

---

# Likelihood Ratio Confidence Sets for Sequential Decision Making

---

Nicolas Emmenegger\*  
ETH Zürich

Mojmír Mutný\*  
ETH Zürich

Andreas Krause  
ETH Zürich

## Abstract

Certifiable, adaptive uncertainty estimates for unknown quantities are an essential ingredient of sequential decision-making algorithms. Standard approaches rely on problem-dependent concentration results and are limited to a specific combination of parameterization, noise family, and estimator. In this paper, we revisit the likelihood-based inference principle and propose to use *likelihood ratios* to construct *any-time valid* confidence sequences without requiring specialized treatment in each application scenario. Our method is especially suitable for problems with well-specified likelihoods, and the resulting sets always maintain the prescribed coverage in a model-agnostic manner. The size of the sets depends on a choice of estimator sequence in the likelihood ratio. We discuss how to provably choose the best sequence of estimators and shed light on connections to online convex optimization with algorithms such as Follow-the-Regularized-Leader. To counteract the initially large bias of the estimators, we propose a reweighting scheme that also opens up deployment in non-parametric settings such as RKHS function classes. We provide a *non-asymptotic* analysis of the likelihood ratio confidence sets size for generalized linear models, using insights from convex duality and online learning. We showcase the practical strength of our method on generalized linear bandit problems, survival analysis, and bandits with various additive noise distributions.

## 1 Introduction

One of the main issues addressed by machine learning and statistics is the estimation of an unknown *model* from noisy observations. For example, in supervised learning, this might concern learning the dependence between an input (covariate)  $x$  and a random variable (observation)  $y$ . In many cases, we are not only interested in an estimate  $\hat{\theta}$  of the true model parameter  $\theta_*$ , but instead in a set of plausible values that  $\theta_*$  could take. Such confidence sets are of tremendous importance in sequential decision-making tasks, where uncertainty is used to drive exploration or risk-aversion needs to be implemented, and covariates are iteratively chosen based on previous observations. This setting includes problems such as bandit optimization, reinforcement learning, or active learning. In the former two, the confidence sets are often used to solve the *exploration-exploitation* dilemma and more generally influence the selection rule (Mukherjee et al., 2022), termination rule (Katz-Samuels and Jamieson, 2020), exploration (Auer, 2002) and/or risk-aversion (Makarova et al., 2021).

When we interact with the environment by gathering data sequentially based on previous confidence sets, we introduce correlations between past noisy observations and future covariates. Data collected in this manner is referred to as *adaptively gathered* (Wasserman et al., 2020). Constructing estimators, confidence sets, and hypothesis tests for such non-i.i.d. data comes with added difficulty. Accordingly, and also for its importance in light of the reproducibility crisis (Baker, 2016), the task has attracted significant attention in the statistics community in recent years (Ramdas et al., 2022).

---

\*Equal contribution.




Figure 1: (a) and (b) show examples of confidence sets defined via level sets of the log-likelihood function in 2D at two dataset sizes, for Gaussian (a) and Laplace (b) likelihoods respectively. The sets inherit the geometry of the likelihood, and are not always ellipsoidal. (c) shows confidence bands on an RKHS function in a bandit game searching for the optimum. We compare prior work on confidence sets (Abbasi-Yadkori et al., 2011), our LR sets, and a common heuristic (orange). Our sets are nearly as small as the commonly used heuristic, but have provable coverage and can vastly improve sequential decision-making tasks such as bandits by quickly eliminating hypotheses.

Instead of deriving explicit concentration inequalities around an online estimator, we construct confidence sets *implicitly* defined by an inclusion criterion that is easy to evaluate in a computationally efficient manner and requires little statistical knowledge to implement. Roughly speaking, given a model  $p_\theta(y|x)$  that describes the conditional dependence of the observation  $y$  given the covariate  $x$  under parameter  $\theta$ , we will build sets based on a *weighted* modification of the sequential likelihood ratio statistic (Robbins et al., 1972; Wasserman et al., 2020)

$$R_t(\theta) := \frac{\mathcal{L}_t(\{\hat{\theta}_s\}_{s=1}^t)}{\mathcal{L}_t(\theta)} := \frac{\prod_{s=1}^t p_{\hat{\theta}_s}^{w_s}(y_s|x_s)}{\prod_{s=1}^t p_\theta^{w_s}(y_s|x_s)}, \quad (1)$$

where  $\{\hat{\theta}_s\}_s$  is a running estimator sequence that we are free to choose, but which may only depend on *previously* collected data. Parameters  $\theta$  for which this statistic is small, i.e., for which  $R_t(\theta) \leq 1/\alpha$  will be included in the set (and considered *plausible*). Examples of sets in a parametric and non-parametric setting are shown in Figure 1. The weighting terms  $w_s \in (0, 1]$  are crucial for dealing with inherent irregularities of many conditional observation models but can be flexibly chosen. Classically, these are set to  $w_s = 1$ . The full exposition of our method with choice of estimators and weights is given in Section 2. Apart from being easy to use and implement, our approach also comes with performance guarantees. These sets maintain a provable  $1 - \alpha$  coverage – a fact we establish using Ville’s inequality for supermartingales (Ville, 1939), which is known to be essentially tight for martingales (see Howard et al., 2018, for a discussion). Therefore, in stark contrast to alternate methods, our confidence sequence is *fully data-dependent*, making it empirically tighter than competing approaches. Despite the rich history of sequential testing and related confidence sets going back to Wald (1945) and Robbins et al. (1972), these sets have found little use in the interactive machine learning community, which is a gap we fill in the present paper.

**Contributions** In this work, we revisit the idea of using likelihood ratios to generate anytime-valid confidence sets. The main insight is that whenever the likelihood of the noise process is known, the likelihood ratio confidence sets are fully specified. They inherit their geometry from the likelihood function, and their size depends on the quality of our estimator sequence. We critically evaluate the likelihood ratio confidence sets and, in particular, we shed light on the following aspects: **Firstly**, for generalized linear models, we *theoretically* analyze the geometry of the LR confidence sets under mild assumptions. We show their geometry is dictated by Bregman divergences of exponential families (Chowdhury et al., 2022). **Secondly**, we show that the size of the confidence set is dictated by an online prediction game. The size of these sets depends on a sequence of estimators  $\{\hat{\theta}_s\}_{s=1}^t$  that one uses to estimate the unknown parameter  $\theta_*$ . We discuss how to pick the estimator sequence in order to yield a provably small radius of the sets, by using the Follow-the-Regularized-Leader algorithm, which implements a regularized maximum-likelihood estimator. We prove that the radius of the confidence sets is nearly-worst-case optimal, and accordingly, they yield nearly-worst-case regret bounds when used in generalized linear bandit applications. However, due to their data-dependent nature, they can be much tighter than this theory suggests. **Thirdly**, we analyze the limitations of classical (un-weighted) LR sets when the underlying conditional observation model is not identifiable. In this case, the resulting (inevitable) estimation bias unnecessarily increases the size of the confidence sets. To mitigate this, we propose an adaptive reweighting scheme that decreases the effect of uninformed early

bias of the estimator sequence on the size of the sets downstream. The reweighting does not affect the coverage guarantees of our sets and utilizes an elegant connection to (robust) powered likelihoods (Wasserman et al., 2020). **Finally**, thanks to the adaptive reweighting scheme, our sets are very practical as we showcase experimentally. We demonstrate that our method works well with exponential and non-exponential family likelihoods, and in parametric as well as in kernelized settings. We attribute their practical benefits to the fact that they *do not depend on (possibly loose) worst-case parameters*.

## 2 The Likelihood Method

The sequential likelihood ratio process (LRP) in (1) is a statistic that compares the likelihood of a given model parameter, with the performance of an adaptively chosen estimator sequence. As noted above, we generalize the traditional definition, which would have  $w_s \equiv 1$ , and define a corresponding confidence set as

$$\mathcal{C}_t = \{\theta \mid R_t(\theta) \leq 1/\alpha\}. \quad (2)$$

The rationale is that the better a parameter  $\theta$  is at explaining the data  $\{(x_s, y_s)\}_s^t$  from the true model  $\theta_*$ , the smaller this statistic will be, thereby increasing its chances to be included in  $\mathcal{C}_t$ . When we construct  $R_t$ , the sequence of  $x_s$ ,  $w_s$  and  $\hat{\theta}_s$  cannot depend on the noisy observation  $y_s$ . Formally, consider the filtration  $(\mathcal{F}_s)_{s=0}^\infty$  with sub- $\sigma$ -algebras  $\mathcal{F}_s = \sigma(x_1, \dots, y_1, \dots, x_s, y_s, x_{s+1})$ . We require that  $\hat{\theta}_s$  and  $w_s$  are  $\mathcal{F}_{s-1}$ -measurable. Under these very mild assumptions and with arbitrary weights  $w_s \in (0, 1]$ , we can show coverage, i.e., our (weighted) confidence sets uniformly track the true parameter with probability  $1 - \alpha$ .

**Theorem 1.** *The stochastic process  $R_t(\theta_*)$  in (1) is a non-negative supermartingale with respect to the filtration  $(\mathcal{F}_t)$  and satisfies  $R_0(\theta_*) \equiv 1$ . In addition, the sequence  $\mathcal{C}_t$  from (2) satisfies  $\mathbb{P}_{\theta_*}(\exists t : \theta_* \notin \mathcal{C}_t) \leq \alpha$ .*

The last statement follows by applying Ville’s inequality for super-martingales on  $R_t(\theta_*)$ . The proof closely follows Wasserman et al. (2020). While coverage is always guaranteed irrespective of the estimator sequence  $\{\hat{\theta}_s\}$ , we would like to make the sets as small as possible at fixed coverage, which we do by picking a well-predicting estimator sequence.

### 2.1 The Estimator Sequence Game

The specification of the LR process (LRP) allows us to choose an arbitrary estimator sequence  $\{\hat{\theta}_s\}_s$ . To understand the importance of the sequence, let us introduce  $\theta_*$  to the definition of  $R_t$  in (1), and divide by  $\mathcal{L}_t(\{\hat{\theta}_s\}_{s=1}^t)$ . This gives the equivalent formulation

$$\mathcal{C}_t := \left\{ \theta \mid \frac{\mathcal{L}_t(\theta_*)}{\mathcal{L}_t(\theta)} \leq \frac{1}{\alpha} \frac{\mathcal{L}_t(\theta_*)}{\mathcal{L}_t(\{\hat{\theta}_s\}_{s=1}^t)} \leftarrow \text{confidence parameter} \right\}.$$

We see that the predictor sequence does not influence the geometry of the confidence set, which is fully specified by the likelihood function. We also observe that the ratio on the right-hand side serves as a confidence parameter controlling the size (radius) of the confidence sets measured under the likelihood ratio distance to  $\theta_*$ . If the confidence parameter goes to zero, only  $\theta_*$  is in the set. The better the estimator sequence is at predicting the data, the smaller the inclusion threshold, and hence the smaller the sets will ultimately be. Specifically, taking the log, we would like to *minimize*

$$\mathcal{R}_t := \log \frac{\mathcal{L}_t(\theta_*)}{\mathcal{L}_t(\{\hat{\theta}_s\}_{s=1}^t)} = \sum_{s=1}^t -\log(p_{\hat{\theta}_s}^{w_s}(y_s \mid x_s)) - \sum_{s=1}^t -\log(p_{\theta_*}^{w_s}(y_s \mid x_s)). \quad (3)$$

The quantity  $\mathcal{R}_t$  corresponds to a regret in an online prediction game, as will become apparent below.

**Online Prediction Game** Online optimization is a mature field in interactive learning (Cesa-Bianchi and Lugosi, 2006; Orabona, 2019). The general goal is to minimize a sequence of loss functions as in Eq. (3) and compete against a baseline, which typically is the best-in-hindsight prediction, or – in our case – given by the performance of the fixed parameter  $\theta_*$ . Specifically, at every timestep  $s$ , iteratively, the agent chooses an action  $\hat{\theta}_s$  based on  $\mathcal{F}_{s-1}$ , and a loss function  $f_s(\theta)$  is revealed. In most of the online optimization literature,  $f_s$  can be chosen adversarially. In our prediction game, we know the whole form of loss function  $f_s(\theta) = -\log(p_{\hat{\theta}_s}^{w_s}(y_s \mid x_s))$ , as can be seen in (3), and not just  $f_s(\hat{\theta}_s)$ . Opposed to traditional assumptions in online prediction, in our case,  $f_s$  are non-adversarial, but have a stochastic component due to  $y_s$ . Also, contrary to most instances of online prediction, we do not compare against the best-in-hindsight predictor, but  $\theta_*$  instead, as this is more meaningful in our setting.

**Online Optimization Algorithms** Generally, we seek an algorithm that incurs low regret. Here, we focus on *Follow-the-Regularized Leader (FTRL)*, which corresponds exactly to using regularized maximum likelihood estimation, making it a natural and computationally practical choice. The update rule is defined in Alg. 1 (Line 3). While other algorithms could be considered, FTRL enjoys the optimal regret rate for generalized linear regression as we show later, and is easily implemented. In order to run the algorithm, one requires a sequence of strongly convex regularizers. For now, let us think of it as  $\psi_s(\theta) = \lambda \|\theta\|_2^2$ , which we use in practice. However, one can derive a tighter analysis for a slightly modified, time-dependent regularization strategy for generalized linear models as we show in Sec. 3.3.

## 2.2 Adaptive Reweighting: Choosing the Right Loss

There is yet more freedom in the construction of the LR, via the selection of the *loss function*. Not only do we select the predictor sequence, but also the *weights* of the losses via  $w_t$ . This idea allows controlling the influence of a particular data point  $(x_t, y_t)$  on the cumulative loss based on the value of  $x_t$ . For example, if we know a priori that for a given  $x_t$  our prediction will be most likely bad, we can opt out of using the pair  $(x_t, y_t)$  by setting  $w_t = 0$ . Below we will propose a weighting scheme that depends on a notion of *bias*, which captures how much of the error in predicting  $y_t$  is due to our uncertainty about  $\hat{\theta}_t$  (compared to the uncertainty we still would have *knowing*  $\theta_*$ ). Sometimes this *bias* is referred to as *epistemic* uncertainty in the literature, while the residual part of the error is referred to as *aleatoric*. Putting large weight on a data point heavily affected by this bias might unnecessarily increase the regret of our learner (and hence blow up the size of the confidence set). Note that, conveniently, even if we put low weight (zero) on a data point, nothing stops us from using this sample point to improve the estimator sequence in the next prediction round. As we will show below, our reweighting scheme is crucial in defining a practical algorithm for Reproducing Kernel Hilbert Space (RKHS) models and in high signal-to-noise ratio scenarios. Since we do not know  $\theta_*$ , our strategy is to compute an *estimate of the bias of the estimator*  $\hat{\theta}_t$  and its effect on the value of the likelihood function for a specific  $x$  that we played. We use the value of the bias to rescale the loss via  $w_t$  such that its effect is of the same magnitude as the statistical error (see Algorithm 1; we call this step BIAS-WEIGHTING).

**Intuition** To give a bit more intuition, suppose we have a Gaussian likelihood. Then the negative log-likelihood of  $(x_t, y_t)$  with weighting is proportional to  $\frac{w_t}{\sigma^2} (y_t - x_t^\top \hat{\theta}_t)^2$ . Now, if  $x_t$  does not lie in the span of the data points  $\{x_s\}_{s=1}^{t-1}$  used to compute  $\hat{\theta}_t$ , it is in general unavoidable to incur large error, inversely proportional to  $\sigma^2$ . To see this, let us decompose the projection onto  $x_t$  as

$$x_t^\top (\hat{\theta}_t - \theta_*) = \underbrace{x_t^\top (\hat{\theta}_t - \mathbb{E}[\hat{\theta}_t])}_{\text{statistical error}} + \underbrace{x_t^\top (\mathbb{E}[\hat{\theta}_t] - \theta_*)}_{\text{bias}_{x_t}(\hat{\theta}_t)}, \quad (4)$$

where the first term represents the statistical error up to time  $t$ , while the second, bias, is deterministic, and independent of the actual realization  $y$ , depending only  $\theta_*$ . Estimators with non-zero bias are *biased*. Plugging this into the likelihood function, we see that in expectation  $\frac{1}{\sigma^2} \mathbb{E}[(y_t - x_t^\top \hat{\theta}_t)^2 | \mathcal{F}_{t-1}] \lesssim \frac{1}{\sigma^2} \text{bias}_{x_t}^2(\hat{\theta}_t) + \epsilon^2 + \frac{C}{t}$ , where  $\epsilon^2$  is the unavoidable predictive error in expectation (due to a noisy objective) and is a constant independent of  $\sigma^2$ .  $\frac{C}{t}$  is the statistical error, and  $C$  is independent of  $\sigma^2$ . Note that the bias term scales inversely with the variance, and leads to unnecessarily big confidence parameters for small  $\sigma^2$ .

In fact, the problem is that we use the likelihood to measure the distance between two parameters, but this is only a “good” distance once the deterministic source of the error (bias) vanishes. For this reason, without weighting, the incurred regret blows up severely in low-noise settings. To counter this, we balance the deterministic estimation bias and noise variance via proper selection of  $w_t$ . In this case, it turns out that  $w_t = \frac{\sigma^2}{\sigma^2 + \text{bias}_{x_t}^2(\hat{\theta}_t)}$  ensures that the overall the scaling is independent of  $\sigma^2$ . While the choice of weights  $\{w_s\}_s^t$  influences the geometry of the confidence sets, with a *good data collection and estimation strategy* the bias asymptotically decreases to zero, and hence the weights converge to 1.

**Bias estimation** In order to generalize this rule beyond Gaussian likelihoods, we need a proper generalization of the bias. Our generalization is motivated by our analysis of generalized linear models, but the method can be applied more broadly. The role of the squared statistical error (variance) is played by the inverse of the smoothness constant of the negative log-likelihood functions  $f_s$ , denoted by  $L$ . This is the usual smoothness, commonly seen in the convex optimization literature.

We consider penalized likelihood estimators with strongly convex regularizers (Alg. 1, line 3). For this estimator class, we define the bias via a hypothetical stochastic-error-free estimate  $\hat{\theta}_t^\times$ , had we access to the expected values of the gradient loss functions (a.k.a. score). We use the first-order optimality conditions and the indicator function of the set  $\Theta$ ,  $i_\Theta$ , to define the error-free-estimate  $\hat{\theta}_t^\times$ , and the bias of the estimator  $\hat{\theta}_t$  as

$$\text{bias}_{x_t}^2(\hat{\theta}_t) = (x_t^\top (\theta_\star - \hat{\theta}_t^\times))^2 \quad \text{with} \quad \mathbb{E} \left[ \sum_{s=1}^{t-1} \nabla \log p_{\hat{\theta}_t^\times}(y_s | x_s) \right] - \nabla \psi_t(\hat{\theta}_t^\times) + i_\Theta(\hat{\theta}_t^\times) = 0, \quad (5)$$

where the expectation denotes a sequence of expectations conditioned on the prior filtration. This notion of bias coincides with the definition of bias in Eq. (4) for the Gaussian likelihood. This quantity cannot be evaluated in general, however, we prove a computable upper bound.

**Theorem 2** (Bias estimate). *Let the negative log-likelihood have the form,  $-\log p_\theta(y_s | x_s) = g(x_s^\top \theta)$ , where  $g : \mathbb{R} \rightarrow \mathbb{R}$  is  $\mu$  strongly-convex and let the regularizer be  $\psi_t(\theta) = \lambda \|\theta\|_2^2$  making the overall objective strongly convex. Then, defining  $\mathbf{V}_t^{\mu;\lambda} = \sum_{s=1}^t \mu x_s x_s^\top + \lambda \mathbf{I}$ , we can bound*

$$\text{bias}_x^2(\hat{\theta}_t) \leq 2\lambda \|\theta_\star\|_2^2 x^\top (\mathbf{V}_t^{\mu;\lambda})^{-1} x. \quad (6)$$

The proof is deferred to App. A.3, and requires elementary convex analysis.

This leads us to propose the weighting scheme  $w_t = \frac{1/L}{\text{bias}_{x_t}^2(\hat{\theta}_t) + 1/L}$ . We justify that this is a sensible choice by analyzing the confidence set on the GLM class in Section 3, which satisfies the smoothness and strong-convexity conditions. We show that this rule properly balances the stochastic and bias components of the error in the regret as in (3). However, this rule is more broadly applicable beyond the canonical representation of GLM or the GLM family altogether.

---

**Algorithm 1** Constructing the LR Confidence Sequence

---

- 1: **Input:** convex set  $\Theta \subset \mathbb{R}^d$ , confidence level  $\alpha > 0$ , likelihood  $p_\theta(y|x)$ , regularizers  $\{\psi_t\}_t$
  - 2: **for**  $t \in \mathbb{N}_0$  **do**
  - 3:    $\hat{\theta}_t = \arg \min_{\theta \in \Theta} \sum_{s=1}^{t-1} -\log p_\theta(y_s | x_s) + \psi_t(\theta)$  ▷ FTRL
  - 4:    $w_t = \begin{cases} \frac{1/L}{1/L + \text{bias}_{x_t}^2(\hat{\theta}_t)} & \text{THIS WORK} \\ 1 & \text{CLASSICAL} \end{cases}$  ▷ BIAS-WEIGHTING  $\text{bias}_{x_t}(\hat{\theta}_t)$  in Eq. (5) or Eq.(6)
  - 5:    $\mathcal{C}_t = \left\{ \theta \in \Theta \mid \prod_{s=1}^t \frac{p_{\hat{\theta}_s}^{w_s}(y_s | x_s)}{p_\theta^{w_s}(y_s | x_s)} \leq \frac{1}{\alpha} \right\}$ . ▷ Confidence set
  - 6: **end for**
- 

### 3 Theory: Linear Models

While the *coverage* (i.e., “correctness”) of the likelihood ratio confidence sets is always guaranteed, their worst-case *size* (affecting the “performance”) cannot be easily bounded in general. We analyze the size and the geometry of the LR confidence sequence in the special but versatile case of generalized linear models.

#### 3.1 Generalized Linear Models

We assume knowledge of the conditional probability model  $p_\theta(y|x)$ , where the covariates  $x \in \mathcal{X} \subset \mathbb{R}^d$ , and the true underlying model parameter lies in a set  $\Theta \subset \mathbb{R}^d$ . If  $t$  is indexing (discrete) time, then  $x_t$  is acquired sequentially, and the – subsequently observed –  $y_t$  is sampled from an exponential family distribution parametrized as

$$p_\theta(y | x_t) = h(y) \exp (T(y) \cdot x_t^\top \theta - A(x_t^\top \theta)). \quad (7)$$

Here,  $A$  is referred to as the *log-partition function* of the conditional distribution, and  $T(y)$  is the sufficient statistic. The function  $h$  is the base measure, and has little effect on our further developments, as it cancels out in the LR. Examples of commonly used exponential families (Gaussian, Binomial, Poisson or Weibull) with their link functions can be found in Table 1 in App. A.1.

