$

$$\leq \sum_{t=1}^n \omega(t, \delta, d) (\|\phi_j(\mathbf{x}^*)\|_2 + \|\phi_j(\tilde{\mathbf{x}}_{t,j})\|_2) + \hat{r}_t(\mathbf{x}^*) - \hat{r}_t(\tilde{\mathbf{x}}_{t,j})$$

with probability  $1 - \delta$ . If the agent selects actions greedily, then  $\hat{r}_t(\mathbf{x}^*) \leq \hat{r}_t(\tilde{\mathbf{x}}_{t,j})$ , and

$$\tilde{R}_{j^*}(n) \leq \sum_{t=1}^n \omega(t, \delta, d) (\|\phi_j(\mathbf{x}^*)\|_2 + \|\phi_j(\tilde{\mathbf{x}}_{t,j})\|_2) \leq \sum_{t=1}^n 2\omega(t, \delta, d)$$

since the feature map is normalized to satisfy  $\|\phi_j(\cdot)\| \leq 1$ . If the agent selects actions optimistically according to the upper confidence bound of Proposition 21, then

$$\hat{r}_t(\tilde{\mathbf{x}}_{t,j}) + \omega(t, \delta, d)\|\phi_j(\tilde{\mathbf{x}}_{t,j})\| \geq \hat{r}_t(\mathbf{x}^*) + \omega(t, \delta, d)\|\phi_j(\mathbf{x}^*)\|$$

which implies,

$$\tilde{R}_{j^*}(n) \leq \sum_{t=1}^n 2\omega(t, \delta, d)\|\phi_j(\tilde{\mathbf{x}}_{t,j})\|_2 \leq \sum_{t=1}^n 2\omega(t, \delta, d).$$

Then due to Proposition 21, with probability greater than  $1 - \delta$ , simultaneously for all  $n \geq 1$ ,

$$\begin{aligned} \sum_{t=1}^n \omega(t, \delta, d) &\leq \sum_{t=1}^n C_1 \sqrt{\frac{\sigma^2 d \log\left(\frac{t}{\lambda d} + 1\right) + 2\sigma^2 \log(1/\delta) + \bar{\lambda} B^2}{\bar{\lambda} + C_{\min} t^{3/4}}} \left(1 + t^{-3/8} \sqrt{\log\left(\frac{M d}{\delta}\right) + (\log \log t)_+}\right) \\ &\leq \tilde{C}_1 n^{5/8} \sqrt{\frac{\sigma^2 d \log\left(\frac{n}{\lambda d} + 1\right) + 2\sigma^2 \log(1/\delta) + \bar{\lambda} B^2}{C_{\min}}} \left(1 + C_{\min}^{-1} n^{-\frac{3}{8}} \sqrt{\log\left(\frac{M d}{\delta}\right) + (\log \log n)_+}\right) \end{aligned}$$

concluding the proof.  $\square$

## Appendix E. Time-Uniform Concentration Inequalities

We will make use of the elegant concentration results in Howard et al. (2021), which analyzes the boundary of sub-Gamma processes.

**Definition 22** (Sub-Gamma process). Let  $(S_t)_{t=0}^\infty, (V_t)_{t=0}^\infty$  be real-valued processes adapted to  $(\mathcal{F}_t)_{t=1}^\infty$  with  $S_0 = V_0 = 0$  and  $V_t$  non-negative. We say that  $S_t$  is sub-Gamma if for each  $\lambda \in [0, 1/c)$ , there exists a supermartingale  $(M_t(\lambda))_{t=0}^\infty$  w.r.t.  $\mathcal{F}_t$ , such that  $\mathbb{E} M_0 = 1$  and for all  $t \geq 1$ :

$$\exp\left\{\lambda S_t - \frac{\lambda^2}{2(1 - c\lambda)} V_t\right\} \leq M_t(\lambda) \quad a.s.$$

The following is a special case of Theorem 1 in Howard et al. (2021). We have simplified it by making a few straightforward choices for the parameters used originally by Howard et al. (2021), which will yield an easier-to-use bound in our scenario.