In order to facilitate theoretical analysis for online algorithms, we make the following assumptions about the covariates  $x \in \mathcal{X}$  and the set of plausible parameters  $\Theta$ .



**Assumption 1.** *The covariates are bounded, i.e.,  $\sup_{x \in \mathcal{X}} \|x\|_2 \leq 1$ , and the set  $\Theta$  is contained in an  $\ell_2$ -ball of radius  $B$ . We will also assume that the log-partition function is strongly convex, that is, that there exists  $\mu := \inf_{z \in [-B, B]} A''(z)$ , and that  $A$  is  $L$ -smooth, i.e.  $L := \sup_{z \in [-B, B]} A''(z)$ .*

These assumptions are common in other works addressing the confidence sets of GLMs (Filippi et al., 2010; Fauray et al., 2020), who remark that the dependence on  $\mu$  is undesirable. However, in contrast to these works, our confidence sets *do not* use these assumptions in the construction of the sets. We only require these for our theoretical analysis. As these are worst-case parameters, the practical performance can be much better for our sets.

### 3.2 Geometry and Concentration

Before stating our results, we need to define a distance notion that the convex negative log-likelihoods induce. For a continuously differentiable convex function  $f$ , we denote the Bregman divergence as  $D_f(a, b) := f(a) - f(b) - \nabla f(b)^\top (a - b)$ . The  $\nu$ -regularized sum of log-partition functions is defined as

$$Z_t^\nu(\theta) := \sum_{s=1}^t w_s A(x_s^\top \theta) + \frac{\nu}{2} \|\theta\|_2^2. \quad (8)$$

This function will capture the geometry of the LR confidence sets. The confidence set size depends mainly on two terms. One refers to a notion of complexity of the space referred to as *Bregman information gain*:  $\Gamma_t^\nu(\tilde{\theta}_t) = \log \left( \frac{\int_{\mathbb{R}^d} \exp(-\frac{\nu}{2} \|\theta\|_2^2) d\theta}{\int_{\mathbb{R}^d} \exp(-D_{Z_t^\nu}(\theta, \tilde{\theta}_t)) d\theta} \right)$ , first defined by Chowdhury et al. (2022) as a generalization of the *information gain* of Srinivas et al. (2009),  $\gamma_t^\nu = \log \left( \det(\sum_{i=1}^t \frac{\mu}{\nu} x_i x_i^\top + \mathbf{I}) \right)$  for Gaussian likelihoods. We will drop the superscript whenever the regularization is clear from context and simply refer to  $\gamma_t$ . This term appears because one can relate the decay of the likelihood as a function of the Bregman Divergence from  $\theta_*$  with the performance of a (regularized) maximum likelihood estimator via convex (Fenchel) duality. In particular, if  $\tilde{\theta}_t$  is a regularized MLE,  $\Gamma_t^\nu := \Gamma_t^\nu(\tilde{\theta}_t)$  will asymptotically scale as  $\mathcal{O}(d \log t)$  (cf. Chowdhury et al., 2022, for further discussion). For Gaussian likelihoods and  $w_s \equiv 1$ , it coincides with the classical information gain independent of  $\tilde{\theta}_t$ . The second term that affects the size is the regret  $\mathcal{R}_t$  of the online prediction game over  $t$  rounds we introduced previously in (3). These two parts together yield the following result:

**Theorem 3.** *Let  $\nu > 0$  and  $\alpha, \delta \in (0, 1)$ . For the level  $1 - \alpha$  confidence set  $\mathcal{C}_t$  defined in (2) under the GLM in (7), with probability  $1 - \delta$ , for all  $t \geq 1$ , any  $\theta \in \mathcal{C}_t$  satisfies*

$$D_{Z_t^\nu}(\theta, \theta_*) \leq \frac{4L}{\mu} \xi_t + 2 \log \left( \frac{1}{\delta} \right) + 2\mathcal{R}_t, \quad (9)$$

where  $\xi_t = \left( \log \left( \frac{1}{\alpha} \right) + \nu B^2 + \Gamma_t^\nu \right)$  and  $L, \mu$  are defined as above and finally  $\mathcal{R}_t$  is the regret of the game in Eq. (3).

The set defined via the above divergence does not coincide with the LR confidence set. It is slightly larger due to a term involving  $\nu$  (as in Eq. (8)). This is a technical consequence of our proof technique, where the gradient of  $Z_t^\nu$  needs to be invertible, and regularization is added to this end. We note that this  $\nu > 0$  can be chosen freely. Note that the theorem involves two confidence levels,  $\alpha$  and  $\delta$ :  $\alpha$  is a bound on the Type I error – coverage of the confidence sets – while  $\delta$  upper bounds the probability of a large radius – and is therefore related to the power and Type II error of a corresponding hypothesis test. The proof of the theorem is deferred to App. B.2.

To give more intuition on these quantities, let us instantiate them for the Gaussian likelihood case with  $w_s \equiv 1$ . In this scenario,  $Z_t^\nu(\theta) = \sum_{s=1}^t \frac{1}{2\sigma^2} \|\theta\|_{x_s x_s^\top}^2 + \frac{\nu}{2} \|\theta\|_2^2$ , and the (in this case symmetric) Bregman divergence is equal to  $D_{Z_t^\nu}(\theta, \theta_*) = \frac{1}{2} \|\theta - \theta_*\|_{\mathbf{V}_t^{\mu, \nu}}^2$ , where  $\mathbf{V}_t^{\mu, \nu} = \sum_{s=1}^t \mu x_s x_s^\top + \nu \mathbf{I}$ , which means that our confidence sets are upper bounded by a ball in the same norm as those in the seminal work on linear bandits (Abbasi-Yadkori et al., 2011).

### 3.3 Online Optimization in GLMs: Follow the Regularized Leader

The size of the confidence sets in Theorem 3 depends on the regret of the online prediction game involving the estimator sequence. We now bound this regret when using the Follow-the-Regularized-Leader (FTRL) algorithm in this setting. This high probability bound is novel to the best of our

knowledge and may be of independent interest. We state in a weight-agnostic manner first, and then with our particular choice. The latter variant uses a specifically chosen regularizer. In this case, we can track the contribution of each time-step towards the regret separately.

**Theorem 4.** *Let  $\psi_t(\theta) = \lambda\|\theta\|_2^2$ . Assume Assumption 1, and additionally that  $A$  is  $L$ -smooth everywhere in  $\mathbb{R}^d$ , and let  $w_t \in [0, 1]$  be arbitrary. Then, with probability  $1 - \delta$  the regret of FTRL (Alg. 1) satisfies for all  $t \geq 1$*

$$\mathcal{R}_t \leq \lambda B^2 + \frac{L}{\mu}(\gamma_t^\lambda + 2\log(1/\delta)) + \frac{2L^2 B^2}{\mu} \gamma_t^\lambda. \quad (10)$$

The regret bounds are optimal in the orders of  $\gamma_t^\lambda$ , matching lower bounds of [Ouhamma et al. \(2021\)](#), as for linear models  $\gamma_t = \mathcal{O}(d \log t)$ . Combining results of Thm. 4 with Thm. 3, we get a confidence parameter that scales with  $\mathcal{O}(\sqrt{\gamma_t})$ , for confidence sets of the form  $\|\theta - \theta_\star\|_{\mathbf{V}_t}$ , which coincides with the best-known confidence sets in this setting in the worst-case ([Abbasi-Yadkori et al., 2012](#)). The requirement of global  $L$ -smoothness can be relaxed to  $L$ -smoothness over  $\Theta$ . With a more elaborate (but less insightful) analysis, we can show that we achieve a  $\tilde{\mathcal{O}}(\gamma_t)$  bound even in this case. The proofs of these results are deferred to App. C.4, App. C.5 and App. C.6 respectively.

**Regret, Weighting and Estimation Bias** Interestingly, the term in Thm. 4 involving the (crude) proxy to the bias – the bound  $B$  – is not scaled by the same  $L/\mu$  factors as the other terms in the regret bound (10) and in Theorem 3. Namely, the prefactor is  $L^2/\mu$  instead of  $L/\mu$ . This extra dependence manifests itself in the unnecessary penalization through the estimation bias we introduced in Sec. 2.2, particularly in low-noise settings. We addressed this issue by picking the weights  $\{w_t\}$ . While the above theorem holds for any valid weighting, it does not exhibit the possible improvement from using specific weights.

We argued earlier that the error in prediction should not be measured by the likelihood function if there is deterministic error, since initially, we are fully uncertain about the value of  $\theta_\star^\top(\cdot)$  outside the span of previous observations. Of course, if our goal would be to purely pick weights to minimize  $\mathcal{R}_t$ , then  $w_s = 0$  would lead to zero regret and hence be optimal. However, the likelihood ratio would then be constant, and uninformative. In other words, the associated log-partition Bregman divergence in Theorem 3 would be trivial and not filter out any hypotheses. Clearly, some balance has to be met. With this motivation in mind, we proposed a *nonzero* weighting that decreases the regret contribution of the bias, namely  $w_t = \frac{1/L}{1/L + \text{bias}_{x_t}^2(\hat{\theta}_t)}$ . The advantage of this choice becomes more apparent when

we use the regularizer  $\psi_t(\theta) = \lambda\|\theta\|^2 + A(x_t^\top \theta)$  to obtain the following result.

**Theorem 5.** *Let  $\psi_s(\theta) = \lambda\|\theta\|^2 + A(x_s^\top \theta)$ . Assume Assumption 1, and additionally that  $A$  is  $L$ -smooth everywhere in  $\mathbb{R}^d$ , and choose  $w_s = \frac{1/L}{1/L + \text{bias}_{x_s}^2(\hat{\theta}_s)}$ . Additionally, let the sequence of  $x_s$  be such that,  $\sum_s (1 - w_s)(f_s(\theta_\star) - f_s(\bar{\theta}_{s+1})) \leq L/\mu\gamma_t^\lambda$ , where  $\bar{\theta}_s$  is the FTRL optimizer with the regularizer  $\lambda\|\theta\|_2^2$  from Theorem 4<sup>2</sup>. Then, with probability  $1 - \delta$  the regret of FTRL (Alg. 1) satisfies for all  $t \geq 1$*

$$\mathcal{R}_t \leq \lambda B^2 + \frac{2L}{\mu} \left( \gamma_t^\lambda + \log\left(\frac{1}{\delta}\right) \right) + \frac{L}{\mu} \sum_{s=1}^t \frac{B^2}{1/L + \text{bias}_{x_s}^2(\hat{\theta}_s)} \Delta\gamma_s^\lambda,$$

where  $\Delta\gamma_s^\lambda = \gamma_{s+1}^\lambda - \gamma_s^\lambda$ .

One can see that for points where the information gain  $\Delta\gamma_s$  is large (corresponding to more unexplored regions of the space, where the deterministic source of error is then large), the weighting scheme will make sure that the multiplicative contribution of  $B^2$  is mitigated, along with having the correct prefactor  $L/\mu$ . The reader may wonder how this result is useful when we replace  $\text{bias}_{x_s}^2(\hat{\theta}_s)$  with the upper bound from Thm. 2. While instructive, our bound still only makes the bias proxy  $B^2$  appear in front of the information gain  $\Delta\gamma_t$ , instead of the more desirable bias itself. In the latter case, we could also directly make use of the upper bound and get an explicit result only using an upper bound on the bias. We leave this for future work.

We point out that this choice of  $\psi_s(\theta)$  in Theorem 5 corresponds to the Vovk-Azoury-Warmuth predictor ([Vovk, 2001](#); [Azoury and Warmuth, 1999](#)) in the online learning literature. This choice is helpful in order to track the bias contribution more precisely in our proof.

<sup>2</sup>Note that this assumption was missing in an earlier version.  $\bar{\theta}_{s+1}$  corresponds to a regularized MLE that did observe the data pair  $(x_s, y_s)$ .

## 4 Application: Linear and Kernelized Bandits

Our main motivation to construct confidence sets is bandit optimization. A prototypical bandit algorithm – the Upper Confidence Bound (UCB) (Auer, 2002) – sequentially chooses covariates  $x_s$  in order to maximize the reward  $\sum_{s=1}^t r_{\theta_*}(x_s)$ , where  $r_{\theta_*}$  is the unknown pay-off function parametrized by  $\theta_*$ . UCB chooses the action  $x_s$  which maximizes the optimistic estimate of the reward in each round, namely

$$x_s = \arg \max_{x \in \mathcal{X}} \max_{\theta \in \mathcal{C}_{s-1}} r_{\theta}(x), \quad (11)$$

where  $\mathcal{C}_{s-1}$  is some confidence set for  $\theta_*$ , and can be constructed with Algorithm 1 from the first  $s - 1$  data points. An important special case is when  $r_{\theta_*}$  is linear (Abe and Long, 1999) or modelled by a generalized linear model (Filippi et al., 2010). In that case, the inner optimization problem is convex as long as  $\mathcal{C}_{s-1}$  is convex. The outer optimization is tractable for finite  $\mathcal{X}$ . In the applications we consider, our confidence sets are convex, and we easily solve the UCB oracle using convex optimization toolboxes.

**Extension to RKHS** We introduced the framework of LR confidence sets only for finite-dimensional Euclidean spaces. However, it can be easily extended to Reproducing Kernel Hilbert Spaces (RKHS) (Cucker and Smale, 2002). The definition of the LR process in (1) is still well-posed, but now the sets are subsets of the RKHS, containing functions  $f \in \mathcal{H}_k$ . An outstanding issue is how to use these sets in downstream applications, and represent them tractably as in Figure 1. Conveniently, even with infinite-dimensional RKHSs, the inner-optimization in (11) admits a Lagrangian formulation, and the generalized representer theorem applies (Schölkopf et al., 2001; Mutný and Krause, 2021). In other words, we can still derive a pointwise upper confidence band as  $\text{ucb}(x) = \max_{f \in \mathcal{H}_k, \|f\|_k \leq B, f \in \mathcal{C}_s} \langle f, k(x, \cdot) \rangle$  in terms of  $\{x_j\}_{j=1}^s \cup \{x\}$ , leading to a  $s + 1$ -dimensional, tractable optimization problem.

We also point out that the weighting is even more paramount in the RKHS setting, as the bias never vanishes for many infinite dimensional Hilbert spaces (Mutný and Krause, 2022). For this purpose, our weighting is of paramount practical importance, as we can see in Figure 2a), where the gray arrow represents the significant improvement from reweighting.

### 4.1 Instantiation of the Theory for Linear Bandits

Before going to the experiments, we instantiate our theoretical results from Sec. 3 to the important and well-studied special case of linear payoffs. In that case,  $r_{\theta}(x) = \langle x, \theta \rangle$  and the agent observes  $y_s = \langle x_s, \theta_* \rangle + \eta_s$  upon playing action  $x_s$ , where  $\eta_s \sim \mathcal{N}(0, \sigma^2)$ . We are interested in minimizing the so-called cumulative pseudo-regret, namely,  $\mathfrak{R}_t = \sum_{s=1}^t [\langle x_*, \theta_* \rangle - \langle x_s, \theta_* \rangle]$ , where  $x_*$  refers to the optimal action. Using the set from (2) along with Theorem 3 and the FTRL result of Theorem 4 we can get a regret bound for the choice  $w_s \equiv 1$ .

**Theorem 6.** *Let  $w_s \equiv 1$ . For any  $\lambda \geq \frac{1}{\sigma^2}$ , with probability at least  $1 - 3\delta$ , for all  $t \in \mathbb{N}$  we have*

$$\mathfrak{R}_t \leq 6\sqrt{t\gamma_t^\lambda} \left( \sigma \sqrt{\log(1/\delta) + \gamma_t^\lambda + \sigma\lambda^{1/2}B} + B\sqrt{\gamma_t^\lambda} \right).$$

Our results are optimal in both  $d$  and  $t$  up to constant and logarithmic factors. The proof is deferred to App. D, but is an instantiation of the aforementioned theorems, along with a standard analysis. There, we also compare to the seminal result of Abbasi-Yadkori et al. (2011), which does not suffer from the dependence on  $B\sqrt{\gamma_t}$ . We attribute this to the incurred bias in the absence of the reweighting scheme.

For the weighted likelihood ratio, we can obtain a result similar to the above, but multiplied by an upper bound on  $\sup_{s \geq 1} w_s^{-1}$ . This is undesirable, as our experiments will show that the reweighting scheme vastly improves performance. While this could be somewhat mitigated by using the Theorem 5 instead of Theorem 4 to bound the FTRL regret, a better result should be achievable using our weighting scheme that improves upon Theorem 6 and possibly even matches Abbasi-Yadkori et al. (2011) exactly in the worst-case. We leave this for future work.

### 4.2 Experimental Evaluation

In this subsection, we demonstrate that the practical applicability goes well beyond the Gaussian theoretical result from the previous subsection. In the examples below, we always use the UCB



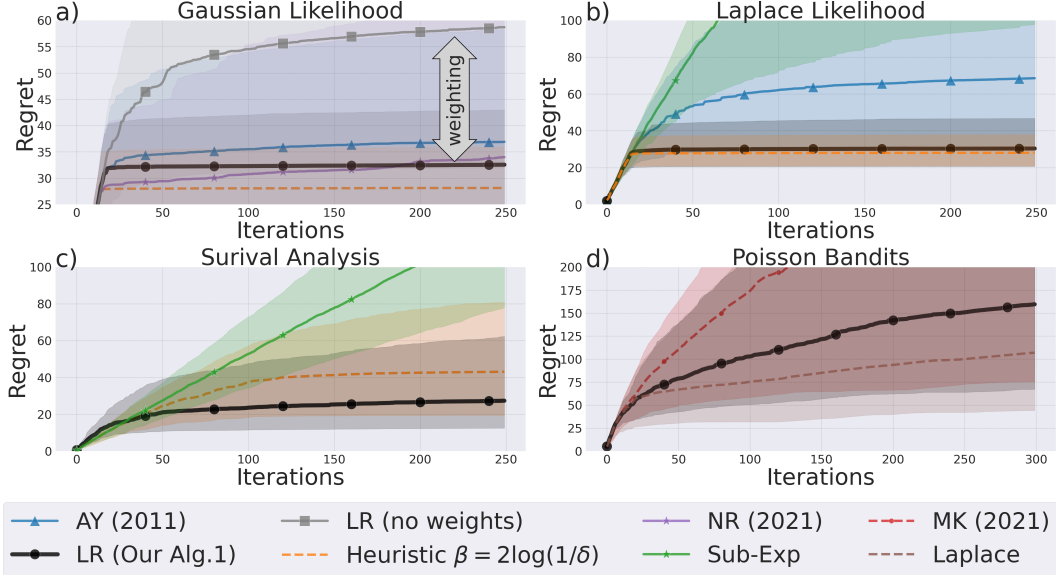


Figure 2: Bandit experiments: On the  $y$ -axis we report cumulative regret, while the  $x$ -axis shows the number of iterations. In a) and b) we report the results for linear models with different parametric additive noise. In c) we report the results on a survival analysis with a log-Weibull distribution ( $p = 2$ ) and in d) we showcase Poisson bandits. See App. E for more details. Heuristic methods are *dashed*, while provable are *solid*. Our sets perform the best among all provable methods. Notice in a) the difference in gray and black represents the improvement due to adaptive weighting over  $w_s = 1$  for all  $s \in [t]$ . For each experiment we did 10 reruns, median values are plotted.

algorithm but employ different confidence sets. In particular, we compare our LR confidence sets for different likelihood families with alternatives from the literature, notably classical sub-family confidence sets (Abbasi-Yadkori et al., 2011; Mutný and Krause, 2021), and the robust confidence set of Neiswanger and Ramdas (2021). In practice, however, the radius of these confidence sets is often tuned heuristically. We include such sets as a baseline *without* provable coverage as well. The main take-home message from the experiments is that among all the estimators and confidence sets that enjoy *provable* coverage, our confidence sets perform the best, on par with successful heuristics. For all our numerical experiments in Figure 2, the true payoff function is assumed to be an infinite dimensional RKHS element. For further details and experiments, please refer to App. E.

**Additive Noise Models** Suppose that  $r_{\theta_*}$  is linear and we observe  $y_s = x_s^\top \theta_* + \eta_s$ , where  $\eta_s$  is additive noise, and  $\theta_*$  is an element of a Hilbert space. We consider classical Gaussian noise as well as Laplace noise in Fig. 2[a), b)]. Notice that in both cases our confidence sets yield lower regret than any other provably valid method. In both cases they are performing as good as *heuristic* confidence sets with confidence parameter  $\beta_t \equiv 2\log(1/\delta)$ . The sub-Gaussian confidence sets of Abbasi-Yadkori et al. (2011) (AY 2011) are invalid for the Laplace distribution as it is not sub-Gaussian but only sub-Exponential. For this reason, we compare also with sub-exponential confidence sets derived similarly to those of (Fauray et al., 2020). The confidence sets of (Neiswanger and Ramdas, 2021) (NR 2021) perform similarly on Gaussian likelihood, but are only applicable to this setting, as their generalization to other likelihood families involves intractable posterior inference. We note also the difference between the unweighted LR and the weighted one. The examples in Fig. 2 use the true payoff functions  $r(x) = -(1.4 - 3x)\sin(18x)$ , which we model as an element of a RKHS with squared exponential kernel lengthscale  $\gamma = 6 \times 10^{-2}$  on  $[0, 1.2]$ , which is the baseline function no. 4 in the global optimization benchmark database *infinity77* (Gavana, 2021). Additional experiments can be found in App. E.

**Poisson Bandits** A prominent example of generalized linear bandits (GLB) are Poisson bandits, where the linear payoff is scaled by an exponential function. We instantiate our results on a common benchmark problem, and report the results in Fig. 2d). We improve the regret of UCB for GLBs compared to two alternative confidence sets: one that uses a Laplace approximation with a heuristic confidence parameter, and one inspired by considerations in Mutný and Krause (2021) (MK 2021), also with a heuristic confidence parameter. Note that we cannot compare directly to their provable results in their original form as they do not state them in the canonical form of the exponential family.

**Survival Analysis** Survival analysis is a branch of statistics with a rich history that models the lifespan of a service or product (Breslow, 1975; Cox, 1997; Kleinbaum and Klein, 2010). The classical approach postulates a well-informed likelihood model. Here, we use a specific hazard model, where the survival time  $T$  is distributed with a Weibull distribution, parametrized by  $\lambda$  and  $p$ . The rate  $\lambda_\theta(x) = \exp(x^\top \theta)$  differs for each configuration  $x$ , and  $p$  – which defines the shape of the survival distribution – is fixed and known. We assume that the unknown part is due to the parameter  $\theta$  which is the quantity we build a confidence set around to use within the UCB Algorithm. In particular, the probability density of the Weibull distribution is  $P(T = t|x) = \lambda_\theta(x) p t^{p-1} \exp(-t^p \lambda_\theta(x))$ . In fact, with  $p = 2$ , the confidence sets are convex and the UCB rule can be implemented efficiently.

Interestingly, this model admits an alternate linear regression formulation. Namely upon using the transformation  $Y = \log T$ , the transformed variables  $Y|x$  follow a Gumbel-type distribution, with the following likelihood that can be obtained by the change of variables  $P(Y = y|x) = \lambda_\theta(x) p \exp(y)^p \exp(-\exp(y)^p \lambda_\theta(x))$ . The expectation over  $Y$  allows us to express it as a linear regression problem since  $\mathbb{E}[Y|x] = -(\theta^\top x + \gamma)/p$ , where  $\gamma$  is the Euler-Mascheroni constant. More importantly,  $Y|x$  is sub-exponential. Hence, this allows us to use confidence sets for sub-exponential variables constructed with the pseudo-maximization technique inspired by Faury et al. (2020). More details on how these sets are derived can be found in App. E. However, these approaches necessarily employ crude worst-case bounds and as can be seen in Figure 2c) the use of our LR-based confidence sequences substantially reduces the regret of the bandit learner.

## 5 Related Work and Conclusion

**Related Work** The adaptive confidence sequences stem from the seminal work of Robbins et al. (1972), who note that these sets have  $\alpha$ -bounded Type I error. The likelihood ratio framework has been recently popularized by Wasserman et al. (2020) for likelihood families without known test statistics under the name *universal inference*. This approach, although long established, is surprisingly uncommon in sequential decision-making tasks like bandits. This might be due to the absence of an analysis deriving the size of the confidence sets (Mutný and Krause, 2021), a necessary ingredient to obtain regret bounds. We address this gap for generalized linear models. Another reason might be that practitioners might be interested in non-parametric *sub*-families – a scenario our method does not cover. That being said, many fields such as survival analysis (Cox, 1997) *do* have well-informed likelihoods. However, most importantly, if used naively, this method tends to fail when one departs from assumptions that our probabilistic model is identifiable (i.e.,  $p_\theta(\cdot|x) = p_{\tilde{\theta}}(\cdot|x)$  even if  $\theta \neq \tilde{\theta}$ ). We mitigate this problem by introducing the scaling parameters  $w_i$  in Eq. (1) to deal with it.

Prevalent constructions of anytime-valid confidence intervals rely on carefully derived concentration results and for a specific estimator such as the least-squares estimator and noise sub-families such as sub-Gaussian, sub-Bernoulli and sub-Poisson Abbasi-Yadkori et al. (2011); Faury et al. (2020); Mutný and Krause (2021). Their constructions involve bounding the suprema of collections of self-normalized stochastic processes (Faury et al., 2020; Mutný and Krause, 2021; Chowdhury et al., 2022). To facilitate closed-form expressions, worst-case parameters are introduced that prohibitively affect the size of the sets – making them much larger than they need to be.

Chowdhury et al. (2022) use the exact form of the likelihood to build confidence sets for parameters of exponential families. However, their approach is restricted to exponential family distributions. They use self-normalization and mixing techniques to explicitly determine the size of the confidence set and do not use an online learning subroutine as we do here. Neiswanger and Ramdas (2021) use likelihood ratios for bandit optimization with possibly misspecified Gaussian processes but is not tractable beyond Gaussian likelihoods. The relation between online convex optimization and confidence sets has been noted in so-called online-to-confidence conversions (Abbasi-Yadkori et al., 2012; Jun et al., 2017; Orabona and Jun, 2021; Zhao et al., 2022), where the existence of a low-regret learner implies a small confidence set. However, these sets still use potentially loose regret bounds to define confidence sets. Our definition is *implicit*. We do not necessarily need a regret bound to run our method, as the radius will depend on the actual, instance-dependent performance of the learner.

**Conclusion** In this work, we generalized and analyzed sequential likelihood ratio confidence sets for adaptive inference. We showed that with well-specified likelihoods, this procedure gives small, any-time valid confidence sets with model-agnostic and precise coverage. For generalized linear models, we quantitatively analyzed their size and shape. We invite practitioners to explore and use this very versatile and practical methodology for sequential decision-making tasks.

## Acknowledgments and Disclosure of Funding

We thank Wouter Koolen and Aaditya Ramdas for helpful discussions as well as for organizing the SAVI workshop where these discussions took place. NE acknowledges support from the Swiss Study Foundation and the Zeno Karl Schindler Foundation. MM has received funding from the Swiss National Science Foundation through NFP75. This publication was created as part of NCCR Catalysis (grant number 180544), a National Centre of Competence in Research funded by the Swiss National Science Foundation.

## References

- Abbasi-Yadkori, Y., Pál, D., and Szepesvari, C. (2011). Improved algorithms for linear stochastic bandits. In *Advances in Neural Information Processing Systems*, pages 2312–2320.
- Abbasi-Yadkori, Y., Pal, D., and Szepesvari, C. (2012). Online-to-confidence-set conversions and application to sparse stochastic bandits. In Lawrence, N. D. and Girolami, M., editors, *Proceedings of the Fifteenth International Conference on Artificial Intelligence and Statistics*, volume 22 of *Proceedings of Machine Learning Research*, pages 1–9, La Palma, Canary Islands. PMLR.
- Abe, N. and Long, P. M. (1999). Associative reinforcement learning using linear probabilistic concepts. In *ICML*, pages 3–11. Citeseer.
- Auer, P. (2002). Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research*, 3(Nov):397–422.
- Azoury, K. S. and Warmuth, M. K. (1999). Relative loss bounds for on-line density estimation with the exponential family of distributions. *Machine Learning*, 43:211–246.
- Baker, M. (2016). 1,500 scientists lift the lid on reproducibility. *Nature*, 533:452–454.
- Breslow, N. E. (1975). Analysis of survival data under the proportional hazards model. *International Statistical Review / Revue Internationale de Statistique*.
- Carpentier, A., Vernade, C., and Abbasi-Yadkori, Y. (2020). The elliptical potential lemma revisited. *ArXiv*, abs/2010.10182.
- Cesa-Bianchi, N. and Lugosi, G. (2006). *Prediction, Learning, and Games*. Cambridge University Press, New York, NY, USA.
- Chowdhury, S. R., Saux, P., Maillard, O.-A., and Gopalan, A. (2022). Bregman deviations of generic exponential families. *arXiv preprint arXiv:2201.07306*.
- Cox, D. R. (1997). Some remarks on the analysis of survival data. In *Proceedings of the First Seattle Symposium in Biostatistics*, pages 1–9. Springer.
- Cucker, F. and Smale, S. (2002). On the mathematical foundations of learning. *Bulletin of the American mathematical society*, 39(1):1–49.
- Faury, L., Abeille, M., Calauzènes, C., and Fercoq, O. (2020). Improved optimistic algorithms for logistic bandits. In *ICML2020: Proceedings of the 37th International Conference on International Conference on Machine Learning*.
- Filippi, S., Cappe, O., Garivier, A., and Szepesvári, C. (2010). Parametric bandits: The generalized linear case. In *Advances in neural information processing systems*, pages 586–594.
- Gavana, A. (2021). infinity global optimization benchmarks and amngo. [http://infinity77.net/global\\_optimization/index.html#](http://infinity77.net/global_optimization/index.html#).
- Hazan, E. (2016). Introduction to online convex optimization. *Found. Trends Optim.*, 2:157–325.
- Hazan, E., Agarwal, A., and Kale, S. (2006). Logarithmic regret algorithms for online convex optimization. *Machine Learning*, 69:169–192.
- Howard, S. R., Ramdas, A., McAuliffe, J., and Sekhon, J. (2018). Time-uniform chernoff bounds via nonnegative supermartingales. *Probability Surveys*, 17:257–317.

- Jun, K.-S., Bhargava, A., Nowak, R., and Willett, R. (2017). Scalable generalized linear bandits: Online computation and hashing. *Advances in Neural Information Processing Systems*, 30.
- Katz-Samuels, J. and Jamieson, K. (2020). The true sample complexity of identifying good arms. In Chiappa, S. and Calandra, R., editors, *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pages 1781–1791. PMLR.
- Kleinbaum, D. G. and Klein, M. (2010). *Survival analysis*, volume 3. Springer.
- Lattimore, T. and Szepesvári, C. (2020). *Bandit Algorithms*. Cambridge University Press.
- Makarova, A., Usmanova, I., Bogunovic, I., and Krause, A. (2021). Risk-averse heteroscedastic bayesian optimization. In *Proc. Neural Information Processing Systems (NeurIPS)*.
- McCullagh, P. (2018). *Generalized linear models*. Chapman and Hall.
- Mukherjee, S., Tripathy, A. S., and Nowak, R. (2022). Chernoff sampling for active testing and extension to active regression. In *International Conference on Artificial Intelligence and Statistics*, pages 7384–7432. PMLR.
- Mutný, M. and Krause, A. (2021). No-regret algorithms for capturing events in poisson point processes. In Meila, M. and Zhang, T., editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 7894–7904. PMLR.
- Mutný, M. and Krause, A. (2022). Experimental design of linear functionals in reproducing kernel hilbert spaces. In *Proc. Neural Information Processing Systems (NeurIPS)*.
- Neiswanger, W. and Ramdas, A. (2021). Uncertainty quantification using martingales for misspecified gaussian processes. In Feldman, V., Ligett, K., and Sabato, S., editors, *Proceedings of the 32nd International Conference on Algorithmic Learning Theory*, volume 132 of *Proceedings of Machine Learning Research*, pages 963–982. PMLR.
- Orabona, F. (2019). A modern introduction to online learning.
- Orabona, F. and Jun, K.-S. (2021). Tight concentrations and confidence sequences from the regret of universal portfolio. *IEEE Transactions on Information Theory*, 70:436–455.
- Ouhamma, R., Maillard, O.-A., and Perchet, V. (2021). Stochastic online linear regression: the forward algorithm to replace ridge. In *Neural Information Processing Systems*.
- Ramdas, A., Grünwald, P., Vovk, V., and Shafer, G. (2022). Game-theoretic statistics and safe anytime-valid inference.
- Robbins, H., Siegmund, D., et al. (1972). A class of stopping rules for testing parametric hypotheses. In *Proceedings of the Sixth Berkeley Symposium on Mathematical Statistics and Probability, Volume 4: Biology and Health*. The Regents of the University of California.
- Schölkopf, B., Herbrich, R., and Smola, A. (2001). A generalized representer theorem. In *Computational learning theory*, pages 416–426. Springer.
- Srinivas, N., Krause, A., Kakade, S. M., and Seeger, M. W. (2009). Gaussian process optimization in the bandit setting: No regret and experimental design. In *International Conference on Machine Learning*.
- Ville, J.-L. (1939). Étude critique de la notion de collectif.
- Vovk, V. (2001). Competitive on-line statistics. *International Statistical Review*, 69.
- Wainwright, M. J. (2019). *High-dimensional statistics: A non-asymptotic viewpoint*. Cambridge university press.
- Wald, A. (1945). Sequential tests of statistical hypotheses. *Annals of Mathematical Statistics*, 16:256–298.

Wasserman, L. A., Ramdas, A., and Balakrishnan, S. (2020). Universal inference. *Proceedings of the National Academy of Sciences*, 117:16880 – 16890.

Zhao, H., Zhou, D., He, J., and Gu, Q. (2022). Bandit learning with general function classes: Heteroscedastic noise and variance-dependent regret bounds. *ArXiv*, abs/2202.13603.



## A Proofs of Theorem 1 and 2

### A.1 GLM Families

Table 1: Examples of exponential family distributions.

Name	$A(z)$	$A'(z)$	$T(y)$	$\mu$	$L$
Gaussian	$z^2/(2\sigma^2)$	$z/\sigma^2$	$y/\sigma$	$1/\sigma^2$	$1/\sigma^2$
Poisson	$\exp(z)$	$\exp(z)$	$y$	$\exp(-B)$	$\exp(B)$
Binomial	$\log(1 + \exp(z))$	$\frac{1}{1+\exp(-z)}$	$y$	$\mathcal{O}(\exp(-B))$	$1/4$
Weibull	$k \log(z) - \log k$	$k/z$	$y^k$	$1/B^2$	$\infty$

### A.2 Proof of Theorem 1 (Coverage)

*Proof.* Starting with  $\mathbb{E}[R_t(\theta_*) | \mathcal{F}_{t-1}]$

$$\begin{aligned}
&= \mathbb{E} \left[ R_{t-1}(\theta_*) \frac{p_{\hat{\theta}_t}^{w_t}(y_t | x_t)}{p_{\theta_*}^{w_t}(y_t | x_t)} \middle| \mathcal{F}_{t-1} \right] \\
&= R_{t-1}(\theta_*) \int \frac{p_{\hat{\theta}_t}^{w_t}(y | x_t)}{p_{\theta_*}^{w_t}(y | x_t)} p_{\theta_*}(y | x_t) dy \\
&= R_{t-1}(\theta_*) e^{((w_t-1)D_{w_t}^r(p_{\theta_*}(x_t), p_{\hat{\theta}_t}(x_t)))} \leq R_{t-1}(\theta_*).
\end{aligned}$$

The second equality is due to the fact that  $R_{t-1}(\theta_*)$  only depends on  $x_1, y_1$  through  $x_{t-1}, y_{t-1}$ . Since  $\hat{\theta}_t$  is  $\mathcal{F}_{t-1}$  measurable by assumption,  $p_{\hat{\theta}_t}$  is a density, and if  $w_t < 1$ , the integral is equal to an exponential of the Rényi-divergence  $D_{w_t}^r(\cdot, \cdot)$ . The negativity of the exponent follows from  $w_t < 1$  and the non-negativity of the divergence. Note that in the "degenerate" case of  $w_t = 1$ , we can easily see that the integral is over a density (cancellation), and hence also bounded by 1. The last part of the statement follows easily by using Ville's inequality for supermartingales.  $\square$

All the elements of the above proof appear in [Wasserman et al. \(2020\)](#) albeit separately, and not with time-varying powered robust likelihoods.

### A.3 Proof of Theorem 2 (Bias)

We will need a gradient characterization of strong-convexity, which we prove in the following lemma.

**Lemma 1** (Convexity: Gradient). *Defining*

$$F_t(\theta) = - \sum_{s=1}^t \mathbb{E}_{\theta_*} [\nabla \log_{\theta} p(y_s | x_s) | \mathcal{F}_{s-1}],$$

under the assumption  $p_{\theta}(y_s | x_s) = -g(x_s^{\top} \theta)$  and  $g$  is  $\mu$ -strongly convex, we have for any  $\theta \in \Theta$ :

$$(F_t(\theta) - F_t(\theta_*))^{\top} (\theta - \theta_*) \geq \|\theta - \theta_*\|_{\mathbf{V}_t^{\mu;0}}^2.$$

*Proof.* We assume that  $g$  is  $\mu$ -strongly convex. Therefore, for any  $s \leq t$ , we get the two inequalities

$$\begin{aligned}
g(x_s^{\top} \theta) - g(x_s^{\top} \theta_*) &\geq g'(x_s^{\top} \theta_*) (x_s^{\top} \theta - x_s^{\top} \theta_*) + \frac{\mu}{2} \|x_s^{\top} \theta - x_s^{\top} \theta_*\|_2^2 \\
g(x_s^{\top} \theta_*) - g(x_s^{\top} \theta) &\geq g'(x_s^{\top} \theta) (x_s^{\top} \theta_* - x_s^{\top} \theta) + \frac{\mu}{2} \|x_s^{\top} \theta_* - x_s^{\top} \theta\|_2^2.
\end{aligned}$$

Adding these two together, we obtain

$$0 \geq (g'(x_s^{\top} \theta_*) - g'(x_s^{\top} \theta)) (x_s^{\top} (\theta - \theta_*)) + \mu \|\theta - \theta_*\|_{x_s x_s^{\top}}^2.$$

Observing that  $-\nabla_{\theta} p_{\theta}(y_s | x_s) = -g'(x_s^{\top} \theta) x_s$ , we can equivalently write

$$0 \geq (\nabla \log p_{\theta}(y_s | x_s) - \nabla \log p_{\theta_*}(y_s | x_s))^{\top} (\theta - \theta_*) + \mu \|\theta - \theta_*\|_{x_s x_s^{\top}}^2.$$

This holds for any realization of  $y_s$ , and hence taking expectations yields

$$(\mathbb{E}[\nabla \log p_{\theta_*}(y_s | x_s) | \mathcal{F}_{s-1}] - \mathbb{E}[\nabla \log p_{\theta}(y_s | x_s) | \mathcal{F}_{s-1}])^\top (\theta - \theta_*) \geq \mu \|\theta - \theta_*\|_{x_s x_s^\top}^2.$$

Summing up over  $s \leq t$  and using the definition of  $F_t$  we get

$$(F_t(\theta) - F_t(\theta_*))^\top (\theta - \theta_*) \geq \|\theta - \theta_*\|_{\mathbf{V}_t^\mu; 0}^2.$$

□

Notice that the estimator in Alg. 1 has to fulfil the KKT conditions. We will denote the condition for belonging to the set as  $h(\theta) \leq B^2$ , where  $h$  is a squared and twice-differentiable norm (there are many choices beyond  $\|\cdot\|_2^2$ ). The KKT conditions are

$$\begin{aligned} \sum_{s=1}^t -\nabla_{\theta} \log p_{\theta}(y_s | x_s) + \nabla \psi_t(\theta) + l \nabla h(\theta) &= 0 \\ l(h(\theta) - B^2) &= 0 \\ l &\geq 0, \end{aligned} \quad (12)$$

where the second and third conditions represent a complementary slackness requirement. Notice that the system of these equations has a unique solution due to the strong-convexity of the objective, and has to attain a unique minimum on a compact convex subset of  $\mathbb{R}^d$ . Adding the same quantity on both sides of (12) yields

$$\begin{aligned} &\sum_{s=1}^t -\mathbb{E}_{\theta_*}[\nabla_{\theta} \log p_{\theta}(y_s | x_s) | \mathcal{F}_{s-1}] + \nabla \psi_t(\theta) + l \nabla h(\theta) \\ &= \underbrace{\sum_{s=1}^t [\nabla_{\theta} \log p_{\theta}(y_s | x_s) - \mathbb{E}_{\theta_*}[\nabla_{\theta} \log p_{\theta}(y_s | x_s) | \mathcal{F}_{s-1}]]}_{:= E_t}. \end{aligned} \quad (13)$$

This line motivates the definition of the error-free estimator in  $\theta_t^\times$  (6), where  $E_t$  is set to zero. We will also make use of a fundamental property of the score (gradient of log-likelihood), namely

$$\mathbb{E}_{y_t \sim p_{\theta_*}(\cdot | x)}[\nabla \log p_{\theta_*}(y_t | x_t) | \mathcal{F}_{t-1}] = 0. \quad (14)$$

A classical textbook reference for this is e.g. [McCullagh \(2018\)](#) but any other classical statistics textbook should contain it. Using these observations, we can already prove Theorem 2.

*Proof of Theorem 2.* Using the optimality conditions of  $\theta_t^\times$ ,  $h(\theta) = \|\theta\|_2^2$  and  $\psi_t(\theta) = \|\theta\|_2^2$ , we obtain the following statements:

$$\begin{aligned} &\sum_{s=1}^t -\mathbb{E}_{\theta_*}[\nabla \log p_{\theta_t^\times}(y_s | x_s) | \mathcal{F}_{s-1}] + \lambda \theta_t^\times + 2l \theta_t^\times = 0 \\ \implies &\sum_{s=1}^t -\mathbb{E}_{\theta_*}[\nabla \log p_{\theta_t^\times}(y_s | x_s) | \mathcal{F}_{s-1}] + \mathbb{E}_{\theta_*}[\nabla \log p_{\theta_*}(y_s | x_s) | \mathcal{F}_{s-1}] + \lambda \theta_t^\times + 2l \theta_t^\times = 0, \end{aligned}$$

where in the last line we used the property (14). Now, notice that since we know  $\theta_*$  is generating the data, the best possible explanation without enforcing the constraint and the regularization would be to set  $\theta_t^\times = \theta_*$  as the cross-entropy is minimized at this point, and the above is just the optimality condition for optimizing the cross-entropy between these two distribution. Of course, this is only in the absence of regularization or constraints i.e.  $\lambda = 0$ . Now, with the regularization constraint, as the true  $\theta_*$  lies inside the constraint  $h(\theta) \leq B^2$ , and both the regularization and constraints induce star-shaped sets, their effect is to make  $\theta_t^\times$  smaller in norm than  $\theta^*$ . This holds generally for any  $h$  which is a norm. As a consequence of this consideration,  $\|\theta_t^\times\|_2 < B$ , and then the complementary slackness dictates that  $l = 0$ .

We can therefore proceed with this simplification. Let us use the shorthand  $F_t(\theta) = \sum_{s=1}^t -\mathbb{E}_{\theta_*}[\nabla \log p_{\theta}(y_s | x_s) | \mathcal{F}_{s-1}]$  and compute

$$F_t(\theta_t^\times) - F_t(\theta_*) + \lambda(\theta_t^\times - \theta_*) = -\lambda \theta_*$$

$$\begin{aligned}
&\implies (\theta_t^\times - \theta_\star)^\top (F_t(\theta_t^\times) - F(\theta_\star) + \lambda(\theta_t^\times - \theta_\star)) &&= -\lambda(\theta_t^\times - \theta_\star)^\top \theta_\star \\
\stackrel{\text{Lemma 1}}{\implies} \|\theta_t^\times - \theta_\star\|_{\mathbf{V}_t^{\mu, \lambda}}^2 &&\leq -\lambda(\theta_t^\times - \theta_\star)^\top \theta_\star. \quad (15)
\end{aligned}$$

It suffices to apply the Cauchy-Schwarz Inequality and invoke (15):

$$\begin{aligned}
\text{bias}_{x_s}(\hat{\theta}_s)^2 &= (x_s^\top (\hat{\theta}_s^\times - \theta_\star))^2 \\
&\leq \|x_s\|_{(\mathbf{V}_t^{\mu, \lambda})^{-1}}^2 \|\theta_t^\times - \theta_\star\|_{\mathbf{V}_t^{\mu, \lambda}}^2 \\
&\leq \lambda \|x_s\|_{(\mathbf{V}_t^{\mu, \lambda})^{-1}}^2 \lambda (\theta_t^\times - \theta_\star)^\top (-\theta_\star) \\
&\leq \lambda \|x_s\|_{(\mathbf{V}_t^{\mu, \lambda})^{-1}}^2 \|(\theta_t^\times - \theta_\star)\|_2 \|\theta_\star\|_2 \\
&\leq 2\lambda \|x_s\|_{(\mathbf{V}_t^{\mu, \lambda})^{-1}}^2 \|\theta_\star\|_2^2,
\end{aligned}$$

where in the last inequality we used  $\|\theta_t^\times\|_2 \leq \|\theta_\star\|_2$ , due to the regularizer, as explained above.  $\square$

**GLM models** Let us define the processes

$$S_t = \sum_{s=1}^t x_s T(y_s) \quad \text{and} \quad W_t = \sum_{s=1}^t x_s A'(x_s^\top \theta_\star).$$

In this scenario, an equivalent of (13) then involves the gradient of the regularized (unweighted) log-partition function  $Z_t^\lambda$  we defined in (8) and is equal to

$$\sum_{s=1}^t A'(x_s^\top \theta) x_s + \nabla \psi_t(\theta) + l \nabla h(\theta) = \tilde{E}_t, \quad (16)$$

where  $\tilde{E}_t = S_t$  for  $\hat{\theta}_t$  and  $\tilde{E}_t = W_t$  for  $\theta_t^\times$ .