**Proposition 23** (Curved Boundary of Sub-Gamma Processes). *Let  $(S_t)$  be sub-Gamma with variance process  $(V_t)$ . Define the boundary*

$$\mathcal{B}_\alpha(v) := \frac{5}{2} \sqrt{\max\{v, 1\} \left( (\log \log ev)_+ + \log\left(\frac{2}{\alpha}\right) \right)} + 3c \left( (\log \log ev)_+ + \log\left(\frac{2}{\alpha}\right) \right),$$

for  $v > 0$ , where  $(x)_+ = \max(0, x)$ . Then,

$$\mathbb{P}(\exists t : S_t \geq \mathcal{B}_\alpha(V_t)) \leq \alpha.$$

*Proof of Proposition 23.* Theorem 1 in Howard et al. (2021) states that if  $(S_t)$  is a sub-Gamma process with variance process  $(V_t)$  then the boundary

$$\mathcal{S}_\alpha(v) = k_1 \sqrt{v \left( s \log \log (\eta v) + \log \left( \frac{\zeta(s)}{\alpha \log^s \eta} \right) \right)} + ck_2 \left( s \log \log (\eta v) + \log \left( \frac{\zeta(s)}{\alpha \log^s \eta} \right) \right).$$

satisfies,

$$\mathbb{P}(\exists t : S_t \geq \mathcal{S}_\alpha(\max(V_t, 1))) \leq \alpha$$

where

$$k_1 := \frac{\eta^{1/4} + \eta^{-1/4}}{\sqrt{2}} \quad \text{and} \quad k_2 := (\sqrt{\eta} + 1)/2.$$

Choosing  $s = 2$  and  $\eta = e$ , we obtain  $\zeta(2) = \pi^2/6 \leq 2$ . Furthermore, we have  $k_1 \leq \frac{3}{2}$  and  $k_2 \leq \frac{3}{2}$ . Then if  $v \geq 1$  (which we will enforce by construction), we compute

$$s \log \log (\eta v) + \log \left( \frac{\zeta(s)}{\alpha \log^s \eta} \right) \leq 2(\log \log ev)_+ + \log \left( \frac{2}{\alpha} \right).$$

Therefore, we can upper bound (using our bounds on  $k_1, k_2$ )

$$\mathcal{S}_\alpha(v) \leq \frac{5}{2} \sqrt{v \left( (\log \log ev)_+ + \log \left( \frac{2}{\alpha} \right) \right)} + 3c \left( (\log \log ev)_+ + \log \left( \frac{2}{\alpha} \right) \right).$$

Now, since the boundary is given by  $\mathcal{S}_\alpha(\max(v, 1))$ , we deduce that

$$\mathcal{B}_\alpha(v) := \frac{5}{2} \sqrt{\max\{v, 1\} \left( (\log \log ev)_+ + \log \left( \frac{2}{\alpha} \right) \right)} + 3c \left( (\log \log ev)_+ + \log \left( \frac{2}{\alpha} \right) \right).$$

is an any-time valid boundary.  $\square$

**Lemma 24** (Time-Uniform Two-sided Bernoulli). *Let  $X_1, \dots, X_s, \dots, X_t$  be a martingale sequence of Bernoulli random variables with conditional mean  $\gamma_s$ . Then for all  $\delta > 0$ ,*

$$\mathbb{P} \left( \exists t : \left| \sum_{s=1}^t (X_s - \gamma_s) \right| \geq \frac{5}{2} \sqrt{t \left( (\log \log t)_+ + \log(4/\delta) \right)} \right) \leq \delta,$$

*Proof of Lemma 24.* By Proposition 23, we know that if  $S_t$  is sub-Gamma with variance process  $V_t$  and scale parameter  $c$ , then

$$\mathbb{P}(\exists t : S_t \geq \mathcal{B}_\delta(V_t)) \leq \delta,$$

where

$$\mathcal{B}_\delta(v) := \frac{5}{2} \sqrt{\max\{1, v\} \left( (\log \log ev)_+ + \log(2/\delta) \right)} + 3c \left( (\log \log ev)_+ + \log(2/\delta) \right).$$

By [Howard et al. \(2020\)](#), we know that if  $(X_t)_{t=1}^\infty$  is a Bernoulli sequence, then  $S_t = \sum_{s=1}^t (X_s - \gamma_s)$  is sub-Gamma with variance process  $V_t = t$  and scale parameter  $c = 0$  (hence, sub-Gaussian). This implies,

$$\mathbb{P} \left( \exists t : \sum_{s=1}^t (X_s - \gamma_s) \geq \frac{5}{2} \sqrt{t ((\log \log t)_+ + \log(2/\delta))} \right) \leq \delta,$$