## B Proof of Theorem 3 (Bregman Ball Confidence Set)

**Proof sketch** We give a quick sketch of the proof. To bound the size of the sets, we will draw inspiration from the i.i.d. parameter estimation analysis of Wasserman et al. (2020) and separate out the likelihood ratio in a part that relates the true parameter with the estimator sequence (i.e. regret), and a part that is independent of the estimator and characterized by a supremum of a stochastic process. We want to show that *any* point which is far away from the true parameter will eventually not be included in the confidence set anymore. Defining  $\mathcal{L}^{(t)}(\{\hat{\theta}_s\}_{s=1}^t)$  as  $\prod_{i=1}^t p_{\hat{\theta}_i}^{w_i}(y_i | x_i)$ , we wish to show that for any  $\theta$  far from  $\theta_\star$ , we have

$$\log \frac{1}{R_t(\theta)} = \log \frac{\mathcal{L}^{(t)}(\theta)}{\mathcal{L}^{(t)}(\theta_\star)} + \log \frac{\mathcal{L}^{(t)}(\theta_\star)}{\mathcal{L}^{(t)}(\{\hat{\theta}_s\}_{s=1}^t)} \leq \log(\alpha),$$

which is equivalent to saying that  $\theta \notin \mathcal{C}_t$ . The second term corresponds to our notion of regret exactly ( $\mathcal{R}_t$ , as discussed above). The first term is what we will focus on. We will bound the supremum of  $\log \frac{\mathcal{L}^{(t)}(\theta)}{\mathcal{L}^{(t)}(\theta_\star)}$  for all  $\theta$  sufficiently far away from  $\theta_\star$ . "Far away" will be measured in the Bregman divergence outlined above. Note that this quantity can be expected to be negative, in general, (especially for "implausible" parameters), since with enough data,  $\theta_\star$  should appear much more likely. Writing this ratio out, we will observe that it is equal to

$$-D_{Z_t^\lambda}(\theta, \theta_\star) + \underbrace{(\theta - \theta_\star, \sum_{s=1}^t w_s x_s (T(y_s) - \mathbb{E}_{\theta_\star}[T(y_s)]))}_{\approx \tilde{S}_t}.$$

At this point, it will be sufficient to bound the cross term (second term) over the whole of  $\Theta$ . We view this supremum as part of the Legendre Fenchel transform of the function  $\mathcal{B}_t(\lambda) = D_{Z_t^\lambda}(\theta_\star + \lambda, \theta_\star)$ :

$$\sup_{\lambda \in \mathbb{R}^d} (\lambda^\top \tilde{S}_t - \mathcal{B}_t(\lambda)) = (\mathcal{B}_t)^\star(\tilde{S}_t)$$

and harness duality properties of the Bregman divergence, along with known concentration arguments (Chowdhury et al., 2022, Theorem A.1).

## B.1 Technical Lemmas

We need to introduce the concept of Legendre functions:

**Definition 1.** Let  $f : \mathbb{R}^d \rightarrow \mathbb{R}$  be a convex function and  $C = \text{int}(\text{dom}(f))$ . Then, a function is called Legendre if it satisfies

1.  $C$  is non-empty.
2.  $f$  is differentiable and strictly convex on  $C$ .
3.  $\lim_{n \rightarrow \infty} \|\nabla f(x_n)\| = \infty$  for any sequence  $(x_n)_n$  with  $x_n \in C$  for all  $n$  and  $\lim_{n \rightarrow \infty} x_n = x$  for some  $x \in \partial C$ .

This means that the gradient has to blow up near the edge of the domain. Note as well that the boundary condition is vacuous if there is no boundary. Legendre functions have some nice properties, most importantly regarding the bijectivity of their gradients (see e.g. [Lattimore and Szepesvári \(2020\)](#)):

**Lemma 2.** For a Legendre function  $f : \mathbb{R}^d \rightarrow \mathbb{R}$

1.  $\nabla f$  is a bijection between  $\text{int}(\text{dom}(f))$  and  $\text{int}(\text{dom}(f^*))$  with the inverse  $(\nabla f)^{-1} = \nabla f^*$ .
2.  $D_f(x, y) = D_{f^*}(\nabla f(y), \nabla f(x))$  for all  $x, y \in \text{int}(\text{dom}(f))$ .
3. The Fenchel conjugate  $f^*$  is also Legendre.

With this, we can prove a slightly extended result, that appears as Lemma 2.1. in [Chowdhury et al. \(2022\)](#).

**Lemma 3.** For a Legendre function  $f$  we have the identity

$$D_f(x, y) = (D_{f,x})^*(\nabla f(y) - \nabla f(x))$$

where we define  $D_{f,x}(\lambda) = D_f(x + \lambda, x)$ .

**Notational Shorthands** Remember the model (7), with log-partition function  $A$ . We define  $A_s(\theta) = w_s A(x_s^\top \theta)$  and  $T_s(y) := w_s x_s T(y)$  to denote the log-partition function and the response function of the same exponential family distribution, but parametrized by  $\theta$  instead of  $x_s^\top \theta$ . That this is a valid parametrization can easily be seen from the likelihood definition. Indeed, denote by  $p_\beta^{EF}$  the exponential family reward distribution with parameter  $\beta$ . Then our model (7) can be seen to satisfy

$$p_\theta(y | x_s) = p_{x_s^\top \theta}(y) = h(y) \exp(T(y)x_s^\top \theta - A(x_s^\top \theta)) = h(y) \exp(T_s(y)^\top \theta - A_s(\theta)).$$

Exponentiating the likelihood with a weighting  $w_s$  gives rise to another exponential family distribution. We can see that

$$p_\theta^{w_s}(y | x_s) = h^{w_s}(y) \exp(w_s T(y)x_s^\top \theta - w_s A(x_s^\top \theta)) = h^{w_s}(y) \exp(T_s(y)^\top \theta - A_s(\theta)).$$

Note that this does not necessarily integrate to one, but it is easy to see that there is a normalization function  $\tilde{h}$  that makes it integrate to one. Therefore, the following is a valid parametrization of an exponential family distribution:

$$\tilde{h}(y) \exp(T_s(y)^\top \theta - A_s(\theta)).$$

Additionally, let  $A_0(\theta) = \frac{\nu}{2} \|\theta\|_2^2$  be defined on  $\mathbb{R}^d$  (i.e. a Legendre Function). We will also define the estimator

$$\tilde{\theta}_t = (\nabla Z_t^\nu)^{-1} \left( \sum_{s=1}^t T_s(y_s) \right).$$

This is a well-defined quantity because the gradient will be invertible, by Lemma 2 above.

Conveniently, [Chowdhury et al. \(2022\)](#) prove the following Theorem 7 using an elegant application of the method of mixtures.

**Proposition 1** (Theorem 7 in Chowdhury et al. (2022)). *With probability  $1 - \delta$ , for all  $t \in \mathbb{N}$*

$$D_{Z_t^\nu}(\theta_\star, \tilde{\theta}_t) \leq \log(1/\delta) + A_0(\theta_\star) + \Gamma_t^\nu,$$

where

$$\Gamma_t^\nu = \log \left( \frac{\int_{\mathbb{R}^d} \exp(-\frac{1}{2}\|\theta\|_2^2) d\theta}{\int_{\mathbb{R}^d} \exp(-D_{Z_t^\nu}(\theta, \tilde{\theta}_t)) d\theta} \right).$$

Lastly, we will need the (one-argument) function

$$\mathcal{B}_t(\lambda) = D_{Z_t^\nu}(\theta + \lambda, \theta),$$

i.e. a shortcut for the Bregman divergence of  $Z_t^\nu$  at  $\theta$ . We use this one-argument function as we will be interested in its dual. We will also need a lemma on the sub-homogeneity properties of this object.

**Lemma 4.** *Under Assumption 1, for  $\theta \in \Theta$  and  $\lambda$  such that  $\theta + \lambda \in \Theta$ , we have for any  $\gamma \leq \frac{\mu}{2L}$*

$$\mathcal{B}_t(\gamma\lambda) \leq \frac{1}{2}\gamma\mathcal{B}_t(\lambda),$$

i.e. function  $g(\gamma) = \mathcal{B}_t(\gamma\lambda)$  is sub-homogeneous with contraction parameter  $\frac{1}{2}$  on  $[0, \frac{\mu}{2L}]$ .

See Appendix B.3 for a proof.

## B.2 Proof of Theorem 3

As mentioned in the main paper, our proof will show that all  $\theta$  sufficiently far from  $\theta_\star$  will be excluded from  $\mathcal{C}_t$  eventually. Equation (B) in the main text specifies the exclusion criterion, i.e.  $\theta \notin \mathcal{C}_t$  if and only if

$$\frac{1}{R_t(\theta)} = \log \frac{\mathcal{L}^{(t)}(\theta)}{\mathcal{L}^{(t)}(\theta_\star)} + \log \frac{\mathcal{L}^{(t)}(\theta_\star)}{\mathcal{L}^{(t)}(\{\hat{\theta}_s\}_{s=1}^t)} \leq \log(\alpha). \quad (17)$$

The second term is bounded by the regret of the online learner. And therefore, a sufficient condition for  $\theta \notin \mathcal{C}_t$  is

$$\log \left( \frac{\mathcal{L}^{(t)}(\theta)}{\mathcal{L}^{(t)}(\theta_\star)} \right) \leq \log(\alpha) - \mathcal{R}_t.$$

Henceforth, we will be interested in having an explicit set  $\tilde{\mathcal{C}}_t$  such that we can upper bound

$$\sup_{\theta \notin \tilde{\mathcal{C}}_t} \log \left( \frac{\mathcal{L}^{(t)}(\theta)}{\mathcal{L}^{(t)}(\theta_\star)} \right). \quad (18)$$

This will imply that that  $\tilde{\mathcal{C}}_t^c \subset \mathcal{C}_t^c$ , or in other words,  $\mathcal{C}_t \subset \tilde{\mathcal{C}}_t$ . Without further ado, let us derive a more convenient form of the ratio in question

$$\begin{aligned} \log \left( \frac{\mathcal{L}^{(t)}(\theta)}{\mathcal{L}^{(t)}(\theta_\star)} \right) &= \log \left( \frac{\prod_{s=1}^t h(y_s) \exp(w_s x_s^\top \theta T(y_s) - w_s A(x_s^\top \theta))}{\prod_{s=1}^t h(y_s) \exp(w_s x_s^\top \theta T(y_s) - w_s A(x_s^\top \theta_\star))} \right) \\ &= \sum_{s=1}^t w_s x_s^\top \theta T(y_s) - w_s A(x_s^\top \theta) - w_s x_s^\top \theta_\star T(y_s) + w_s A(x_s^\top \theta_\star) \\ &= \sum_{s=1}^t \langle \theta - \theta_\star, w_s T(y_s) x_s \rangle + w_s A(x_s^\top \theta_\star) - w_s A(x_s^\top \theta) \\ &= \sum_{s=1}^t \langle \theta - \theta_\star, w_s T(y_s) x_s \rangle - (w_s A(x_s^\top \theta) - w_s A(x_s^\top \theta_\star)) - x_s^\top (\theta - \theta_\star) w_s A'(x_s^\top \theta_\star) \\ &\quad + x_s^\top (\theta - \theta_\star) w_s A'(x_s^\top \theta_\star) \\ &= \sum_{s=1}^t \langle \theta - \theta_\star, w_s T(y_s) x_s \rangle - w_s D_A(x_s^\top \theta, x_s^\top \theta_\star) - x_s^\top (\theta - \theta_\star) w_s A'(x_s^\top \theta_\star) \end{aligned}$$



$$= - \sum_{s=1}^t w_s D_A(x_s^\top \theta, x_s^\top \theta_\star) + \sum_{s=1}^t \langle \theta - \theta_\star, w_s T(y_s) x_s - x_s w_s A'(x_s^\top \theta_\star) \rangle.$$

We can switch parametrizations as described above:

$$\begin{aligned} \log \left( \frac{\mathcal{L}^{(t)}(\theta)}{\mathcal{L}^{(t)}(\theta_\star)} \right) &= -D_{Z_t^0}(\theta, \theta_\star) + \sum_{s=1}^t \langle \theta - \theta_\star, T_s(y_s) - \nabla A_s(\theta_\star) \rangle \\ &= -D_{Z_t^0}(\theta, \theta_\star) + \sum_{s=1}^t \langle \theta - \theta_\star, T_s(y_s) - \mathbb{E}_{\theta_\star}[T_s(y_s)] \rangle \\ &= -D_{Z_t^0}(\theta, \theta_\star) + \langle \theta - \theta_\star, S_t \rangle, \end{aligned} \quad (19)$$

where we define  $S_t := \sum_{s=1}^t (T_s(y_s) - \mathbb{E}_{\theta_\star}[T_s(y_s)])$ .

$Z_t^\nu$  is strictly convex whenever  $\nu \neq 0$ , and convex otherwise (it might also be strictly convex otherwise, corresponding to some cases where the  $x_s$  span the full  $d$ -dimensional Euclidean space and  $w_s > 0$ , which will be satisfied uniformly). We note that since  $\text{dom}(Z_t^\nu) = \mathbb{R}^d$ ,  $Z_t^\nu$  is, therefore, Legendre, and its gradient is invertible. We will relate our problem to this estimator via the duality properties developed above. First, note that by the well-known fact  $\mathbb{E}_{\theta_\star}[T_s(y_s)] = \nabla A_s(\theta_\star)$  and by the definition of  $\tilde{\theta}_t$ , we have

$$\begin{aligned} S_t &= \nabla Z_t^\nu \left( (\nabla Z_t^\nu)^{-1} \left( \sum_{s=1}^t T_s(y_s) \right) \right) - \nabla Z_t^0(\theta_\star) \\ &= \underbrace{\nabla Z_t^\nu(\tilde{\theta}_t) - \nabla Z_t^\nu(\theta_\star)}_{=: \tilde{S}_t} + \underbrace{\nabla A_0(\theta_\star)}_{=: \nu \theta_\star}. \end{aligned} \quad (20)$$

Now, we leverage the duality properties: We can write

$$\begin{aligned} \sup_{\lambda \in \mathbb{R}^d} \left( \lambda^\top \tilde{S}_t - \mathcal{B}_t(\lambda) \right) &\stackrel{(i)}{=} (\mathcal{B}_t)^\star(\tilde{S}_t) \\ &\stackrel{(20)}{=} (\mathcal{B}_t)^\star(\nabla Z_t^\nu(\tilde{\theta}_t) - \nabla Z_t^\nu(\theta_\star)) \\ &\stackrel{\text{Lemma 3}}{=} D_{Z_t^\nu}(\theta_\star, \tilde{\theta}_t), \end{aligned} \quad (21)$$

where (i) is simply the definition of the Legendre-Fenchel transform. Why did we do all this work? Well, we are interested in the supremum in Equation (18). It is sufficient to bound the supremum over all  $\theta \in \Theta$  of terms of the form (see Equation (19))

$$\langle \theta - \theta_\star, S_t \rangle.$$

While we could do a covering type argument (carefully relaxing the i.i.d. data assumptions typical in empirical process theory), it is much easier to relate this supremum to the estimator via duality.

With probability at least  $1 - \delta$ , Proposition 1 gives us a high-probability time-uniform bound on

$$D_{Z_t^\nu}(\theta_\star, \tilde{\theta}_t) \leq \log(1/\delta) + A_0(\theta_\star) + \Gamma_t^\nu,$$

and therefore, by plugging into Equation (21) and making the reparametrization  $\lambda = \gamma(\theta - \theta_\star)$  for some positive  $\gamma$ , it gives us

$$\forall t \geq 0 \forall \theta \in \mathbb{R}^d \forall \gamma \in \mathbb{R}_+ : \gamma \tilde{S}_t^\top(\theta - \theta_\star) - \mathcal{B}_t(\gamma(\theta - \theta_\star)) \leq \log(1/\delta) + A_0(\theta_\star) + \Gamma_t^\nu.$$

Therefore, for all  $t \geq 0$  and all  $\theta \in \mathbb{R}^d$ , the following holds:

$$\begin{aligned} S_t^\top(\theta - \theta_\star) &= \tilde{S}_t^\top(\theta - \theta_\star) + \nabla A_0(\theta_\star)^\top(\theta - \theta_\star) \\ &\leq \frac{1}{\gamma} \log(1/\delta) + \frac{1}{\gamma} A_0(\theta_\star) + \frac{1}{\gamma} \Gamma_t^\nu + \frac{1}{\gamma} \mathcal{B}_t(\gamma(\theta - \theta_\star)) + \nabla A_0(\theta_\star)^\top(\theta - \theta_\star). \end{aligned}$$

Since  $A_0(\theta) = \frac{\nu}{2} \|\theta\|_2^2$ , restricting our uniform bound over  $\theta \in \Theta$  gives us  $\forall t \geq 0 \forall \theta \in \Theta$ :

$$S_t^\top(\theta - \theta_\star) \leq \frac{1}{\gamma} \log(1/\delta) + \frac{\nu}{2\gamma} B^2 + \frac{1}{\gamma} \Gamma_t^\nu + \frac{1}{\gamma} \mathcal{B}_t(\gamma(\theta - \theta_\star)) + \nu B^2.$$

Now, we note that under Assumption 1, Lemma 4 kicks in and we have for any  $t \geq 0, \theta \in \Theta$  and  $\gamma = \frac{\mu}{2L}$

$$S_t^\top (\theta - \theta_\star) \leq \frac{1}{\gamma} \log(1/\delta) + \frac{\nu}{2\gamma} B^2 + \frac{1}{\gamma} \Gamma_t^\nu + \frac{1}{2} \mathcal{B}_t(\theta - \theta_\star) + \nu B^2. \quad (22)$$

Finally, we can use this in (19) to obtain

$$\begin{aligned} \log \left( \frac{\mathcal{L}^{(t)}(\theta)}{\mathcal{L}^{(t)}(\theta_\star)} \right) &\leq -\mathcal{B}_t(\theta - \theta_\star) + \langle \theta - \theta_\star, S_t \rangle \\ &\stackrel{(22)}{\leq} -\frac{1}{2} \mathcal{B}_t(\theta - \theta_\star) + \frac{1}{\gamma} \log(1/\delta) + \frac{\nu}{2\gamma} B^2 + \frac{1}{\gamma} \Gamma_t^\nu + \nu B^2 \\ &\leq -\frac{1}{2} \mathcal{B}_t(\theta - \theta_\star) + \frac{2L}{\mu} \left( \log(1/\delta) + \frac{\nu B^2}{2} + \Gamma_t^\nu \right) + \nu B^2. \end{aligned} \quad (23)$$

It remains to investigate the full likelihood ratio in (17):

$$\begin{aligned} &\frac{1}{R_t(\theta)} - \log(\alpha) \\ &= \log \frac{\mathcal{L}^{(t)}(\theta)}{\mathcal{L}^{(t)}(\theta_\star)} + \log \frac{\mathcal{L}^{(t)}(\theta_\star)}{\mathcal{L}^{(t)}(\{\hat{\theta}_s\}_{s=1}^t)} + \log(1/\alpha) \\ &\stackrel{(23) \& (3)}{\leq} -\frac{1}{2} \mathcal{B}_t(\theta - \theta_\star) + \frac{2L}{\mu} \left( \log(1/\delta) + \frac{\nu B^2}{2} + \Gamma_t^\nu \right) + \nu B^2 + \log(1/\alpha) + \mathcal{R}_t. \end{aligned} \quad (24)$$

Note that crucially for  $\theta \in \Theta$ , we have

$$\theta \notin \mathcal{C}_t \iff \frac{1}{R_t(\theta)} - \log(\alpha) \leq 0.$$

This is implied by

$$\mathcal{B}_{Z_t^\nu}(\theta, \theta_\star) \geq \frac{4L}{\mu} \left( \log(1/\delta) + \frac{\nu B^2}{2} + \Gamma_t^\nu \right) + 2\nu B^2 + 2\log(1/\alpha) + 2\mathcal{R}_t,$$

or, since  $L \geq \mu$ , more compactly by

$$\mathcal{B}_{Z_t^\nu}(\theta, \theta_\star) \geq \frac{4L}{\mu} (\log(1/\delta) + \nu B^2 + \Gamma_t^\nu) + 2\log(1/\alpha) + 2\mathcal{R}_t.$$

### B.3 Proof of Technical Lemmas

First, we will prove Lemma 3. The proof exactly follows Chowdhury et al. (2022), we include it here for convenience because it is very short.

*Proof.* By definition

$$\begin{aligned} &(D_{f,x})^*(\nabla f(y) - \nabla f(x)) \\ &= \sup_{a \in \mathbb{R}^d} (\langle a, \nabla f(y) - \nabla f(x) \rangle - D_{f,x}(a)) \\ &= \sup_{a \in \mathbb{R}^d} (\langle a, \nabla f(y) - \nabla f(x) \rangle - D_f(x+a, x)) \\ &= \sup_{a \in \mathbb{R}^d} (\langle a, \nabla f(y) - \nabla f(x) \rangle - f(x+a) + f(x) + \langle \nabla f(x), a \rangle) \\ &= \sup_{a \in \mathbb{R}^d} (\langle a, \nabla f(y) \rangle - f(x+a) + f(x)). \end{aligned}$$

Since  $f$  is strictly convex and differentiable, first-order optimality conditions imply that the optimal  $a$  satisfies  $\nabla f(y) - \nabla f(x+a) = 0$  ( $a$  is unconstrained). Since the gradient is invertible, we must have  $a = y - x$ . If we plug this into the above, we have

$$\begin{aligned} (D_{f,x})^*(\nabla f(y) - \nabla f(x)) &= \langle y - x, \nabla f(y) \rangle - f(y) + f(x) \\ &= f(x) - f(y) - \langle \nabla f(y), x - y \rangle \\ &= D_f(x, y). \end{aligned}$$

□

Now we prove Lemma 4. To this end, we will do a reduction to the one-dimensional case, and prove the one-dimensional result below.