The above arguments also holds for the sequence  $Z_s = -X_s$ . Then taking a union bound and adjusting  $\delta \leftarrow \delta/2$  concludes the proof.  $\square$

**Lemma 25** (Time-Uniform Bernstein). *Let  $(\xi_i)_{i=1}^\infty$  be a sequence of conditionally standard sub-gaussian variables, where each  $\xi_i$  is  $\mathcal{F}_{i-1} = \sigma(\xi_1, \dots, \xi_i)$  measurable. Then, for  $v_i \in \mathbb{R}$  and  $\delta \in (0, 1]$*

$$\mathbb{P} \left( \exists t : \sum_{i=1}^t (\xi_i^2 - 1)v_i \geq \frac{5}{2} \sqrt{\max \left\{ 1, 4\|\mathbf{v}_t\|_2^2 \right\} \omega_\delta(\|\mathbf{v}_t\|_2) + 12\omega_\delta(\|\mathbf{v}_t\|_2) \max_{i \geq 1} v_i} \right) \leq \delta$$

where,  $\mathbf{v}_t = (v_1, \dots, v_t) \in \mathbb{R}^t$  and  $\omega_\delta(v) := (\log \log(4ev^2))_+ + \log(2/\delta)$ .

*Proof of Lemma 25.* From Lemma 4,  $S_t = \sum_{i=1}^t (\xi_i^2 - 1)v_i$  is sub-Gamma with variance process  $V_t = 4 \sum_{i=1}^t v_i^2$  and  $c = 4 \max_{i \geq 1} v_i$ . By Proposition 23, we know that if  $S_t$  is sub-Gamma with variance process  $V_t$  and scale parameter  $c$ , then

$$\mathbb{P}(\exists t : S_t \geq \mathcal{B}_\delta(V_t)) \leq \delta,$$

where

$$\mathcal{B}_\delta(v) := \frac{5}{2} \sqrt{\max\{1, v\} ((\log \log ev)_+ + \log(2/\delta))} + 3c((\log \log ev)_+ + \log(2/\delta)).$$

$\square$

**Lemma 26** (Time-Uniform Azuma-Hoeffding). *Let  $X_1, \dots, X_n$  be a martingale difference sequence such that  $|X_t| \leq B$  for all  $t > 1$  almost surely. Then for all  $\delta > 0$ ,*

$$\mathbb{P} \left( \exists t : \sum_{s=1}^t X_s \geq \frac{5B}{2} \sqrt{t ((\log \log etB^2)_+ + \log(2/\delta))} \right) \leq \delta,$$

*Proof of Lemma 26.* By Proposition 23, we know that if  $S_t$  is sub-Gamma with variance process  $V_t$  and scale parameter  $c$ , then

$$\mathbb{P}(\exists t : S_t \geq \mathcal{B}_\delta(V_t)) \leq \delta,$$

where

$$\mathcal{B}_\delta(v) := \frac{5}{2} \sqrt{\max\{1, v\} ((\log \log ev)_+ + \log(2/\delta))} + 3c((\log \log ev)_+ + \log(2/\delta)).$$

By ([Howard et al., 2020](#)), we know that if  $(X_t)_{t=1}^\infty$  is  $B$ -bounded martingale difference sequence, then  $S_t = \sum_{s=1}^t X_s$  is sub-Gamma with variance process  $V_t = tB^2$  and scale parameter  $c = 0$ . This implies,

$$\mathbb{P} \left( \exists t : S_t \geq \frac{5B}{2} \sqrt{t ((\log \log etB^2)_+ + \log(2/\delta))} \right) \leq \delta,$$

concluding the proof.  $\square$

## Appendix F. Experiments

**Experiment Setup.** We create a synthetic dataset based on our data model (Section 3), and choose the domain to be 1-dimensional  $\mathcal{X} = [-1, 1]$ . As a natural choice of features, we consider the set of degree- $p$  Legendre polynomials, since they form an orthonormal basis for  $L^2(\mathcal{X})$  if  $p$  grows unboundedly. We construct each feature map, by choosing  $s$  different polynomials from this set, and therefore obtaining  $M = \binom{p+1}{s}$  different models. More formally, we let  $\phi_j(x) = (P_{j_1}(x), \dots, P_{j_s}(x)) \in \mathbb{R}^s$  where  $\{j_1, \dots, j_s\} \subset \{1, \dots, p\}$  and  $P_{j'}$  denotes a degree  $j'$  Legendre polynomial. To construct the reward function, we randomly sample  $j^*$  from  $[M]$ , and draw  $\theta_{j^*}$  from an i.i.d. standard gaussian distribution. We then normalize  $\|\theta_{j^*}\| = 1$ . When sampling from the reward, we add Gaussian noise with standard deviation  $\sigma = 0.01$ . Figure 1 in the appendix shows how the random reward functions may look. For all experiments we set  $n = 100$ , and plot the cumulative regret  $R(n)$  averaged over 20 different random seeds, the shaded areas in all figures show the standard error across these runs.

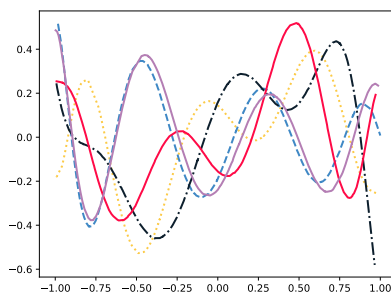


Figure 1: Examples of possible reward functions  $r(\cdot)$  in our experiments.

**Algorithms.** We perform experiments on two UCB algorithms, one with oracle knowledge of  $j^*$ , and a naive one which takes into account all  $M$  feature maps. We run Explore-then-Commit (ETC) by Hao et al. (2020), which explores for a horizon of  $n_0$  steps, performs Lasso once, and then selects actions greedily for the remaining steps. As another baseline, we introduce Explore-then-Select (ETS) that explores for  $n_0$  steps, performs model selection using the sparsity pattern of the Lasso estimator, and then performs UCB on the selected features for the remaining steps. Performance of ETC and ETS depends highly on  $n_0$ , so we tune this hyperparameter separately for each experiment. We also run CORRAL as proposed by Agarwal et al. (2017), with only UCB agents similar to ALEXP. We tune the hyper-parameters of CORRAL as well. Appendix F.1 presents the pseudo-code to all baseline algorithms. To initialize ALEXP we set the rates of  $\lambda_t$ ,  $\gamma_t$  and  $\eta_t$  according to Theorem 1, and perform a light hyper-parameter tuning to choose the scaling constants. For solving (1), we use CELER, a fast solver for the group Lasso (Massias et al., 2018). Every time a UCB policy is used, we set the exploration coefficient  $\beta_t = 2$ . We have included the details and results of our hyper-parameter tuning in Appendix F.1.

**Easy vs. Hard cases.** We construct an easy problem instance, where  $s = 2$ ,  $p = 10$ , and thus  $M = 55$ . Models are lightly correlated since each two model can have at most one Legendre polynomial in common. We also generate an instance with highly correlated feature maps where  $s = 8$  and  $p = 10$ , which will be a harder problem, since out of the total  $M = 55$  models, there are 36 models which have at least 6 Legendre polynomials in common with

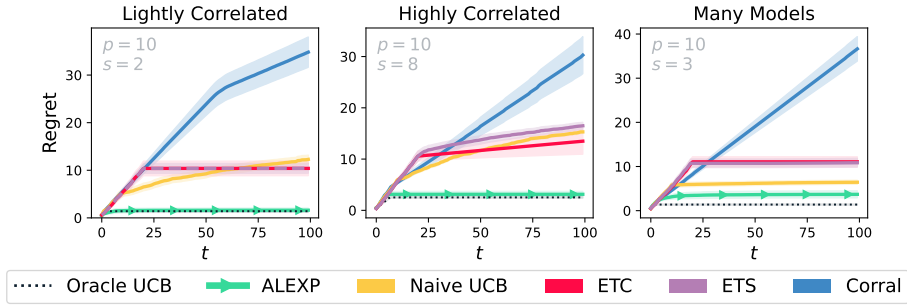


Figure 2: ALEXP can model-select in both orthogonal and correlated classes ( $M = 55$ )

Figure 3: ALEXP performs well on a large class ( $M = 165$ )