**Lemma 5.** *Under Assumption 1, for any  $a \in [B, B]$ , any  $\gamma \in (0, \frac{\mu}{2L}]$  and any  $\Delta$  with  $a + \Delta \in [B, B]$*

$$A(a + \gamma\Delta) - A(a) - A'(a)\gamma\Delta \leq \frac{1}{2}\gamma[A(a + \Delta) - A(a) - A'(a)\Delta].$$

We prove that this implies the desired sublinearity of the full Bregman difference.

*Proof.* (of Lemma 4). Let  $\theta$  and  $\lambda$  be such that  $\theta, \theta + \lambda \in \Theta$ . We will first show that for any  $s \in \{0, \dots, t\}$ ,  $B_{A_s(\theta, \theta + \cdot)}$  is sublinear, and then the result follows by the linearity of the Bregman divergence. Define  $a_s = x_s^\top \theta$  and  $\Delta_s = x_s^\top \lambda$ . Then we have  $|\Delta_s| = |x_s^\top \lambda| \leq \|x_s\| \|\lambda\| \leq \|x_s\|(\|\theta\| + \|\theta + \lambda\|) \leq (B + B) \leq 2B$ . Similarly we have  $|a_s| \leq B$ . Hence we satisfy the premise of Lemma 5 and we deduce that

$$\begin{aligned} D_{A_s}(\theta + \gamma\lambda, \theta) &= w_s A(x_s^\top \theta + \gamma x_s^\top \lambda) - w_s A(x_s^\top \theta) + \langle x_s w_s A'(x_s^\top \theta), \gamma\lambda \rangle \\ &= w_s (A(a_s + \gamma\Delta_s) - A(a_s) + A'(a_s)\gamma\Delta_s) \\ &\leq \frac{w_s}{2} \gamma [A(a_s + \gamma\Delta_s) - A(a_s) + A'(a_s)\gamma\Delta_s] \\ &= \frac{1}{2} \gamma [w_s A(x_s^\top \theta + x_s^\top \lambda) - w_s A(x_s^\top \theta) + w_s A'(x_s^\top \theta) x_s^\top \lambda] \\ &= \frac{1}{2} \gamma D_{A_s}(\theta + \lambda, \theta). \end{aligned}$$

We also note that for  $\gamma \leq \frac{\mu}{2L} \leq \frac{1}{2}$ ,

$$D_{A_0}(\theta + \gamma\lambda, \theta) = \frac{\nu}{2} \|\gamma\lambda\|^2 = \frac{\gamma^2 \nu}{2} \|\lambda\|^2 \leq \frac{\gamma \nu}{4} \|\lambda\|^2 = \frac{1}{2} \gamma D_{A_0}(\theta + \lambda, \theta). \quad (25)$$

Therefore, by summing up the terms, we obtain

$$\mathcal{B}_t(\gamma\lambda) \leq \frac{1}{2} \gamma \mathcal{B}_t(\lambda).$$

□

Then it remains to prove that Assumption 1 implies Lemma 5.

*Proof.* (Lemma 5)  $L$ -Lipschitzness of  $A'$  implies smoothness of  $A$ . Additionally,  $\mu$ 's existence implies strong convexity of  $A$ . With this, we can write for any  $a$  and  $\Delta$  with  $a + \Delta \in [-B, B]$

$$\begin{aligned} A(a + \Delta) &\geq A(a) + A'(a)\Delta + \frac{\mu}{2}\Delta^2 \\ \implies A(a + \Delta) - A(a) - A'(a)\Delta &\geq \frac{\mu}{2}\Delta^2. \end{aligned}$$

Similarly,

$$A(a + \gamma\Delta) - A(a) - A'(a)\gamma\Delta \leq \frac{L}{2}\gamma^2\Delta^2.$$

Putting this together, we have

$$A(a + \gamma\Delta) - A(a) - A'(a)\gamma\Delta \leq \frac{L}{2}\gamma^2\Delta^2 = \frac{L\gamma^2}{\mu} \frac{\mu}{2}\Delta^2 \leq \frac{L\gamma}{\mu} \gamma [A(a + \Delta) - A(a) - A'(a)\Delta].$$

The question is therefore: when is  $\frac{L\gamma}{\mu} \leq \frac{1}{2}$ ? Clearly, choosing  $\gamma_0 = \frac{\mu}{2L}$  makes  $\frac{L\gamma}{\mu} \leq \frac{1}{2}$  for all  $\gamma \leq \gamma_0$ . □

## C FTRL Results: Proofs

### C.1 Technical Lemmas I: Exponential Families

**Lemma 6** (MGF for Exponential family).

$$\mathbb{E}[\exp(T(y)u)|x] = \exp(A(\theta_\star^\top x + u) - A(\theta_\star^\top x)).$$

*Proof.*

$$\begin{aligned} \mathbb{E}[\exp(T(y)u)|x] &= \int_y \exp(T(y)u)h(y) \exp(T(y)\theta_\star^\top x - A(\theta_\star^\top x))dy \\ &= \int \exp(T(y)(\theta_\star^\top x + u))h(y) \exp(-A(\theta_\star^\top x)) \\ &\quad \times \exp(-A(\theta_\star^\top x + u)) \exp(A(\theta_\star^\top x + u))dy \\ &= \exp(A(\theta_\star^\top x + u) - A(\theta_\star^\top x)), \end{aligned}$$

where the last step follows because the density of a new exponential family distribution with parameter  $\theta_\star^\top x + u$  also integrates to 1.  $\square$

### C.2 Technical Lemmas II: Elliptical Potential Lemma

We will repeatedly use instantiations of the following key lemma, known as the elliptical potential lemma. We will use the version from [Hazan et al. \(2006\)](#). Other variants are stated in [Abbasi-Yadkori et al. \(2011\)](#) or [Carpentier et al. \(2020\)](#).

**Lemma 7** (Lemma 11 in [Hazan et al. \(2006\)](#)). *Let  $u_s \in \mathbb{R}^d$  be a sequence of vectors such that  $\|u_s\| \leq r$ . Define  $\bar{\mathbf{V}}_t = \sum_{s=1}^t u_s u_s^\top + \lambda \mathbf{I}$ . Then*

$$\sum_{s=1}^t \|u_s\|_{\bar{\mathbf{V}}_s^{-1}}^2 \leq \log \left( \frac{\det \bar{\mathbf{V}}_t}{\det \lambda \mathbf{I}} \right) \leq d \log \left( \frac{r^2 t}{\lambda} + 1 \right).$$

We will also need a result where the time indices of the matrix are shifted. For this, note that if  $\lambda \geq r^2$ , then  $u_s u_s^\top \preceq r^2 \mathbf{I} \preceq \lambda \mathbf{I}$ , and so we get  $\bar{\mathbf{V}}_s \leq \bar{\mathbf{V}}_{s-1} + u_s u_s^\top \preceq \bar{\mathbf{V}}_{s-1} + \lambda \mathbf{I} \preceq 2\bar{\mathbf{V}}_{s-1}$ . Under our conditions, it follows that

$$\sum_{s=1}^t \|u_s\|_{\bar{\mathbf{V}}_{s-1}^{-1}}^2 \leq 2 \sum_{s=1}^t \|u_s\|_{\bar{\mathbf{V}}_s^{-1}}^2$$

**Corollary 1.** *We have the following bounds:*

$$\gamma_t^\lambda = \log \left( \frac{\det(\sum_{s=1}^t \mu x_s x_s^\top + \lambda \mathbf{I})}{\det(\lambda \mathbf{I})} \right) \leq d \log \left( \frac{\mu t}{\lambda} + 1 \right),$$

and

$$\sum_{s=1}^t \|x_s\|_{(\mathbf{V}_{s-1}^{\mu;\lambda})^{-1}}^2 \leq \frac{2}{\mu} \gamma_t^\lambda.$$

*Proof.* The first bound is trivial by instantiating  $u_s = \sqrt{\mu} x_s$ . The second bound is by noting

$$\sum_{s=1}^t \|x_s\|_{(\mathbf{V}_{s-1}^{\mu;\lambda})^{-1}}^2 = \frac{1}{\mu} \sum_{s=1}^t \|u_s\|_{(\mathbf{V}_{s-1}^{\mu;\lambda})^{-1}}^2 \leq \frac{2}{\mu} \sum_{s=1}^t \|u_s\|_{\bar{\mathbf{V}}_s^{-1}}^2 \leq \frac{2}{\mu} \gamma_t^\lambda.$$

$\square$

### C.3 Technical Lemmas III: Supermartingales

**Lemma 8** (Martingale Increment). *Define the parametrized random processes*

$$\mathcal{M}_j(r) = \exp(\nabla f_j(\theta_\star)^\top r - A'(x_j^\top \theta_\star) x_j^\top r - A(x_j^\top \theta_\star - x_j^\top r) + A(x_j^\top \theta_\star))$$

and

$$\mathcal{N}_j(r) = \exp(\nabla f_j(\theta_\star)^\top r - \frac{L}{2} r^\top x_j x_j^\top r).$$

Then, under Assumption 1 we have for any  $r \in \mathbb{R}^d$  that  $\mathbb{E}[\mathcal{M}_j(r) | \mathcal{F}_{j-1}] = 1$  and  $\mathbb{E}[\mathcal{N}_j(r) | \mathcal{F}_{j-1}] \leq 1$ .

*Proof.* First, using the form of the exponential family and recalling that  $\nabla_\theta f_j(\theta) = -\nabla_\theta \log p_\theta(y_j | x_j) = \nabla_\theta [A(x_j^\top \theta) - T(y_j) x_j^\top \theta]$  we obtain

$$\begin{aligned} & \mathbb{E}[\exp(\nabla f_j(\theta_\star)^\top r) | \mathcal{F}_{j-1}] \\ &= \int_y \exp(\nabla f_j(\theta_\star)^\top r) \times h(y) \exp(T(y) x_j^\top \theta_\star - A(x_j^\top \theta_\star)) dy \\ &= \int_y \exp(-T(y) x_j^\top r + A'(x_j^\top \theta_\star) x_j^\top r) \times h(y) \exp(T(y) x_j^\top \theta_\star - A(x_j^\top \theta_\star)) dy \\ &= \exp(A'(x_j^\top \theta_\star) x_j^\top r) \underbrace{\int_y h(y) \exp(T(y) (x_j^\top \theta_\star - x_j^\top r)) \exp(-A(x_j^\top \theta_\star - x_j^\top r)) dy}_{=1} \\ & \quad \times \exp(A(x_j^\top \theta_\star - x_j^\top r)) \exp(-A(x_j^\top \theta_\star)) \\ &= \exp(A'(x_j^\top \theta_\star) x_j^\top r) \exp(A(x_j^\top \theta_\star - x_j^\top r)) \exp(-A(x_j^\top \theta_\star)), \end{aligned}$$

which finishes the proof. The second statement follows by using  $L$ -smoothness on the last equation and therefore noting that  $\mathcal{N}_j(r) \leq \mathcal{M}_j(r)$ .  $\square$

**Lemma 9.** (Sequential Mixing) *Define the martingale process,*

$$M_t(r_1, \dots, r_t) = \prod_{s=1}^t \mathcal{N}_s(r_s),$$

and recursively define the mixture martingale,

$$\bar{M}_s = \bar{M}_{s-1} \times \int_r \mathcal{N}_s(r) p_s(r) dr,$$

where  $p_s$  is a probability distribution equal  $\mathcal{N}(0, \mathbf{H}_s^{-1})$ ,  $\mathbf{H}_s = \sum_{j=1}^{s-1} L x_j x_j^\top + \mathbf{I} \lambda \frac{L}{\mu}$ , and  $\bar{M}_0 = 1$ . Then the following statements hold

- $\{\bar{M}_s\}_s$  is an adapted super-martingale with respect to the usual filtration.
- $\bar{M}_t = \exp\left(\frac{\mu}{L} \sum_{s=1}^t \nabla f_s(\theta_\star)^\top (\mathbf{V}_s^{\mu;\lambda})^{-1} \nabla f_s(\theta_\star)\right) \sqrt{\frac{\det(\mathbf{I}\lambda)}{\det(\mathbf{V}_s^{\mu;\lambda})}}$ .

where

$$\mathbf{V}_s^{\mu;\lambda} = \sum_{j=1}^s \mu x_j x_j^\top + \lambda \mathbf{I}.$$

*Proof.* The first point follows from the fact that  $p_s(r)$  is deterministic conditioned on the sub- $\sigma$ -algebra  $\mathcal{F}_{s-1}$  (since  $p_s$  makes use of  $x_s$  but not  $x_{s+1}$ ). Therefore, under mild regularity conditions

$$\mathbb{E}[\bar{M}_s | \mathcal{F}_{s-1}] = \mathbb{E}\left[\bar{M}_{s-1} \int p_s(r) \mathcal{N}_s(r) dr | \mathcal{F}_{s-1}\right] = \bar{M}_{s-1} \int p_s(r) \mathbb{E}[\mathcal{N}_s(r) | \mathcal{F}_{s-1}] dr \leq \bar{M}_{s-1}.$$



In other words, mixing does not affect the supermartingale properties. For the second point, we derive an explicit form of the mixture martingale. Note that we can write out

$$\int_r \mathcal{N}_s(r) p_s(r) dr = \frac{1}{\sqrt{(2\pi)^d \det(\mathbf{H}_s^{-1})}} \int_r \exp\left(\nabla f_s(\theta_\star)^\top r - \frac{L}{2} \|r\|_{x_s x_s^\top}^2 - \frac{1}{2} r^\top \mathbf{H}_s r\right) dr. \quad (26)$$

We can complete the square to obtain

$$\begin{aligned} & \nabla f_s(\theta_\star)^\top r - \frac{L}{2} \|r\|_{x_s x_s^\top}^2 - \frac{1}{2} r^\top \mathbf{H}_s r \\ &= \frac{1}{2} \|\nabla f_s(\theta_\star)\|_{(\mathbf{H}_s + Lx_s x_s^\top)^{-1}}^2 - \frac{1}{2} \|r - (\mathbf{H}_s + Lx_s x_s^\top)^{-1} \nabla f_s(\theta_\star)\|_{\mathbf{H}_s + Lx_s x_s^\top}^2. \end{aligned}$$

The second term is the exponent of a Gaussian integral with covariance  $\mathbf{H}_{s+1}^{-1}$ , and therefore results in

$$\int_r \exp\left(-\frac{1}{2} \|r - (\mathbf{H}_s + Lx_s x_s^\top)^{-1} \nabla f_s(\theta_\star)\|_{\mathbf{H}_s + Lx_s x_s^\top}^2\right) dr = \sqrt{(2\pi)^d \det(\mathbf{H}_{s+1}^{-1})}.$$

Plugging this into (26) we get

$$\int_r \mathcal{N}_s(r) p_s(r) dr = \sqrt{\frac{\det \mathbf{H}_s}{\det \mathbf{H}_{s+1}}} \exp\left(\frac{1}{2} \|\nabla f_s(\theta_\star)\|_{\mathbf{H}_{s+1}^{-1}}^2\right).$$

By multiplying the individual steps, we can see that the determinant terms cancel in a telescoping product. This leads to the formulation

$$\bar{M}_t = \exp\left(\sum_{s=1}^t \nabla f_s(\theta_\star)^\top \mathbf{H}_{s+1}^{-1} \nabla f_s(\theta_\star)\right) \sqrt{\frac{\det \frac{\lambda L}{\mu} \mathbf{I}}{\det \mathbf{H}_{t+1}}}.$$

To conclude the proof, note that  $\mathbf{H}_{s+1} = \frac{L}{\mu} \mathbf{V}_s^{\mu; \lambda}$ . □

**Lemma 10.** *Under assumption of Lemma 9,*

$$\mathbb{P}\left(\sum_{s=1}^t \|\nabla f_s(\theta_\star)\|_{(\mathbf{V}_s^{\mu; \lambda})^{-1}}^2 \leq \frac{L}{\mu} \log\left(\frac{\det(\mathbf{V}_s^{\mu; \lambda})}{\det(\mathbf{I}\lambda)}\right) + \frac{L}{\mu} \log\left(\frac{1}{\delta}\right)\right) \leq \delta. \quad (27)$$

with probability  $1 - \delta$ .

*Proof.* The statement, follows by applying Ville's inequality for supermartingales, applying the logarithm, and rearranging. Namely,

$$\mathbb{P}(\bar{M}_t \geq \delta) = \mathbb{P}(\log(\bar{M}_t) \geq \log(\delta)) \leq \delta. \quad \square$$

The following results allow us to upper bound the weighted regret by the unweighted regret:

**Lemma 11** (Weighting Reduction). *Let  $\{\theta_s\}_{s=1}^t$  be a sequence of vectors adapted to the filtration  $\{\mathcal{F}_{s-1}\}_s$ . Define*

$$\Delta_t(\{\theta_s\}) = \sum_{s=1}^t w_s (f_s(\theta_s) - f_s(\theta_\star)) - f_s(\theta_s) + f_s(\theta_\star) = \sum_{s=1}^t (1 - w_s) (f_s(\theta_\star) - f_s(\theta_s)).$$

Then,  $P_t = \exp(\Delta_t(\{\theta_s\}_s))$  is a non-negative super-martingale for any choice of adapted  $\{w_s\}$ , and hence,

$$\sum_{s=1}^t w_s (f_s(\theta_s) - f_s(\theta_\star)) \leq \sum_{s=1}^t (f_s(\theta_s) - f_s(\theta_\star)) + \log\left(\frac{1}{\delta}\right)$$

with probability  $1 - \delta$  for all  $t \geq 0$ .

*Proof.*

$$\begin{aligned}
\mathbb{E}[P_t | \mathcal{F}_{t-1}] &= \mathbb{E}_{\theta_\star} \left[ \exp \left( \sum_{s=1}^t -(1-w_s)f_s(\theta_s) + (1-w_s)f_s(\theta_\star) \right) \middle| \mathcal{F}_{t-1} \right] \\
&= P_{t-1} \mathbb{E}_{y_t \sim \mathbb{P}_\star} \exp(-(1-w_t)f_t(\theta_t) + (1-w_t)f_t(\theta_\star)) \\
&= P_{t-1} \int_{y_t} \exp(-(1-w_t)f_t(\theta_t) + (1-w_t)f_t(\theta_\star)) \exp(-f_t(\theta_\star)) dy_t \\
&= P_{t-1} \int_{y_t} \exp(-(1-w_t)f_t(\theta_t) - w_t f_t(\theta_\star)) dy_t \\
&= P_{t-1} \int_{y_t} p_{\theta_t}(y_t | x_t)^{1-w_t} p_{\theta_\star}(y_t | x_t)^{w_t} dy_t \\
&= P_{t-1} \exp(-(1-w_t)D_{w_t}(\theta_\star, \theta_t)) \leq P_{t-1}.
\end{aligned}$$

We have used here the definition of the Renyi-divergence and the fact that it is always non-negative, namely

$$D_w(\theta_1, \theta_2) = \frac{1}{w-1} \log \int_y p_{\theta_1}(y|x)^{1-w} p_{\theta_2}(y|x)^w dy \geq 0,$$

for  $0 < w \neq 1$ .<sup>3</sup> The rest follows by the application of Ville's inequality.  $\square$

#### C.4 FTRL Proof: the Unweighted Case

*Proof of Theorem 4 (first part).* We define the function that FTRL minimizes in each step (to pick  $\hat{\theta}_t$ ) as,  $g_t(\theta) = \sum_{s=1}^{t-1} -\log p_\theta(y_s | x_s) + \lambda \|\theta\|_2^2$ . We can rewrite this objective as

$$\hat{\theta}_t = \arg \min_{\theta \in \Theta} g_t(\theta) = \arg \min_{\theta \in \Theta} \sum_{s=1}^{t-1} f_s(\theta) + \lambda \|\theta\|_2^2 = \arg \min_{\theta \in \Theta} \sum_{s=1}^{t-1} m_s(\theta) + \phi_t(\theta),$$

where we recall that<sup>4</sup>

$$f_s(\theta) = A(x_s^\top \theta) - T(y_s) x_s^\top \theta - \log h(y_s),$$

and we have introduced the shorthands

$$m_s(\theta) = -T(y_s) x_s^\top \theta$$

and

$$\phi_t(\theta) = \sum_{s=1}^{t-1} A(\theta^\top x_s) + \lambda \|\theta\|_2^2.$$

In essence, we have shifted some of the objective into what is commonly looked at as the regularizer. By a standard telescoping sum argument, we obtain for any  $u$

$$\begin{aligned}
&\sum_{s=1}^t (m_s(\hat{\theta}_s) - m_s(u)) \\
&= \phi_{t+1}(u) - \min_{\theta} \phi_1(\theta) + \sum_{s=1}^t [g_s(\hat{\theta}_s) - g_{s+1}(\hat{\theta}_{s+1}) + m_s(\hat{\theta}_s)] + \underbrace{g_{t+1}(\hat{\theta}_{t+1}) - g_{t+1}(u)}_{\leq 0} \\
&\leq \phi_{t+1}(u) + \sum_{s=1}^t [g_s(\hat{\theta}_s) - g_{s+1}(\hat{\theta}_{s+1}) + m_s(\hat{\theta}_s)] \\
&= \phi_{t+1}(u) + \sum_{s=1}^t [g_s(\hat{\theta}_s) - g_{s+1}(\hat{\theta}_{s+1}) + g_{s+1}(\hat{\theta}_s) - \phi_{s+1}(\hat{\theta}_s) - g_s(\hat{\theta}_s) + \phi_s(\hat{\theta}_s)]
\end{aligned}$$

<sup>3</sup>The case  $w_t = 1$  is trivial for us.

<sup>4</sup>The  $\log h(y_s)$  term does not play any role in the regret nor the FTRL objective.

$$= \phi_{t+1}(u) + \sum_{s=1}^t [g_{s+1}(\hat{\theta}_s) - g_{s+1}(\hat{\theta}_{s+1}) - \phi_{s+1}(\hat{\theta}_s) + \phi_s(\hat{\theta}_s)].$$

Now we use the strong-convexity of  $g_{s+1}$  under the norm  $\|\cdot\|_{\mathbf{V}_s^{\mu;\lambda}}$  where  $\mathbf{V}_s^{\mu;\lambda} = \sum_{j=1}^s \mu x_j x_j^\top + \lambda \mathbf{I}$ ,

$$\begin{aligned} & \sum_{s=1}^t (m_s(\hat{\theta}_s) - m_s(u)) \\ & \leq \phi_{t+1}(u) + \sum_{s=1}^t [(\hat{\theta}_s - \hat{\theta}_{s+1})^\top \nabla g_{s+1}(\hat{\theta}_s) \\ & \quad - \frac{1}{2}(\hat{\theta}_s - \hat{\theta}_{s+1})^\top \mathbf{V}_s^{\mu;\lambda}(\hat{\theta}_s - \hat{\theta}_{s+1}) - \phi_{s+1}(\hat{\theta}_s) + \phi_s(\hat{\theta}_s)] \\ & \leq \phi_{t+1}(u) + \sum_{s=1}^t [(\hat{\theta}_s - \hat{\theta}_{s+1})^\top \nabla f_s(\hat{\theta}_s) \\ & \quad - \frac{1}{2}(\hat{\theta}_s - \hat{\theta}_{s+1})^\top \mathbf{V}_s^{\mu;\lambda}(\hat{\theta}_s - \hat{\theta}_{s+1}) - \phi_{s+1}(\hat{\theta}_s) + \phi_s(\hat{\theta}_s)] \\ & \leq \phi_{t+1}(u) + \sum_{s=1}^t \left[ \frac{1}{2} \|\nabla f_s(\hat{\theta}_s)\|_{(\mathbf{V}_s^{\mu;\lambda})^{-1}}^2 - \phi_{s+1}(\hat{\theta}_s) + \phi_s(\hat{\theta}_s) \right], \end{aligned}$$

where in the second inequality we used that  $\nabla g_s(\hat{\theta}_s)^\top (x - \hat{\theta}_s) \geq 0$  due to the first-order optimality conditions for convex constrained minimization. Lastly, we optimized the resulting quadratic function over  $\hat{\theta}_{s+1}$  (over  $\mathbb{R}^d$ ) to get a worst case bound involving the dual-norm.