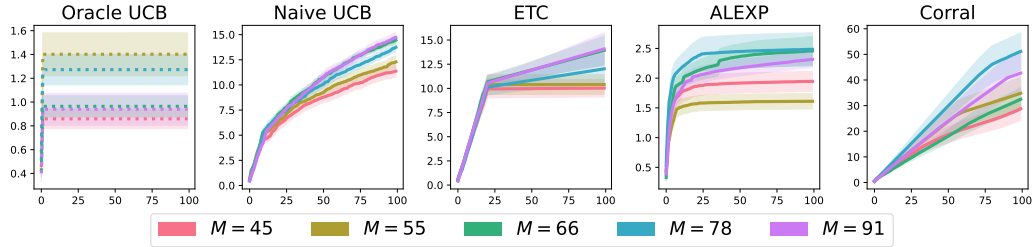


Figure 4: ALEXP is hardly affected by increasing the number of models (y-axis have various scales)

the oracle model  $j^*$ . Figure 2 shows that not only ALEXP is not affected by the correlations between the models, but also it achieves oracle-optimal performance in both cases, showing that our exponential weights technique for model selection is robust to choice of features. ETC and ETS rely on Lasso for model selection, which performs poorly in the case of correlated features. CORRAL uses log-barrier-OMD with an importance-weighted estimator, which has a significantly high variance. We construct another hard instance (Figure 3), where the model class is large ( $s = 3, p = 10, M = 165$ ). ALEXP continues to outperform all baselines with a significant gap. It is clear in the regret curves how explore-then-commit style algorithms are inherently horizon-dependent, and may exhibit linear regret, if stopped at an arbitrary time. This is not an issue with the other algorithms.

**Scaling with  $M$ .** Figure 4 shows how the algorithms scale as  $M$  grows. For this experiment we set  $s = 2$  and change  $p \in \{9, \dots, 13\}$ . While increasing  $M$  hardly affects ALEXP and Oracle UCB, other baselines become less sample efficient.

**Learning Dynamics of ALEXP.** Figure 5 gives some insight into the dynamics of ALEXP when  $M = 165$ . In particular, it shows how ALEXP can rule out sub-optimal agents without ever having queried them. Figure (a) shows the distribution  $\mathbf{q}_t$ , at  $t = 20$  which is roughly equal to the optimal  $n_0$  for ETC in this configuration. The oracle model  $j^*$  is annotated with a star, and has the highest probability of selection. We observe that, already at this time step, more than 80% of the agents are practically ruled out, due to small probability of selection. However, according to Figure (b), which shows  $M_t$  the total number of visited models, less than

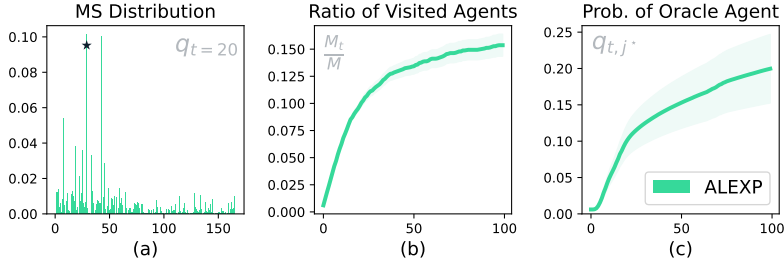


Figure 5: ALEXP can rule out models without ever having queried them ( $M = 165$ )

10% of the models are queried at  $t = 20$ . This is the key practical benefit of ALEXP compared to black-box algorithms such as CORRAL. Lastly, Figure (c) shows how  $q_{t,j^*}$  the probability of selecting the oracle agent changes with time. While this probability is higher than that of the other agents, Figure (c) shows that  $q_{t,j^*}$  is not exceeding 0.25, therefore there is always a probability of greater than 0.75 that we sample another agent, making ALEXP robust to hard problem instances where many agents perform efficiently. We conclude that ALEXP seems to rapidly recognize the better performing agents, and select among them with high probability.

### F.1 Hyper-Parameter Tuning Details

We implement 6 algorithms in our experiments, ETC (Algorithm 4, Hao et al., 2020), ETS (Algorithm 5), CORRAL (Algorithm 6, Agarwal et al., 2017), ALEXP (Algorithm 1), and Lastly UCB (Algorithm 3) with the oracle feature map  $\phi_{j^*}$  and the concatenated feature map  $\phi$ . When algorithms require exploration, e.g., in the case of ETC or ALEXP, we simply set  $\pi = \text{Unif}(\mathcal{X})$ . To ensure that the curves are valid, we run each configuration for 20 different random seeds, i.e. on different random environments.

**UCB.** For all the experiments, we set the exploration coefficient of UCB to  $\beta_t = 2^2$  and choose the regression regularizer from  $\tilde{\lambda} \in \{0.01, 0.1, 0.5\}$ . We use PYTORCH (Paszke et al., 2017) for updating the Upper Confidence bounds, which requires more regularization for longer feature maps (e.g. when  $s = 8, p = 2$ ), to be computationally stable.

**Lasso.** Every time we need to solve Equation (1), we set  $\lambda_t$  according to the *rate* suggested by Theorem 3. To find a suitable constant scaling coefficient, we perform a hyper-parameter tuning experiment sampling 20 values in  $[10^{-5}, 10^0]$ . We choose  $\lambda_0 = 0.009$ , and scale  $\lambda_t$  with it across all experiments.

**ALEXP.** We set the rates for  $\gamma_t$  and  $\eta_t$  as prescribed by Theorem 1. For the scaling constants, we perform a hyper-parameter tuning experiment log-uniformly sampling 20 different configurations from  $\gamma_0 \in [10^{-4}, 10^{-1}]$  and  $\eta_0 \in [10^0, 10^2]$ . For each problem instance (i.e. as  $s$  and  $p$  change) we repeat this process. However we observe that the optimal hyper-parameters work well across all problem instances.

**ETC/ETS.** For these algorithms, we separately tune  $n_0$  for each problem instance. We set  $\lambda_1 \propto \sqrt{\log M/n_0}$  according to Theorem 4.2 of (Hao et al., 2020) and scale it with  $\lambda_0 = 0.009$ , as stated before. We uniformly sample 10 different values where  $n_0 \in [2, 80]$  since the horizon is  $n = 100$ . The optimal value often happens around  $n_0 = 20$ .

2. To achieve the  $\sqrt{dT \log T}$  regret, one has to set  $\beta_t = \mathcal{O}(\sqrt{d \log T})$  as shown in Proposition 21.



---

**Algorithm 1** ALEXP

---

Inputs:  $\pi, (\gamma_t, \eta_t, \lambda_t)$  for  $t \geq 1$   
 Let  $\mathbf{q}_0 \leftarrow \text{Unif}(M)$   
 Initialize base agents  $(p_{0,1}, \dots, p_{0,M})$ .  
**for**  $t \geq 1$  **do**  
   Draw  $\alpha_t \sim \text{Bernoulli}(\gamma_t)$ .  
   Draw  $\mathbf{x}_t \sim \alpha_t \pi + (1 - \alpha_t) \sum_{j=1}^M q_{t-1,j} p_{t-1,j}$ .  
   Observe  $y_t = r(\mathbf{x}_t) + \varepsilon_t$ .  
    $H_t = H_{t-1} \cup \{(\mathbf{x}_t, y_t)\}$ , report to all agents.  
   Get updated policies  $p_{t,1}, \dots, p_{t,M}$ .  
    $\hat{\boldsymbol{\theta}}_t \leftarrow \arg \min \mathcal{L}(\boldsymbol{\theta}; H_t, \lambda_t)$ .  
   Update estimated average return of every agent

$$\hat{r}_{t,j} \leftarrow \mathbb{E}_{\mathbf{x} \sim p_{t,j}} [\hat{\boldsymbol{\theta}}_t^\top \boldsymbol{\phi}(\mathbf{x})], \quad j = 1, \dots, M$$

Update agent probabilities

$$q_{t,j} \leftarrow \frac{\exp(\eta_t \sum_{s=1}^t \hat{r}_{s,j})}{\sum_{i=1}^M \exp(\eta_t \sum_{s=1}^t \hat{r}_{s,i})}$$