Note that for the shorthands we defined above:

$$\sum_{s=1}^t [-\phi_{s+1}(\hat{\theta}_s) + \phi_s(\hat{\theta}_s)] = \sum_{s=1}^t -A(\hat{\theta}_s^\top x_s).$$

Using our previous observations and the definition of  $\phi_{t+1}(\theta_*)$ , we get for the overall regret:

$$\begin{aligned} & \mathcal{R}_t \\ & = \sum_{s=1}^t f_s(\hat{\theta}_s) - f_s(\theta_*) \\ & = \sum_{s=1}^t m_s(\hat{\theta}_s) - m_s(\theta_*) + \sum_{s=1}^t A(x_s^\top \hat{\theta}_s) - A(x_s^\top \theta_*) \\ & = \sum_{s=1}^t A(x_s^\top \theta_*) - A(x_s^\top \hat{\theta}_s) + \sum_{s=1}^t A(x_s^\top \hat{\theta}_s) - A(x_s^\top \theta_*) + \frac{1}{2} \sum_{s=1}^t \|\nabla f_s(\hat{\theta}_s)\|_{(\mathbf{V}_s^{\mu;\lambda})^{-1}}^2 + \lambda \|\theta_*\|^2 \\ & \leq \frac{1}{2} \sum_{s=1}^t \|\nabla f_s(\hat{\theta}_s)\|_{(\mathbf{V}_s^{\mu;\lambda})^{-1}}^2 + \lambda \|\theta_*\|^2 \\ & \leq \frac{1}{2} \sum_{s=1}^t \|T(y_s)x_s - A'(x_s^\top \hat{\theta}_s)x_s\|_{(\mathbf{V}_s^{\mu;\lambda})^{-1}}^2 + \lambda \|\theta_*\|^2 \\ & \leq \sum_{s=1}^t \left[ \|T(y_s)x_s - A'(x_s^\top \theta_*)x_s\|_{(\mathbf{V}_s^{\mu;\lambda})^{-1}}^2 + \|(A'(x_s^\top \hat{\theta}_s) - A'(x_s^\top \theta_*))x_s\|_{(\mathbf{V}_s^{\mu;\lambda})^{-1}}^2 \right] + \lambda \|\theta_*\|^2 \\ & \leq \sum_{s=1}^t \left[ \|T(y_s)x_s - A'(x_s^\top \theta_*)x_s\|_{(\mathbf{V}_s^{\mu;\lambda})^{-1}}^2 + 2L^2 B^2 \|x_s\|_{(\mathbf{V}_s^{\mu;\lambda})^{-1}}^2 \right] + \lambda \|\theta_*\|^2 \\ & \leq \sum_{s=1}^t \left[ \|\nabla f_s(\theta_*)\|_{(\mathbf{V}_s^{\mu;\lambda})^{-1}}^2 + 2L^2 B^2 \|x_s\|_{(\mathbf{V}_s^{\mu;\lambda})^{-1}}^2 \right] + \lambda B^2 \end{aligned}$$

$$\begin{aligned}
&\leq \lambda B^2 + \frac{L}{\mu} \left( \gamma_t^\lambda + \log \left( \frac{1}{\delta} \right) \right) + \sum_{s=1}^t 2L^2 B^2 \|x_s\|_{(\mathbf{V}_s^{\mu;\lambda})^{-1}}^2 \\
&\leq \lambda B^2 + \frac{L}{\mu} \left( \gamma_t^\lambda + \log \left( \frac{1}{\delta} \right) \right) + \frac{2L^2 B^2}{\mu} \gamma_t^\lambda.
\end{aligned}$$

The last line follows because of Lemma 7, and the second to last one follows because of Lemma 10. Notice that if we wish to deal with arbitrary weights  $\{w_t\}$ , we can simply resort to Lemma 11 and bound the weighted case with the unweighted case. In that case, we incur an additional additive  $\log(1/\delta)$  term.  $\square$

### C.5 FTRL Analysis: the Weighted Case (Vovk-Azoury-Warmuth Forecaster)

*Proof.* We define the function that FTRL minimizes in each step (to pick  $\hat{\theta}_t$ ) as  $\tilde{g}_t(\theta) = \sum_{s=1}^{t-1} [A(x_s^\top \theta) - T(y_s)x_s^\top \theta] + \psi_t(\theta)$  with  $\psi_t(\theta) = A(x_t^\top \theta) + \lambda \|\theta\|_2^2$ . We can rewrite this objective as

$$\hat{\theta}_t = \arg \min_{\theta \in \Theta} \tilde{g}_t(\theta) = \arg \min_{\theta \in \Theta} \sum_{s=1}^{t-1} m_s(\theta) + \phi_t(\theta),$$

by introducing the shorthands

$$m_s(\theta) = -T(y_s)x_s^\top \theta$$

and (notice the difference in time index of the second sum when compared to the proof in the previous subsection):

$$\phi_t(\theta) = \sum_{s=1}^t A(\theta^\top x_s) + \lambda \|\theta\|_2^2.$$

In addition consider the objective  $g_t$  from the classical FTRL analysis in Section C.4. It is not used to run the online algorithm, but is helpful in our analysis. With our new components, it is equal to

$$g_t(\theta) = \sum_{s=1}^{t-1} m_s(\theta) + \sum_{s=1}^{t-1} A(\theta^\top x_s) + \lambda \|\theta\|_2^2 = \sum_{s=1}^{t-1} m_s(\theta) + \phi_{t-1}(\theta),$$

and its minimizer is  $\bar{\theta}_t = \arg \min_{\theta \in \Theta} g_t(\theta)$ . Also, consider a weighted version of the regularizer

$$\bar{\phi}_t(\theta) = \sum_{s=1}^t w_s A(x_s^\top \theta) + \lambda \|\theta\|_2^2,$$

which will be useful. We use a variant of a similar telescoping sum argument as in the previous proof of Section C.4. We specifically use  $\theta_\star$  as the comparator to compete against. Notice that we insert a telescoping sum involving the objective  $g_s$ , which is not the objective that our estimator is minimizing:

$$\begin{aligned}
&\sum_{s=1}^t w_s (m_s(\hat{\theta}_s) - m_s(\theta_\star)) \\
&\stackrel{(*)}{=} \bar{\phi}_t(\theta_\star) - \phi_0(\bar{\theta}_1) + \sum_{s=1}^t [g_s(\bar{\theta}_s) - g_{s+1}(\bar{\theta}_{s+1}) + w_s m_s(\hat{\theta}_s)] + \underbrace{g_{t+1}(\bar{\theta}_{t+1}) - g_{t+1}(\theta_\star)}_{\leq 0} \\
&\quad + \sum_{s=1}^t (1 - w_s) f_s(\theta_\star) \\
&\leq \bar{\phi}_t(\theta_\star) + \sum_{s=1}^t [w_s (g_s(\bar{\theta}_s) - g_{s+1}(\bar{\theta}_{s+1})) + w_s m_s(\hat{\theta}_s)] \\
&\quad + \sum_{s=1}^t (1 - w_s) (g_s(\bar{\theta}_s) - g_{s+1}(\bar{\theta}_{s+1}) + f_s(\theta_\star))
\end{aligned}$$

$$\begin{aligned}
&\stackrel{(**)}{=} \bar{\phi}_t(\theta_\star) + \sum_{s=1}^t w_s [g_s(\bar{\theta}_s) - g_{s+1}(\bar{\theta}_{s+1}) + g_{s+1}(\hat{\theta}_s) - \phi_s(\hat{\theta}_s) - g_s(\hat{\theta}_s) + \phi_{s-1}(\hat{\theta}_s)] \\
&\quad \sum_{s=1}^t (1-w_s) [g_s(\bar{\theta}_s) - g_{s+1}(\bar{\theta}_{s+1}) + f_s(\theta_\star)] \\
&\stackrel{(***)}{\leq} \bar{\phi}_t(\theta_\star) + \sum_{s=1}^t w_s [-g_{s+1}(\bar{\theta}_{s+1}) + g_{s+1}(\hat{\theta}_s) - \phi_s(\hat{\theta}_s) + \phi_{s-1}(\hat{\theta}_s)] + \tilde{\Delta}_t.
\end{aligned}$$

In (\*), we used the shorthands and definitions introduced above. In (\*\*), we used the identity  $g_{s+1}(\hat{\theta}_s) - \phi_s(\hat{\theta}_s) - g_s(\hat{\theta}_s) + \phi_{s-1}(\hat{\theta}_s) = m_s(\hat{\theta}_s)$ . Finally, for (\*\*\*), recall that  $\bar{\theta}_s$  is the minimizer of  $g_s$ , and hence,  $g_s(\bar{\theta}_s) - g_s(\hat{\theta}_s) \leq 0$ . Next, define  $\tilde{\Delta}_t = \sum_{s=1}^t (1-w_s) [g_s(\bar{\theta}_s) - g_{s+1}(\bar{\theta}_{s+1}) + f_s(\theta_\star)]$ . We will bound this term later.

Now, we use the strong-convexity of  $g_{s+1}(\theta)$  under the norm  $\|\cdot\|_{\mathbf{V}_s^{\mu;\lambda}}$  where  $\mathbf{V}_s^{\mu;\lambda} = \sum_{j=1}^s \mu x_j x_j^\top + \lambda \mathbf{I}$ , namely

$$g_{s+1}(\bar{\theta}_{s+1}) \geq g_{s+1}(\hat{\theta}_{s+1}) + \nabla g_{s+1}(\hat{\theta}_s)^\top (\bar{\theta}_{s+1} - \hat{\theta}_s) + \frac{1}{2} \|\bar{\theta}_{s+1} - \hat{\theta}_s\|_{\mathbf{V}_s^{\mu;\lambda}}^2.$$

We can then proceed as follows:

$$\begin{aligned}
&\sum_{s=1}^t w_s (m_s(\hat{\theta}_s) - m_s(u)) \\
&\leq \tilde{\Delta}_t + \bar{\phi}_t(\theta_\star) + \sum_{s=1}^t w_s [\nabla g_{s+1}(\hat{\theta}_s)^\top (\hat{\theta}_s - \bar{\theta}_{s+1}) \\
&\quad - \frac{1}{2} \|\bar{\theta}_{s+1} - \hat{\theta}_s\|_{\mathbf{V}_s^{\mu;\lambda}}^2 - \phi_s(\hat{\theta}_s) + \phi_{s-1}(\hat{\theta}_s)] \\
&\leq \tilde{\Delta}_t + \bar{\phi}_t(\theta_\star) + \sum_{s=1}^t w_s [(\nabla \tilde{g}_s(\hat{\theta}_s) + \nabla m_s(\hat{\theta}_s))^\top (\hat{\theta}_s - \bar{\theta}_{s+1}) \\
&\quad - \frac{1}{2} \|\bar{\theta}_{s+1} - \hat{\theta}_s\|_{\mathbf{V}_s^{\mu;\lambda}}^2 - \phi_s(\hat{\theta}_s) + \phi_{s-1}(\hat{\theta}_s)] \\
&\leq \tilde{\Delta}_t + \bar{\phi}_t(\theta_\star) + \frac{1}{2} \sum_{s=1}^t w_s \|\nabla m_s(\hat{\theta}_s)\|_{(\mathbf{V}_s^\mu)^{-1}}^2 - w_s (\phi_s(\hat{\theta}_s) + \phi_{s-1}(\hat{\theta}_s)), \tag{28}
\end{aligned}$$

where in the second to last line we used that  $\nabla \tilde{g}_s(\hat{\theta}_s)^\top (x - \hat{\theta}_s) \geq 0$  for any  $x$ , due to the optimality of  $\hat{\theta}_s$  for the FTRL objective. In the last line, we optimized over  $\bar{\theta}_{s+1}$  to get a worst-case bound on the quadratic function involving it. Also, note that for the shorthands we defined above:

$$\sum_{s=1}^t w_s (-\phi_s(\hat{\theta}_s) + \phi_{s-1}(\hat{\theta}_s)) = \sum_{s=1}^t w_s (-A(\hat{\theta}_s^\top x_s)).$$

Our goal here is to upper bound the overall regret:

$$\begin{aligned}
\mathcal{R}_t &= \sum_{s=1}^t w_s (f_s(\hat{\theta}_s) - f_s(\theta_\star)) \\
&= \sum_{s=1}^t w_s (m_s(\hat{\theta}_s) - m_s(\theta_\star)) + \sum_{s=1}^t w_s (A(x_s^\top \hat{\theta}_s) - A(x_s^\top \theta_\star)) \\
&\leq \bar{\phi}_t(\theta_\star) - \sum_{s=1}^t w_s A(x_s^\top \hat{\theta}_s) + \frac{1}{2} \sum_{s=1}^t w_s \|\nabla m_s(\hat{\theta}_s)\|_{(\mathbf{V}_s^\mu)^{-1}}^2 + \tilde{\Delta}_t \\
&\quad + \sum_{s=1}^t w_s (A(x_s^\top \hat{\theta}_s) - A(x_s^\top \theta_\star))
\end{aligned}$$



$$\begin{aligned}
&= \sum_{s=1}^t w_s (A(x_s^\top \theta_\star)) - \sum_{s=1}^t w_s A(x_s^\top \hat{\theta}_s) + \frac{1}{2} \sum_{s=1}^t w_s \|\nabla m_s(\hat{\theta}_s)\|_{(\mathbf{V}_s^\mu)^{-1}}^2 + \lambda \|\theta_\star\|^2 + \tilde{\Delta}_t \\
&\quad + \sum_{s=1}^t w_s (A(x_s^\top \hat{\theta}_s) - A(x_s^\top \theta_\star)) \\
&= \tilde{\Delta}_t + \frac{1}{2} \sum_{s=1}^t w_s \|\nabla m_s(\hat{\theta}_s)\|_{(\mathbf{V}_s^\mu)^{-1}}^2 + \lambda \|\theta_\star\|^2. \tag{29}
\end{aligned}$$

The first inequality follows by plugging in (28).

We return back to the term  $\tilde{\Delta}_t$ ,

$$\begin{aligned}
\tilde{\Delta}_t &= \sum_{s=1}^t (1 - w_s) [g_s(\bar{\theta}_s) - g_{s+1}(\bar{\theta}_{s+1}) + f_s(\theta_\star)] \\
&= \sum_{s=1}^t (1 - w_s) [g_s(\bar{\theta}_s) - g_s(\bar{\theta}_{s+1}) - f_s(\bar{\theta}_{s+1}) + f_s(\theta_\star)] \\
&\leq \sum_{s=1}^t (1 - w_s) (f_s(\theta_\star) - f_s(\bar{\theta}_{s+1})) \\
&\leq \frac{L}{\mu} (\gamma_t^\lambda + \log(1/\delta)), \tag{30}
\end{aligned}$$

where the second to last line is by the optimality of  $\bar{\theta}_s$  and the last one by the assumption in the theorem.

Carrying on with the analysis, i.e. with (29), we insert the definition of  $m_s$  and obtain

$$\begin{aligned}
&\mathcal{R}_t \\
&\leq \Delta_t + \frac{1}{2} \sum_{s=1}^t w_s \|-T(y_s)x_s\|_{(\mathbf{V}_s^{\mu;\lambda})^{-1}}^2 + \lambda \|\theta_\star\|^2 \\
&\leq \Delta_t + \sum_{s=1}^t \left[ w_s \|T(y_s)x_s - A'(x_s^\top \theta_\star)x_s\|_{(\mathbf{V}_s^{\mu;\lambda})^{-1}}^2 + w_s \|(A'(x_s^\top \theta_\star))x_s\|_{(\mathbf{V}_s^{\mu;\lambda})^{-1}}^2 \right] + \lambda \|\theta_\star\|^2 \\
&\stackrel{(*)}{\leq} \Delta_t + \sum_{s=1}^t \|T(y_s)x_s - A'(x_s^\top \theta_\star)x_s\|_{(\mathbf{V}_s^{\mu;\lambda})^{-1}}^2 + L^2 B^2 \sum_{s=1}^t w_s \|x_s\|_{(\mathbf{V}_s^{\mu;\lambda})^{-1}}^2 + \lambda \|\theta_\star\|^2 \\
&\stackrel{(**)}{=} \Delta_t + \sum_{s=1}^t \|\nabla f_s(\theta_\star)\|_{(\mathbf{V}_s^\mu)^{-1}}^2 + L \sum_{s=1}^t \frac{B^2}{1/L + \text{bias}_{x_s}^2(\hat{\theta}_s)} \|x_s\|_{(\mathbf{V}_s^{\mu;\lambda})^{-1}}^2 + \lambda B^2 \\
&\stackrel{(***)}{\leq} \lambda B^2 + \frac{2L}{\mu} \left( \gamma_t^\lambda + \log\left(\frac{1}{\delta}\right) \right) + L \sum_{s=1}^t \frac{B^2}{1/L + \text{bias}_{x_s}^2(\hat{\theta}_s)} \|x_s\|_{(\mathbf{V}_s^{\mu;\lambda})^{-1}}^2.
\end{aligned}$$

In (\*), we use  $w_s \leq 1$  and the Lipschitzness of  $A'$ . In (\*\*), we use the definition of the weights. Finally, in (\*\*\*), we used Lemma 10 and (30). By substituting  $\Delta\gamma_s = \mu \|x_s\|_{(\mathbf{V}_s^{\mu;\lambda})^{-1}}^2$ , we finish the proof. The event in Lemma 10 holds with probability  $1 - \delta$ , completing the proof.  $\square$

## C.6 FTRL Analysis: Beyond Global Smoothness

In this subsection, we give alternative analysis which avoids the necessity to impose a global smoothness condition our likelihood; instead strong convexity within a bounded domain suffices, and we will only assume that

$$\epsilon_s := \mathbb{E}_{x_s^\top \theta_\star} [T(y_s)] - T(y_s)$$

are sub-Exponential random variables, setting us apart from Zhao et al. (2022) which assume sub-Gaussianity. In particular, we can show the following theorem

**Theorem 7.** *With probability  $1 - \delta$ , uniformly over time  $t \in \mathbb{N}$ , we have*

$$\mathcal{R}_t \leq cd \log^2(t/\delta) \log(t),$$

where the universal constant  $c$  hides all constants independent of  $t, d$  and  $\delta$ .

### C.6.1 Lemmas

We state the following result on sub-Exponential random variables.

**Proposition 2** (Theorem 2.13 in [Wainwright \(2019\)](#)). *If  $X$  is a centered sub-Exponential variable with some finite variance proxy, then there exist constants  $c_1, c_2 > 0$  such that for any  $t > 0$*

$$\mathbb{P}(|X| \geq a) \leq c_1 e^{-c_2 a}.$$

By some careful union bounds (akin to a stitching argument), we can also provide upper bounds on anytime-valid upper bounds on the process  $S_t = \max_{s \leq t} \epsilon_s$ .

**Lemma 12.** *For any sequence  $(\epsilon_s)_{s=1}^\infty$  of sub-Exponential-variables, there exists a constant  $\tilde{c}$  independent of  $t$  such that*

$$\mathbb{P}(\exists t : \max_{s \leq t} |\epsilon_s| \geq \tilde{c} \log(s/\delta)) \leq \delta.$$

*Proof.* (of Lemma 12) By Proposition 2, there exists  $c_1, c_2 > 0$  such that we have  $\mathbb{P}(|\epsilon_s| \geq a) \leq c_1 e^{-c_2 a}$  for any fixed  $s$ . Note that  $c_1 e^{-c_2 a} \leq \delta$  is satisfied for  $a \geq \frac{1}{c_2} \log(c_1/\delta) =: c_3 \log(c_1/\delta)$ . Let us denote by  $\mathcal{E}_i$  the event all  $j \in [2^i, 2^{i+1}) \cap \mathbb{N}$  satisfy the inequality

$$|\epsilon_j| < c_3 \log(c_1(2^{2i+1})/\delta).$$

For a single  $j$ , this happens with probability at least  $1 - \delta/2^{2i+1}$ . Therefore, by a union bound, as  $|[2^i, 2^{i+1}) \cap \mathbb{N}| = 2^i$ , we can bound the probability of the complement, namely  $\mathbb{P}(\mathcal{E}_i^c) \leq 2^i \frac{\delta}{2^{2i+1}} = \frac{\delta}{2^{i+1}}$ . Now, by another union bound, we can conclude that

$$\mathbb{P}(\cup_{i=0}^\infty \mathcal{E}_i^c) \leq \sum_{i=0}^\infty \frac{\delta}{2^{i+1}} = \frac{\delta}{2} \frac{1}{1 - \frac{1}{2}} = \delta.$$

Now we also have for any  $j$  in this range that  $2^{2i+1} \leq 2j^2$ , and therefore, if  $\mathcal{E}_i$  holds, we have for any  $j \in [2^i, 2^{i+1}) \cap \mathbb{N}$ :

$$\epsilon_j \leq c_3 \log(c_1(2j^2)/\delta) \leq 2c_3 \log(2c_1 j/\delta) \leq \tilde{c} \log(j/\delta).$$

We can immediately see that this implies

$$\mathbb{P}(\exists t : \max_{s \leq t} |\epsilon_s| \geq \tilde{c} \log(s/\delta)) \leq \delta.$$

as desired. □

### C.6.2 Proof of Theorem 7

Our proof initially follows the FTRL regret bound proofs in the adversarial setting [Hazan \(2016\)](#); [Orabona \(2019\)](#). It also has overlap with the proof in [Zhao et al. \(2022\)](#). We define the function that FTRL minimizes in each step as (to pick  $\hat{\theta}_t$ )

$$g_t(\theta) = \sum_{s=1}^{t-1} -\log p_\theta(y_s | x_s) + \phi(\theta)$$

for convenience. We initially use the same steps as in Theorem 4 to see that for any  $u \in \Theta$

$$\begin{aligned} & \sum_{s=1}^t (f_s(\hat{\theta}_s) - f_s(u)) \\ & \leq \phi(u) - \min_{\theta} \phi(\theta) + \sum_{s=1}^t [g_s(\hat{\theta}_s) - g_{s+1}(\hat{\theta}_{s+1}) + f_s(\hat{\theta}_s)] + g_{t+1}(\hat{\theta}_{t+1}) - g_{t+1}(u) \end{aligned}$$

$$\leq \lambda B^2 + \sum_{s=1}^t [g_s(\hat{\theta}_s) - g_{s+1}(\hat{\theta}_{s+1}) + f_s(\hat{\theta}_s)].$$

Similarly to the proof of Theorem 4 in Appendix C.4 we bound these increments by the dual norm of the gradient of the objective.

$$g_s(\hat{\theta}_s) - g_{s+1}(\hat{\theta}_{s+1}) + f_s(\hat{\theta}_s) \leq \frac{\|\nabla f_s(\hat{\theta}_s)\|_{(\mathbf{V}_s^{\mu;\lambda})^{-1}}^2}{2}. \quad (31)$$

Now we note that

$$\nabla f_s(\theta) = A'(x_s^\top \theta) x_s - T(y_s) x_s.$$

Using properties of the exponential family, we deduce that

$$\begin{aligned} \nabla f_s(\theta) &= \left( \mathbb{E}_{x_s^\top \theta} [T(y_s)] - T(y_s) \right) x_s \\ &= \left( \mathbb{E}_{x_s^\top \theta} [T(y_s)] - \mathbb{E}_{x_s^\top \theta_*} [T(y_s)] + \mathbb{E}_{x_s^\top \theta_*} [T(y_s)] - T(y_s) \right) x_s. \end{aligned}$$

From here on out, we proceed more crudely than in our previous analyses, since we are only concerned with asymptotic behavior when  $d$  and  $t$  are large. Let us define

$$U := \sup_{\theta \in \Theta} \left| \mathbb{E}_{x_s^\top \theta} [T(y_s)] - \mathbb{E}_{x_s^\top \theta_*} [T(y_s)] \right|,$$

which is a model-dependent, deterministic quantity. Let us define the noise variables

$$\epsilon_s := \mathbb{E}_{x_s^\top \theta_*} [T(y_s)] - T(y_s).$$

We bound

$$\|\nabla f_s(\theta)\|_{(\mathbf{V}_s^{\mu;\lambda})^{-1}}^2 \leq 2(U^2 + \epsilon_s^2) \|x_s\|_{(\mathbf{V}_s^{\mu;\lambda})^{-1}}^2.$$

Note that the  $\epsilon_s$  are centered, independent sub-Exponential variables, and as such are guaranteed to satisfy

$$\mathbb{P}(\exists t : \max_{s \leq t} |\epsilon_s| \geq c_4 \log(s/\delta)) \leq \delta,$$

by Lemma 12. This tells us that conditional on this event, we can upper bound for any  $t$

$$\mathcal{R}_t \leq \lambda B^2 + \frac{1}{\mu} (U^2 + \tilde{c}^2 \log^2(t/\delta)) \sum_{s=1}^t \|x_s\|_{(\mathbf{V}_s^{\mu;\lambda})^{-1}}^2.$$

for some constant  $\tilde{c}$  independent of  $t$ . By Lemma 7, there is thus a constant  $c'$  independent of  $t$  and  $d$  such that

$$\mathcal{R}_t \leq c' d \log^2(t/\delta) \log(t) = \mathcal{O}(d \log^3(t)),$$

with probability  $1 - \delta$  uniformly over  $t \in \mathbb{N}$ .

## D Regret Consequences for Stochastic Linear Bandits

As a corollary of our analysis, we provide the regret for stochastic linear bandits that use our confidence sets within the LinUCB algorithm.

*Proof.* We proceed in two parts: first, we instantiate Theorem 3 and then we follow the classical regret analysis for stochastic linear bandits.

**Specializing the Bregman divergence results** By Theorem 3, we know that for any  $\nu > 0$ , we have that with probability  $1 - \delta$ ,

$$D_{Z_t^\nu}(\theta, \theta_\star) \leq \frac{4L}{\mu} \xi_t + 2 \log\left(\frac{1}{\delta}\right) + 2\mathcal{R}_t, \quad (32)$$

for all  $t$ , where  $\xi_t = (\log(\frac{1}{\alpha}) + \nu B^2 + \Gamma_t^\nu)$  and  $\mathcal{R}_t$  is the online convex optimization regret. We also recall that

$$Z_t^\nu(\theta) = \sum_{s=1}^t w_s A(x_s^\top \theta) + \frac{\nu}{2} \|\theta\|_2^2.$$

In the Gaussian case, where  $A(z) = z^2/(2\sigma^2)$ . This implies that

$$\nabla Z_t^\nu(\theta) = \sum_{s=1}^t \frac{w_s}{\sigma^2} x_s x_s^\top \theta + \nu \theta = \mathbf{W}_t^{\sigma^{-2}; \nu} \theta,$$

where we have defined a weighted version of  $\mathbf{V}_t^{\sigma^{-2}; \nu}$  as  $\mathbf{W}_t^{\sigma^{-2}; \nu} = \sum_{s=1}^t \frac{w_s x_s x_s^\top}{\sigma^2} + \nu \mathbf{I}$  and therefore the Bregman divergence is given by  $D_{Z_t^\nu}(\theta, \theta_\star) = \frac{1}{2} \|\theta - \theta_\star\|_{\mathbf{W}_t^{\sigma^{-2}; \nu}}^2$ . We can also see that the Bregman information gain is given by

$$\Gamma_t^\nu = \log\left(\frac{\int_{\mathbb{R}^d} \exp(-\frac{1}{2} \|\theta\|_2^2) d\theta}{\int_{\mathbb{R}^d} \exp(-D_{Z_t^\nu}(\theta, \tilde{\theta}_t)) d\theta}\right) = \log\left(\frac{\int_{\mathbb{R}^d} \exp(-\frac{1}{2} \|\theta\|_2^2) d\theta}{\int_{\mathbb{R}^d} \exp(-\frac{1}{2} \|\theta - \theta_\star\|_{\mathbf{W}_t^{\sigma^{-2}; \nu}}^2) d\theta}\right).$$

These Gaussian integrals are straightforward to evaluate. We know that

$$\int_{\mathbb{R}^d} \exp\left(-\frac{1}{2} \|\theta\|_2^2\right) d\theta = (2\pi)^{d/2} \sqrt{\det((\nu \mathbf{I}_d)^{-1})}.$$

Similarly,

$$\int_{\mathbb{R}^d} \exp\left(-\frac{1}{2} \|\theta - \theta_\star\|_{\mathbf{W}_t^{\sigma^{-2}; \nu}}^2\right) d\theta = (2\pi)^{d/2} \sqrt{\det((\mathbf{W}_t^{\sigma^{-2}; \nu})^{-1})}.$$

Then, we can compute

$$\Gamma_t^\nu = \log\left(\frac{\det(\mathbf{W}_t^{\sigma^{-2}; \nu})}{\det(\nu \mathbf{I})}\right) = \log\left(\det\left(\sum_{s=1}^t \frac{w_s x_s x_s^\top}{\sigma^2 \nu} + \mathbf{I}\right)\right).$$

In the unweighted case, with which we proceed, we have  $\Gamma_t^\nu = \gamma_t^\nu$ , that is we recover the classical upper bound on the information gain (Srinivas et al., 2009). To summarize, we have specialized the bound (32) to say that for any  $\theta \in \mathcal{C}_t$ , we have (since  $L = \mu = 1/\sigma^2$ )

$$\|\theta - \theta_\star\|_{\mathbf{V}_t^{\sigma^{-2}; \nu}}^2 \leq 8 (\log(1/\alpha) + \nu B^2 + \gamma_t^\nu) + 4 \log(3/\delta) + 4\mathcal{R}_t.$$

Now, we instantiate the regret of the online learner using Theorem 4. With probability  $1 - \delta$ , uniformly over  $t$ , we have

$$\mathcal{R}_t \leq \lambda B^2 + \frac{L}{\mu} \left(\gamma_t^\lambda + \log\left(\frac{1}{\delta}\right)\right) + \frac{2L^2 B^2}{\mu} \gamma_t^\lambda. \quad (33)$$

We get by choosing  $\alpha = \delta$ , and setting  $\nu = \lambda$  that

$$\begin{aligned} & \|\theta - \theta_\star\|_{\mathbf{V}_t^{\sigma^{-2}; \lambda}}^2 \\ & \leq 8 (\log(1/\delta) + \nu B^2 + \gamma_t^\nu) + 4 \log(1/\delta) + 4\lambda B^2 + 4(\gamma_t^\lambda + \log(1/\delta)) + \frac{8B^2}{\sigma^2} \gamma_t^\lambda \\ & \leq 16 \log(1/\delta) + 12\lambda B^2 + 8 \left(\frac{B^2}{\sigma^2} + 1\right) \gamma_t^\lambda =: \beta_t. \end{aligned}$$

**Linear bandit regret analysis** We are ready to proceed with the bandit analysis for the UCB Algorithm. We follow [Lattimore and Szepesvári \(2020\)](#) and bound the pseudo-regret, letting  $x_*$  be the optimal action. We bound the instantaneous regret at step  $0 \leq s \leq t$  as

$$\begin{aligned}
r_s &= \langle \theta_*, x_* - x_s \rangle \\
&= \langle \theta_*, x_* \rangle - \langle \theta_*, x_s \rangle \\
&\leq \max_{\theta \in \mathcal{C}_{s-1}} \langle \theta, x_* \rangle - \langle \theta_*, x_s \rangle \\
&\leq \max_{x \in \mathcal{X}} \max_{\theta \in \mathcal{C}_{s-1}} \langle \theta, x \rangle - \langle \theta_*, x_s \rangle \\
&\stackrel{(*)}{=} \max_{\theta \in \mathcal{C}_{s-1}} \langle \theta, x_s \rangle - \langle \theta_*, x_s \rangle \\
&\stackrel{(**)}{\leq} \|\tilde{\theta}_s - \theta_*\|_{\mathbf{V}_{s-1}^{\sigma^{-2}; \lambda}} \|x_s\|_{(\mathbf{V}_{s-1}^{\sigma^{-2}; \lambda})^{-1}} \\
&\leq \sqrt{\beta_{s-1}} \|x_s\|_{(\mathbf{V}_{s-1}^{\sigma^{-2}; \lambda})^{-1}} \\
&\leq \sqrt{\beta_t} \|x_s\|_{(\mathbf{V}_{s-1}^{\sigma^{-2}; \lambda})^{-1}}.
\end{aligned}$$

The first inequality replaces  $\theta_*$  by the upper confidence bound for action  $x_*$ , which is valid with probability  $1 - \alpha = 1 - \delta$  uniformly over time. Then,  $(*)$  uses the fact that  $x_t$  is chosen to maximize the upper confidence bound. Finally  $(**)$  defines the UCB parameter  $\tilde{\theta}_t$ . By Corollary 1, we have

$$\sum_{s=1}^t \|x_s\|_{(\mathbf{V}_{s-1}^{\sigma^{-2}; \lambda})^{-1}}^2 \leq 2\sigma^2 \gamma_t^\lambda.$$

Plugging all this together and using an  $\ell_1/\ell_2$ -norm inequality, we get

$$\begin{aligned}
\mathfrak{R}_t &= \sum_{s=1}^t r_s \\
&\leq \sqrt{\beta_t} \sqrt{t \sum_{s=1}^t \|x_s\|_{(\mathbf{V}_{s-1}^{\sigma^{-2}; \lambda})^{-1}}^2} \\
&\leq \sqrt{2t\beta_t\sigma^2\gamma_t^\lambda} \\
&\leq \sqrt{2\sigma^2 t (16 \log(1/\delta) + 12\lambda B^2 + 8\gamma_t^\lambda + 8B^2/\sigma^2\gamma_t^\lambda)\gamma_t^\lambda} \\
&\leq 6\sqrt{t\gamma_t} \left( \sigma \sqrt{\log(1/\delta)} + \gamma_t^\lambda + \sigma\lambda^{1/2}B + B\sqrt{\gamma_t^\lambda} \right).
\end{aligned}$$

To summarize and to justify why this bound holds with probability  $1 - 3\delta$  uniformly over time, note that we have bounded the probability of the FTRL bound (33) not holding for some  $t$  by  $\delta$ . Then, the probability of (32) not holding for some  $t$  is at most  $\delta$ . Finally, the anytime Type I error of our sets is also bounded by  $\delta$ . A union bound therefore concludes the proof.  $\square$

## D.1 Comparison to Abbasi-Yadkori et. al. (2011)

We compare our result to the one from [Abbasi-Yadkori et al. \(2011\)](#). Under the assumption that  $\lambda \geq 1$ , they show that the regret satisfies

$$\mathfrak{R}_t \leq 4\sqrt{td \log(\lambda + t/d)} \left( \sqrt{\lambda}B + \sigma\sqrt{2\log(1/\delta) + d \log(1 + t/(\lambda d))} \right).$$

Observe that there is a reparametrization for the regularizer to get even more similar bounds. If we take  $\lambda = \tilde{\lambda}/\sigma^2$  for some  $\tilde{\lambda} \geq 1$ , our bound reads as

$$6\sqrt{t\gamma_t^{\tilde{\lambda}/\sigma^2}} \left( \sigma\sqrt{\log(1/\delta)} + \gamma_t^{\tilde{\lambda}/\sigma^2} + \tilde{\lambda}^{1/2}B + B\sqrt{\gamma_t^{\tilde{\lambda}/\sigma^2}} \right).$$

Given that by Corollary 1, we have

$$\gamma_t^{\tilde{\lambda}/\sigma^2} \leq d \log \left( \frac{t}{\tilde{\lambda}} + 1 \right),$$

we get almost matching bounds, up to an additional  $B\sqrt{\gamma_t^{\tilde{\lambda}/\sigma^2}}$  term blowing up the regret, which we attribute to the accumulation of bias without the reweighting scheme. The remaining differences are down to using slightly different versions of the elliptical potential lemma, trading off generality and tightness (Abbasi-Yadkori et al., 2011; Hazan et al., 2006; Lattimore and Szepesvári, 2020).

## E Experimental Details

### E.1 Calibration Plots

In Figure 3 we report the calibration of heuristics as well as other theoretically motivated works. The other theoretically motivated works are very pessimistic and are not appropriately calibrated. Note that one caveat of reporting calibration is that it is very much influenced by the data collection scheme in the sequential regime. In our case we use a bandit algorithm to collect the data. Arguably, in this setting, regret might be a better measure rather than looking at the calibration of the confidence sets. Additionally, the calibration depends on the true value  $\theta_*$ . We report the results for zero parameter and a random parameter from a unit ball. We also report results for i.i.d. data.

### E.2 Baselines and Details

In the following section, we describe the baseline we used in the comparison. The details of parameters used in the experiments can be found in Table 2. As there is no explicit statement for sub-exponential variables, we give a formal derivation in Appendix F.

#### E.2.1 Sub-Exponential Random Variables: Confidence Sets

For this baseline, we will assume a linear model with additive sub-exponential noise, namely that there is  $\theta_* \in \Theta$  such that  $y_t = \langle \theta_*, x_t \rangle + \eta_t$ , where  $\eta_t$  is  $(\nu, \gamma)$ -conditionally sub-exponential (Wainwright, 2019). We let as usual

$$\mathbf{V}_t^{\nu^{-2}; \lambda} = \sum_{s=1}^t \frac{x_s x_s^\top}{\nu^2} \lambda \mathbf{I}.$$

With this, one can prove the following time-uniform concentration result:

**Proposition 3.** *For any  $k \in (0, 1)$ , the following holds:*

$$\mathbb{P} \left( \exists t : \|\hat{\theta}_t - \theta_*\|_{\mathbf{V}_t^{\nu^{-2}; \lambda}} \geq \sqrt{\beta_{SE}} \right) \leq \delta,$$

where

$$\sqrt{\beta_{SE}} = \sqrt{\lambda} \|\theta_*\|_2 + \sqrt{\lambda} k B + \frac{d}{\sqrt{\lambda} k B} \log \left( \frac{1}{1-k} \right) + \frac{1}{\sqrt{\lambda} k B} \log \left( \frac{(\det(\mathbf{V}_t^{\nu^{-2}; \lambda}))^{1/2}}{\delta \det(\sqrt{\lambda} \mathbf{I})} \right).$$

The proof is very similar to prior work Faury et al. (2020); Mutný and Krause (2021), and can be found in Appendix F. This can readily be applied in the survival analysis (after a suitable transformation explained in the main paper) and Laplace noise experiments.

#### E.2.2 Poisson Heuristics

We implement two heuristics. One is a Bayesian formulation due to the Laplace method, and one is due to considerations in Mutný and Krause (2021) of how to construct a valid confidence set using the Fisher information. The Laplace method creates an ellipsoidal confidence set using the second-order information evaluated at the penalized maximum likelihood estimator. Namely, the second derivative of the likelihood evaluated at the maximum penalized likelihood  $\hat{\theta}_t$  is

$$\mathbf{V}_{\text{laplace}} = \sum_{s=1}^t \exp(\hat{\theta}_t^\top x_s) x_s x_s^\top.$$

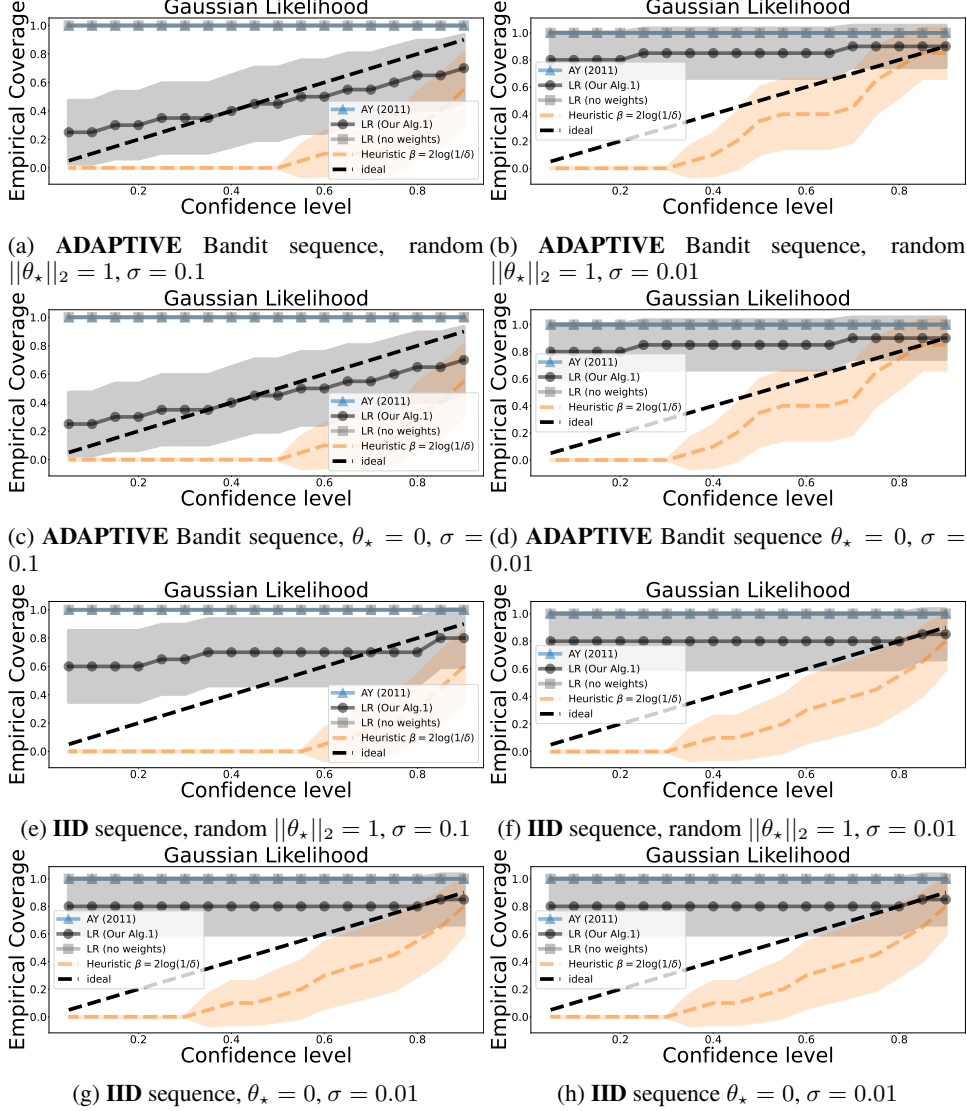


Figure 3: We plot the calibration diagram for data collected from a bandit game trying to optimize a ground truth function using the same model as in Fig. 2a). Instead of the test function, we use an explicit member of the confidence set to avoid a potential mismatch between models. We check after  $T = 15$  whether  $\theta_* \in C_t$  and average over 200 runs. We see that the (LR) are more conservative than the ideal calibration, however, they are provably valid and substantially better than any theoretically motivated confidence sets. We also see that the heuristic is not calibrated and fails many times. We see that for i.i.d. data, our sets are somewhat conservative since the data is not adapted, and our approach is not necessary. We note that the results depend on  $\theta_*$ .

We use this to define ellipsoidal confidence sets as  $\|\hat{\theta}_t - \theta\|_{\mathbf{V}_{\text{laplace}}}^2 \leq 2 \log(1/\delta)$ . The other heuristic suggests using the worst-case parameter instead, namely

$$\mathbf{V}_{\text{mutny}} = \sum_{s=1}^t \exp(B) x_s x_s^\top.$$

This method would have provable coverage with a proper confidence parameter. Its derivation is beyond the scope of this paper.

Benchmark function	dim	$ \mathcal{X} $	$\gamma$	$B$	$\lambda$	Gaussian/Laplace $\sigma/b$
1D	1	$2^6$	0.06	4	1	0.15
Camelback	2	$10^2$	0.2	2	1	0.10

Table 2: Summary of experimental parameters

### E.2.3 NR (2021)

This method follows from [Neiswanger and Ramdas \(2021\)](#). Per se, this method was not developed to be versatile in terms of likelihood but instead provides a confidence set on  $f$ , even if it originates from a misspecified prior. Nevertheless, it provides a likelihood-aware confidence sequence that is anytime-valid and doesn't employ worst-case parameters, and hence is a good benchmark for our analysis. The confidence sets are of the form

$$\{\theta \mid \log \mathcal{L}(\theta) \leq \log(1/\delta) + \log(p(\mathcal{D}))\},$$

where  $\log(p(\mathcal{D}))$  is the current log-evidence of the data given the Gaussian prior. For more information, see [Neiswanger and Ramdas \(2021\)](#).