**end for**

---

**Algorithm 2** GetPosterior

---

Inputs:  $H_t, \boldsymbol{\phi}, \tilde{\lambda}$   
 Let  $K_t \leftarrow [\boldsymbol{\phi}^\top(\mathbf{x}_i) \boldsymbol{\phi}(\mathbf{x}_j)]_{i,j \leq t}$ , and  $V_t \leftarrow (K_t + \tilde{\lambda}^2 \mathbf{I})$ , and  $\mathbf{k}(\cdot) \leftarrow [\boldsymbol{\phi}^\top(\mathbf{x}_i) \boldsymbol{\phi}(\cdot)]_{i \leq t}$   
 Calculate  $\boldsymbol{\mu}_t(\cdot) \leftarrow \mathbf{k}^T(\cdot) V_t^{-1} \mathbf{y}_t$   
 Calculate  $\sigma_t(\cdot) \leftarrow \sqrt{\boldsymbol{\phi}^\top(\cdot) \boldsymbol{\phi}(\cdot) - \mathbf{k}^\top(\cdot) V_t^{-1} \mathbf{k}(\cdot)}$   
 Return:  $\boldsymbol{\mu}_t, \sigma_t$

---

**CORRAL.** We set the rates of the parameters as  $\gamma = \mathcal{O}(1/n)$  and  $\eta = \mathcal{O}(\sqrt{M/n})$  according to Agarwal et al. (2017, Theorem 5). Then similar to ALEXP, we tune the scaling constants. The procedure for tuning the constants is identical to ALEXP, as in we use the same search interval, and try 10 different configurations for  $\gamma$  and  $\eta$ .

---

**Algorithm 3** UCB

---

Inputs:  $\tilde{\lambda}, \beta_t, \boldsymbol{\phi}$   
**for**  $t = 1, \dots, n$  **do**  
   Choose  $\mathbf{x}_t \arg \max u_{t-1}(\mathbf{x}) = \boldsymbol{\mu}_{t-1}(\mathbf{x}) + \beta_t \sigma_{t-1}(\mathbf{x})$ .    $\triangleright$  Choose actions optimistically  
   Observe  $y_t = r(\mathbf{x}_t) + \varepsilon_t$ .    $\triangleright$  Receive reward  
    $H_t \leftarrow H_{t-1} \cup \{(\mathbf{x}_t, y_t)\}$     $\triangleright$  Append history  
   Update  $\boldsymbol{\mu}_t, \sigma_t \leftarrow \text{GetPosterior}(H_t, \boldsymbol{\phi}, \tilde{\lambda})$   
**end for**

---

---

**Algorithm 4** ETC (Hao et al., 2020)

---

Inputs:  $n_0, n, \lambda_1, \pi$   
Let  $H_0 = \emptyset$   
**for**  $t = 1, \dots, n_0$  **do**  
    Draw  $\mathbf{x}_t \sim \pi$ , observe  $y_t = r(\mathbf{x}_t) + \varepsilon_t$  ▷ Explore  
     $H_t \leftarrow H_{t-1} \cup \{(\mathbf{x}_t, y_t)\}$  ▷ Append history  
**end for**  
 $\hat{\boldsymbol{\theta}}_{n_0} \leftarrow \mathcal{L}(\boldsymbol{\theta}, H_{n_0}, \lambda_1)$  ▷ Perform Lasso once  
**for**  $t = n_0 + 1, \dots, n$  **do**  
    Choose  $\mathbf{x}_t = \arg \max \hat{\boldsymbol{\theta}}_{n_0}^\top \boldsymbol{\phi}(\mathbf{x})$  ▷ Choose actions greedily  
**end for**