### E.3 Additive Models

We implemented two likelihoods, namely Gaussian and Laplace. We implemented the discretization of the domain  $|\mathcal{X}|$ , and in the implementation we used Nystrom features defined on  $|\mathcal{X}|$  providing the exact representation of the RKHS on the  $\mathcal{X}$ . The points were chosen to be on a uniform grid. Notice that for the regularized estimator, we chose the rule of thumb  $\lambda = 1/B$  as is motivated in ([Mutný and Krause, 2022](#)).

The laplace parameter was picked as  $b = 0.15$  likewise. Note that Laplace distribution is sub-exponential with parameters  $(b, 2b^2)$ . We use  $1/\sigma^2$  or  $1/b$  respectively for the value  $L$ . Strictly speaking, the Laplace likelihood is not smooth, but a smoothed objective would most likely inherit a value depending on  $b$ . As we maintain coverage with any choice of weighting, we do not risk invalidity by using a heuristic choice for  $L$ .

### E.4 Survival Analysis

We implemented the method exactly as specified, using the Weibull likelihood with parameter  $p = 2$ . Upon log-transformation, the Gumbel distribution is sub-exponential. To determine the parameter, consider the moment-generating function of the Gumbel distribution ( $\beta = 1/p$  in the canonical parameterization):

$$\mathbb{E}[e^{Xt}] = \Gamma(1 - t/2) \exp(t) \leq \exp(t^2/2) \quad \text{for } t < 1/2,$$

hence, the sub-exponentiality parameter is 1, and we can use the above sub-exponential confidence sets with value  $b = 1$ . For the likelihood ratio code, we used  $L = \exp(B)$ , as this is the leading term of the Hessian of log-likelihood. The function is not smooth everywhere, but on a bounded domain, this turns out to be an appropriate scaling.

### E.5 Poisson Bandits

In this case, we implemented a bandit game, where we used the parametrization  $r_\theta(x) \sim \text{Poisson}(\exp(-\theta^\top \Phi(x)))$ , where  $\Phi(x)$  is the RKHS evaluation operator, and  $\theta$  is the unknown value.

We used  $L = \exp(B)$ , as this is the leading term of the Hessian of log-likelihood in this parametrization. The function is not smooth everywhere but on a bounded domain this is an appropriate scaling.

### E.6 Additional Benchmark Functions

We focus on an additional baseline function: Camelback in 2D, a standard BO benchmark function. The results can be see in [Figure 4](#).



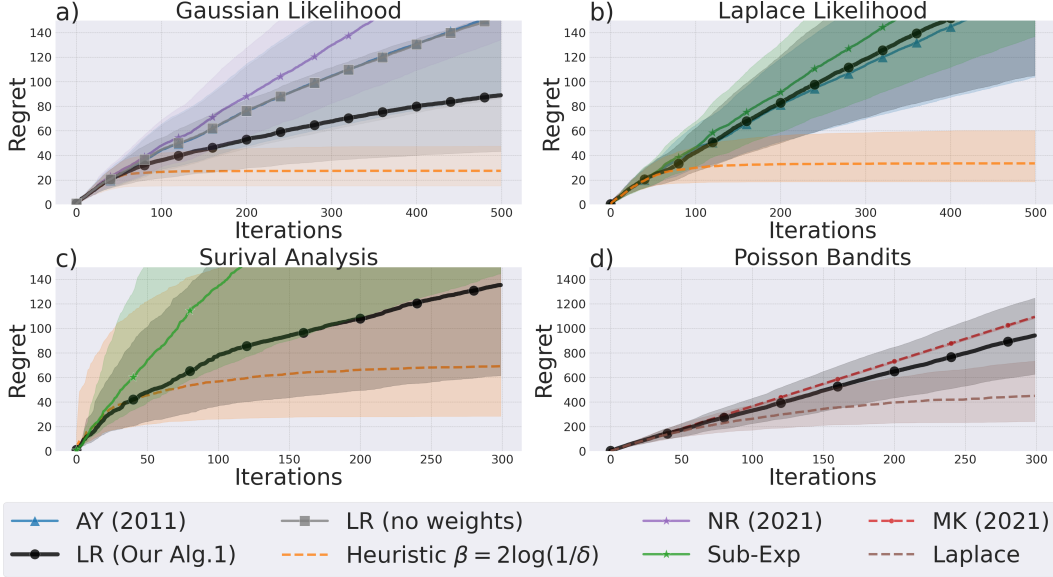


Figure 4: Camelback function. 10 repeats with median and standard quantiles plotted. Note that our method is the best method with provable coverage.

### F Proof of Proposition 3

Define  $S_t$  and the shorthand  $\mathbf{V}_t$

$$S_t = \sum_{s=1}^t \eta_s \frac{x_s}{\nu^2} \quad \text{and} \quad \mathbf{V}_t := \mathbf{V}_t^{\nu^{-2}; \lambda} = \sum_{s=1}^t \frac{x_s x_s^\top}{\nu^2} + \lambda \mathbf{I}$$

and the parametrized process

$$\mathcal{M}_t(x) = \exp(\langle x, S_t \rangle - \frac{1}{2} \|x\|_{\mathbf{V}_t}^2).$$

**Lemma 13.** *If  $\eta_t$  is  $(\nu, \gamma)$ -conditionally sub-Exponential, then  $\mathcal{M}_t(x)$  is a super-martingale on the ball  $\{x \in \mathbb{R}^d \mid \|x\|_2 \leq \frac{\nu^2}{\gamma}\}$  with  $\mathcal{M}_0(x) \leq 1$ .*

Note that for  $\gamma \rightarrow 0$  (sub-Gaussian case), this recovers Lemma 20.2. in [Lattimore and Szepesvári \(2020\)](#).

*Proof.* It is easy to observe that for any  $x$ , we have

$$\exp\left(S_0^\top x - \frac{1}{2} \|x\|_{\mathbf{V}_0}^2\right) = \exp\left(-\frac{1}{2} \|x\|_{\lambda \mathbf{I}}^2\right) \leq 1.$$

For the first part, we can write

$$\begin{aligned} \mathbb{E}[\mathcal{M}_t(x) | \mathcal{F}_{t-1}] &= \mathbb{E}[\exp(\langle x, S_t \rangle - \frac{1}{2} \|x\|_{\mathbf{V}_t}^2) | \mathcal{F}_{t-1}] \\ &= \mathbb{E}[\exp(\langle x, S_{t-1} \rangle - \frac{1}{2} \|x\|_{\mathbf{V}_{t-1}}^2) \exp(\frac{1}{\nu^2} \langle x, \eta_t x_t \rangle - \frac{1}{2\nu^2} \|x\|_{x_t x_t^\top}^2) | \mathcal{F}_{t-1}] \\ &= \mathcal{M}_{t-1}(x) \mathbb{E}[\exp(\frac{1}{\nu^2} \eta_t \langle x, x_t \rangle) | \mathcal{F}_{t-1}] \exp(-\frac{1}{2\nu^2} \|x\|_{x_t x_t^\top}^2), \end{aligned}$$

where in the last step we use that  $x_t$  is  $\mathcal{F}_{t-1}$ -measurable. Now, as long as  $\frac{1}{\nu^2} |\langle x, x_s \rangle| \leq \frac{1}{\gamma}$ , we can apply our definition of conditional sub-Exponential noise to bound

$$\mathbb{E}[\exp(\eta_t \frac{1}{\nu^2} \langle x, x_t \rangle)] \leq \exp\left(\frac{(\langle x, x_t \rangle)^2 \nu^2}{2\nu^4}\right) \leq \exp\left(\frac{\|x\|_{x_t x_t^\top}^2}{2\nu^2}\right).$$

From this we directly conclude

$$\mathbb{E}[\mathcal{M}_t(x)|\mathcal{F}_{t-1}] \leq \mathcal{M}_{t-1}(x).$$

By Cauchy-Schwarz, a sufficient condition is  $\|x\|_2 \leq \frac{\nu^2}{\gamma}$  as this implies (with our assumptions on the actions)

$$|\langle x, x_t \rangle| \leq \|x\|_2 \|x_t\|_2 \leq \|x\|_2 \leq \frac{\nu^2}{\gamma}.$$

□

In the following, will use this result to prove any-time confidence estimates for the parameter  $\theta$  using the technique of pseudo-maximization, closely following [Muný and Krause \(2021\)](#).

Recall that  $\mathcal{M}_t(x)$  is defined on the ball of radius  $\frac{\nu^2}{\gamma L}$ . This allows us some freedom in choosing the radius of the ball on which we integrate. In particular, let  $K$  be this radius. While we need  $K \leq \frac{\nu^2}{\gamma L}$ , we can make  $K$  larger by choosing larger  $\nu^2$  (increasing  $\nu^2$  only makes the set of noise distributions larger).

Ultimately, we wish to bound (following [Lattimore and Szepesvári \(2020\)](#))

$$\begin{aligned} \|\hat{\theta}_t - \theta^*\|_{\mathbf{V}_t} &= \|\mathbf{V}_t^{-1} X_{1:t} (X_{1:t} \theta_* + \eta_{1:t}) - \theta_*\|_{\mathbf{V}_t} \\ &= \|\mathbf{V}_t^{-1} X_{1:t} X_{1:t} \theta_* - \theta_* + \mathbf{V}_t^{-1} S_t\|_{\mathbf{V}_t} \\ &\leq \|S_t\|_{\mathbf{V}_t^{-1}} + \sqrt{\lambda} \|\theta_*\|_2. \end{aligned}$$

We can not control the second term, so we focus on the first: the self-normalized residuals. Via fenchel duality, one can motivate that the right object to study is the supremum of the martingale  $\mathcal{M}_t(x)$  over all  $x \in \mathbb{R}^5$ . Define  $\bar{M}_t$  to be the martingale  $\mathcal{M}_t$  from above but with  $\lambda = 0$ , i.e. no regularisation term. Similarly, let  $\tilde{\mathbf{V}}_t = \mathbf{V}_t - \lambda \mathbf{I}$  be the design matrix without the regularisation term. Slightly counterintuitively, we will study

$$\bar{M}_t = \int_{\|x\|_2 \leq K} \tilde{M}_t(x) dh(x),$$

where  $h$  is the probability density function of a truncated normal distribution with inverse variance  $\lambda$ , that is with covariance matrix  $\frac{1}{\lambda} \mathbf{I}$ . By Lemma 20.3 in [Lattimore and Szepesvári \(2020\)](#),  $\bar{M}_t$  is also a super-martingale with  $\bar{M}_0 \leq 1$ . Then we have

$$\begin{aligned} \bar{M}_t &= \frac{1}{N(h)} \int_{\|x\|_2 \leq K} \exp\left(x^\top S_t - \frac{1}{2} \|x\|_{\tilde{\mathbf{V}}_t}^2\right) \exp\left(-\frac{1}{2} x^\top \lambda \mathbf{I} x\right) dx \\ &= \frac{1}{N(h)} \int_{\|x\|_2 \leq K} \exp\left(x^\top S_t - \frac{1}{2} \|x\|_{\tilde{\mathbf{V}}_t}^2\right) dx. \end{aligned}$$

We will define the shorthand  $f_t(x) = x^\top S_t - \frac{1}{2} x^\top \mathbf{V}_t x = \bar{f}_t(x^*) + \nabla \bar{f}_t(x^*)^\top (x - x^*) - \frac{1}{2} (x - x^*)^\top \mathbf{V}_t (x - x^*)$  (by Taylor's theorem), where  $x^* = \arg \max_{\|x\| \leq kK} f_t(x)$ ,  $k \in (0, 1)$  will be chosen later. We can lower bound  $\bar{M}_t$  by

$$\begin{aligned} \bar{M}_t &= \frac{1}{N(h)} \int_{\|x\|_2 \leq K} \exp\left(x^\top S_t - \frac{1}{2} \|x\|_{\tilde{\mathbf{V}}_t}^2\right) dx \\ &= \frac{\exp(f_t(x^*))}{N(h)} \int_{\|x\|_2 \leq K} \exp\left(\nabla f_t(x^*)^\top (x - x^*) - \frac{1}{2} (x - x^*)^\top \mathbf{V}_t (x - x^*)\right) dx \\ &= \frac{\exp(f_t(x^*))}{N(h)} \int_{\|y+x^*\|_2 \leq K} \exp\left(\nabla f_t(x^*)^\top y - \frac{1}{2} y^\top \mathbf{V}_t y\right) dy \end{aligned} \tag{34}$$

$$\geq \frac{\exp(f_t(x^*))}{N(h)} \int_{\|y\|_2 \leq (1-k)K} \exp\left(\nabla f_t(x^*)^\top y - \frac{1}{2} y^\top \mathbf{V}_t y\right) dy \tag{35}$$

<sup>5</sup>But that is in our case ill-defined

$$\begin{aligned}
&= \frac{\exp(f_t(x^*))}{N(h)} \int_{\|y\|_2 \leq (1-k)K} \exp(\nabla f_t(x^*)^\top y) \exp\left(-\frac{1}{2}y^\top \mathbf{V}_t y\right) dy \\
&= \frac{\exp(f_t(x^*))N(g)}{N(h)} \mathbb{E}_{y \sim g} [\exp(\nabla f_t(x^*)^\top y)] \\
&\geq \frac{\exp(f_t(x^*))N(g)}{N(h)} \exp\left(\mathbb{E}_{y \sim g} [\nabla f_t(x^*)^\top y]\right) \\
&= \frac{\exp(f_t(x^*))N(g)}{N(h)}.
\end{aligned} \tag{36}$$

where in step (34) we used the change of variables  $x = y + x^*$ . In (35) we use that if  $\|y\|_2 \leq (1-k)K$ , then  $\|x^* + y\|_2 \leq \|x^*\|_2 + \|y\|_2 \leq (1-k)K + kK = K$ . Finally, in (36), we used Jensen's inequality. The last inequality follows from symmetry. Note that we implicitly defined  $g$  to be a truncated normal distribution with covariance matrix  $\mathbf{V}_t^{-1}$  on the ball of radius  $(1-k)K$ .

This puts us in a position to put Ville's inequality to good use:

$$\begin{aligned}
\delta &\geq \mathbb{P}(\exists t : \log(\bar{M}_t) \geq \log(1/\delta)) \\
&\geq \mathbb{P}\left(\exists t : f_t(x^*) + \log\left(\frac{N(g)}{N(h)}\right) \geq \log(1/\delta)\right) \\
&\geq \mathbb{P}\left(\exists t : f_t(x^*) \geq \log\left(\frac{N(h)}{N(g)\delta}\right)\right).
\end{aligned}$$

We now wish to recover  $\|S_t\|_{\mathbf{V}_t}$ . Recall the definition of  $f_t(x^*)$  as the maximum over all  $x$  in a ball of radius  $kK$ . Consequently, we can choose  $x = \frac{\mathbf{V}_t^{-1} S_t}{\|S_t\|_{\mathbf{V}_t^{-1}}} \sqrt{\lambda} k K$ , which has norm bounded by  $kK$ .

We have

$$f_t(x^*) \geq f_t\left(\frac{\mathbf{V}_t^{-1} S_t}{\|S_t\|_{\mathbf{V}_t^{-1}}} \sqrt{\lambda} k K\right) = \|S_t\|_{\mathbf{V}_t^{-1}} \sqrt{\lambda} k K - \lambda k^2 K^2,$$

which immediately yields

$$\mathbb{P}\left(\|S_t\|_{\mathbf{V}_t^{-1}} \geq \sqrt{\lambda} k K + \frac{1}{\sqrt{\lambda} k K} \log\left(\frac{N(h)}{N(g)\delta}\right)\right) \leq \delta.$$

The only thing that remains is bounding  $\log\left(\frac{N(h)}{N(g)}\right)$ .

We give the following Lemma that is a slightly generalized version of [Mutný and Krause \(2021\)](#) and originally inspired by [Faury et al. \(2020\)](#).

**Lemma 14.** *The normalizing constants satisfy*

$$\log\left(\frac{N(h)}{N(g)}\right) \leq d \log\left(\frac{1}{1-k}\right) + \log\left(\frac{(\det(\mathbf{V}_t))^{1/2}}{\det(\sqrt{\lambda} \mathbf{I})}\right).$$

We can use the bound from Lemma 14 to conclude that

$$\mathbb{P}\left(\|S_t\|_{\mathbf{V}_t^{-1}} \geq \sqrt{\lambda} k K + \frac{d}{\sqrt{\lambda} k K} \log\left(\frac{1}{1-k}\right) + \frac{1}{\sqrt{\lambda} k K} \log\left(\frac{(\det(\mathbf{V}_t))^{1/2}}{\delta \det(\sqrt{\lambda} \mathbf{I})}\right)\right) \leq \delta.$$

We stated earlier that

$$\|\hat{\theta}_t - \theta^*\|_{\mathbf{V}_t} \leq \|S_t\|_{\mathbf{V}_t^{-1}} + \sqrt{\lambda} \|\theta_\star\|_2.$$

Combining this with our analysis, we get the Proposition 3.

We may now choose the parameters  $k$ ,  $K$  and  $\lambda$ . Note that to get sub-Gaussian rates as in [Abbasi-Yadkori](#), one needs to pick a regularization parameter of the order of  $\lambda = d \log(T)$ .

**Proof of Lemma 14** We give a proof of the Lemma for completeness, and because the additional generality makes for a slightly different proof, even though the bound stays the same.

*Proof.* We have

$$\begin{aligned}
N(h) &= \int_{\|x\|_2 \leq K} \exp(-\lambda \|x\|_2^2) dx \\
&= \frac{1}{|\det(\sqrt{2\lambda}\mathbf{I})|} \int_{\|x\|_2 \leq K} \exp\left(-\frac{1}{2} \|\sqrt{2\lambda}x\|_2^2\right) |\det(\sqrt{2\lambda}\mathbf{I})| dx \\
&= \frac{1}{|\det(\sqrt{2\lambda}\mathbf{I})|} \int_{\|x\|_2 \leq \sqrt{2\lambda}K} \exp\left(-\frac{1}{2} \|x\|_2^2\right) dx.
\end{aligned}$$

Further we have

$$\begin{aligned}
N(g) &= \int_{\|x\|_2 \leq (1-k)K} \exp\left(-\frac{1}{2} x^\top \mathbf{V}_t x\right) dx \\
&= \frac{1}{|\det(\mathbf{V}_t^{1/2})|} \int_{\|x\|_2 \leq (1-k)K} \exp\left(-\frac{1}{2} \|\mathbf{V}_t^{1/2} x\|_2^2\right) |\det(\mathbf{V}_t^{1/2})| dx \\
&= \frac{1}{|\det(\mathbf{V}_t^{1/2})|} \int_S \exp\left(-\frac{1}{2} \|x\|_2^2\right) dx,
\end{aligned}$$

where  $S = \{\mathbf{V}_t^{1/2} x \mid \|x\|_2 \leq (1-k)K\} = \{x \mid \|\mathbf{V}_t^{-1/2} x\|_2 \leq (1-k)K\} = \{x \mid x^\top \mathbf{V}_t^{-1} x \leq (1-k)K\}$ . Note that  $\mathbf{V}_t \succeq \lambda \mathbf{I}$  and so  $\mathbf{V}_t^{-1} \preceq \frac{1}{\lambda} \mathbf{I}$ . Therefore if  $\|x\|_2^2 \leq (1-k)K\sqrt{\lambda}$ , we have

$$\sqrt{x^\top \mathbf{V}_t^{-1} x} \leq \frac{1}{\sqrt{\lambda}} \|x\|_2 \leq (1-k)K \implies x \in S.$$

Thus  $\{x \mid \|x\|_2 \leq (1-k)\sqrt{\lambda}K\} \subseteq S$  and

$$N(g) \geq \frac{1}{|\det(\mathbf{V}_t^{1/2})|} \int_{\|x\|_2 \leq (1-k)\sqrt{\lambda}K} \exp\left(-\frac{1}{2} \|x\|_2^2\right) dx.$$

We may therefore bound

$$\frac{N(g)}{N(h)} \leq \frac{(\det \mathbf{V}_t)^{1/2} \int_{\|x\|_2 \leq \sqrt{2\lambda}K} \exp\left(-\frac{1}{2} \|x\|_2^2\right) dx}{(\det \sqrt{2\lambda}\mathbf{I}) \int_{\|x\|_2 \leq (1-k)\sqrt{\lambda}K} \exp\left(-\frac{1}{2} \|x\|_2^2\right) dx}.$$

By a rather crude bound (as  $1-k \leq \sqrt{2}$  in any case) we get

$$\begin{aligned}
&\frac{\int_{\|x\|_2 \leq \sqrt{2\lambda}K} \exp\left(-\frac{1}{2} \|x\|_2^2\right) dx}{\int_{\|x\|_2 \leq (1-k)\sqrt{\lambda}K} \exp\left(-\frac{1}{2} \|x\|_2^2\right) dx} \\
&\leq \frac{\int_{\|x\|_2 \leq (1-k)\sqrt{\lambda}K} \exp\left(-\frac{1}{2} \|x\|_2^2\right) dx + \int_{(1-k)\sqrt{\lambda}K \leq \|x\|_2 \leq \sqrt{2\lambda}K} \exp\left(-\frac{1}{2} \|x\|_2^2\right) dx}{\int_{\|x\|_2 \leq (1-k)\sqrt{\lambda}K} \exp\left(-\frac{1}{2} \|x\|_2^2\right) dx} \\
&= 1 + \frac{\int_{(1-k)\sqrt{\lambda}K \leq \|x\|_2 \leq \sqrt{2\lambda}K} \exp\left(-\frac{1}{2} \|x\|_2^2\right) dx}{\int_{\|x\|_2 \leq (1-k)\sqrt{\lambda}K} \exp\left(-\frac{1}{2} \|x\|_2^2\right) dx} \\
&\leq 1 + \frac{\exp\left(-\frac{1}{2}(1-k)^2\lambda K^2\right) \int_{(1-k)\sqrt{\lambda}K \leq \|x\|_2 \leq \sqrt{2\lambda}K} dx}{\exp\left(-\frac{1}{2}(1-k)^2\lambda K^2\right) \int_{\|x\|_2 \leq (1-k)\sqrt{\lambda}K} dx} \\
&= 1 + \frac{\text{vol}_d(\sqrt{2\lambda}K) - \text{vol}_d((1-k)\sqrt{\lambda}K)}{\text{vol}_d((1-k)\sqrt{\lambda}K)} \\
&= \frac{\text{vol}_d(\sqrt{2\lambda}K)}{\text{vol}_d((1-k)\sqrt{\lambda}K)}
\end{aligned}$$

$$= (1 - k)^{-d} \sqrt{2}^d.$$

We can put this together to obtain

$$\frac{N(h)}{N(g)} \leq (1 - k)^{-d} \sqrt{2}^d \frac{(\det(\mathbf{V}_t))^{1/2}}{\det(\sqrt{2\lambda}\mathbf{I})} = (1 - k)^{-d} \frac{(\det(\mathbf{V}_t))^{1/2}}{\det(\sqrt{\lambda}\mathbf{I})}.$$

□