---

---

**Algorithm 5** ETS

---

Inputs:  $n_0, n, \lambda_1, \tilde{\lambda}, \beta_t, \pi$   
Let  $H_0 = \emptyset$   
**for**  $t = 1, \dots, n_0$  **do**  
    Draw  $\mathbf{x}_t \sim \pi$ , observe  $y_t = r(\mathbf{x}_t) + \varepsilon_t$  ▷ Explore  
     $H_t \leftarrow H_{t-1} \cup \{(\mathbf{x}_t, y_t)\}$  ▷ Append history  
**end for**  
 $\hat{\boldsymbol{\theta}}_{n_0} \leftarrow \mathcal{L}(\boldsymbol{\theta}, H_{n_0}, \lambda_1)$  ▷ Perform Lasso once  
 $\hat{J} \leftarrow \{j \mid \hat{\boldsymbol{\theta}}_{n_0, j} \neq \mathbf{0}, j \in [M]\}$  ▷ Get sparsity pattern  
 $\boldsymbol{\phi}_{\hat{j}}(\cdot) \leftarrow [\boldsymbol{\phi}_j(\cdot)]_{j \in \hat{J}}$  ▷ Model-select acc. to  $\hat{J}$   
**for**  $t = n_0 + 1, \dots, n$  **do**  
    Choose  $\mathbf{x}_t = \arg \max u_{t-1}(\mathbf{x}) = \mu_{t-1}(\mathbf{x}) + \beta_t \sigma_{t-1}(\mathbf{x})$  ▷ Choose actions optimistically  
    Observe  $y_t = r(\mathbf{x}_t) + \varepsilon_t$   
     $H_t \leftarrow H_{t-1} \cup \{(\mathbf{x}_t, y_t)\}$   
    Update  $\mu_t, \sigma_t \leftarrow \text{GetPosterior}(H_t, \boldsymbol{\phi}_{\hat{j}}, \tilde{\lambda})$   
**end for**

---

---

**Algorithm 6** CORRAL (Agarwal et al., 2017)

---

Inputs:  $n, \gamma, \eta$   
Initialize  $\beta = e^{1/\ln n}$ ,  $\eta_{1,j} = \eta$ ,  $\rho_{1,j} = 2M$  for all  $j = 1, \dots, M$   
Set  $\mathbf{q}_0 = \bar{\mathbf{q}}_0 = \frac{\mathbf{1}}{M}$  and initialize base agents  $(p_{0,1}, \dots, p_{0,M})$ .  
**for**  $t = 1, \dots, n$  **do**  
    Choose  $j_t \sim \bar{\mathbf{q}}_{t-1}$ . ▷ Sample Agent  
    Draw  $\mathbf{x}_t \sim p_{t-1, j_t}$ . ▷ Play action according to agent  $j_t$   
    Observe  $y_t = r(\mathbf{x}_t) + \varepsilon_t$ .  
    Calculate IW estimates  $\hat{r}_{t,j} = \frac{y_t}{\bar{q}_{t-1,j}} \mathbb{I}\{j = j_t\}$  for all  $j = 1, \dots, M$ .  
    Send  $\hat{r}_{t,j} = \frac{y_t}{\bar{q}_{t-1,j}} \mathbb{I}\{j = j_t\}$  to agents and get updated policies  $p_{t,j}$ .  
     $\mathbf{q}_t = \text{LOG-BARRIER-OMD}(\mathbf{q}_{t-1}, \hat{r}_{t, j_t} \mathbf{e}_{j_t}, \boldsymbol{\eta}_t)$  ▷ Update agent probabilities  
     $\bar{\mathbf{q}}_t = (1 - \gamma)\mathbf{q}_t + \gamma \frac{\mathbf{1}}{M}$  ▷ Mix with exploratory distribution  
    **for**  $j = 1, \dots, M$  **do** ▷ Update parameters  
        **if**  $\frac{1}{\bar{q}_{t,j}} > \rho_{t-1,j}$  **then**  $\rho_{t,j} \leftarrow \frac{2}{\bar{q}_{t-1,j}}$ , and  $\eta_{t,j} \leftarrow \beta \eta_{t-1,j}$   
        **else**  $\rho_{t,j} \leftarrow \rho_{t-1,j}$  and  $\eta_{t,j} \leftarrow \eta_{t-1,j}$   
        **end if**  
    **end for**  
**end for**

---

---

**Algorithm 7** LOG-BARRIER-OMD

---

Inputs:  $\mathbf{q}_t, \boldsymbol{\ell}_t, \boldsymbol{\eta}_t$   
Find  $\xi \in [\min_j \ell_{t,j}, \max_j \ell_{t,j}]$  such that  $\sum_{j=1}^M \left( q_{t,j}^{-1} + \eta_{t,j}(\ell_{t,j} - \xi) \right)^{-1} = 1$   
Return:  $\mathbf{q}_{t+1}$  where  $q_{t+1,j}^{-1} = q_{t,j}^{-1} + \eta_{t,j}(\ell_{t,j} - \xi)$  for all  $j \in [M]$

---