

Vector-valued self-normalized concentration inequalities beyond sub-Gaussianity

Diego Martinez-Taboada

Department of Statistics & Data Science, Carnegie Mellon University

DIEGOMAR@ANDREW.CMU.EDU

Tomas Gonzalez

Machine Learning Department, Carnegie Mellon University

TCGONZAL@ANDREW.CMU.EDU

Aaditya Ramdas

Department of Statistics & Data Science, Carnegie Mellon University

Machine Learning Department, Carnegie Mellon University

ARAMDAS@ANDREW.CMU.EDU

Editors: Matus Telgarsky and Jonathan Ullman

Abstract

The study of self-normalized processes plays a crucial role in a wide range of applications, from sequential decision-making to econometrics. While the behavior of self-normalized concentration has been widely investigated for scalar-valued processes, vector-valued processes remain comparatively underexplored, especially outside of the sub-Gaussian framework. In this contribution, we provide concentration bounds for self-normalized processes with light tails beyond sub-Gaussianity (such as Bennett or Bernstein bounds). We illustrate the relevance of our results in the context of online linear regression, with applications in (kernelized) linear bandits.

Keywords: self-normalized processes, concentration inequalities, online linear regression.

1. Introduction

Self-normalized processes naturally arise in a variety of applications, ranging from econometrics (Shao, 2015; Agarwal et al., 2018) and finance (Darolles et al., 2006; Pacurar, 2008) to sequential decision making (Abbasi-Yadkori et al., 2011; Chowdhury and Gopalan, 2017; Yang et al., 2020). Concentration inequalities for self-normalized processes have been widely studied in the case of scalar-valued random variables (de la Peña et al., 2004, 2007, 2009b). Nonetheless, important theoretical challenges arise when working in higher dimensions and concentration inequalities for vector-valued self-normalized processes are scarce in the literature, the majority of them assuming sub-Gaussianity of the random variables. If the tails are sub-Gaussian, mixture arguments lead to closed-form, powerful inequalities (Abbasi-Yadkori et al., 2011; Chowdhury and Gopalan, 2017). The same methods do not seem easy to generalize to other regimes.

In this work, we aim to establish general concentration inequalities that hold for light tails. More precisely, let us define

$$V_t = \sum_{i \leq t} X_i X_i^T, \quad M_t = \sum_{i \leq t} X_i \epsilon_i, \quad (1.1)$$

where X_t are sequentially randomly drawn or even adversarially chosen, and ϵ_t is a real-valued noise with zero expectation and light tails (these definitions will be formalized in Section 3). The ultimate goal of this contribution is to provide new probabilistic guarantees of the form

$$\|(\rho I + V_t)^{-\frac{1}{2}} M_t\| \leq J_t(\delta) \text{ simultaneously for all } t \text{ with probability } 1 - \delta \quad (1.2)$$

for different assumptions on the noises ϵ_t . Importantly, note that (1.2) establishes a probabilistic guarantee that holds uniformly over time; this sort of guarantee is usually exploited in the fully sequential scenario, where one may peek at the data at any given point of the procedure (Ramdas et al., 2020, 2023). Minimizing $J_t(\delta)$ for a fixed $n = t$ corresponds to the (more classical) batch setting, where the sample size is fixed prior to data collection. In contrast, one may be interested in utilizing (1.2) without specifying a target sample size in advance, which is useful when developing procedures that may be continuously monitored and adaptively stopped (Howard et al., 2021).

Furthermore, we aim for concentration inequalities that hold in arbitrary Hilbert spaces and so do not have an explicit dependence on the dimension of the Hilbert space. In this regime, much less is known beyond sub-Gaussian ϵ_t . For example, can alternative concentration inequalities be provided when ϵ_t attains the Bernstein’s condition, which allows for heavier tails than sub-Gaussian? When ϵ_t is bounded, can we derive concentration inequalities that adapt to the unknown variance of the random variables, and not only to conservative upper bounds that dictate their sub-Gaussian behavior? This work provides a positive answer to these questions.

Related work. Few contributions have proposed dimension-free self-normalized concentration inequalities for light-tailed noises beyond sub-Gaussianity; we defer an extended presentation of related work to Appendix B. The closest work is that of Zhou et al. (2021, Theorem 4.1), which presented a Bernstein-type self-normalized concentration inequality for vector-valued martingales, whose proof exploits solely univariate concentration inequalities (similarly to Dani et al. (2008, Theorem 5)). In contrast to our Bernstein-type inequality, this results in substantially looser constants and an extra logarithmic term in the sample size. Furthermore, our results are significantly more general, leading to Bennett-type inequalities and generalizing to unbounded random variables.

The recent work of Whitehouse et al. (2025) does handle (for example) sub-exponential and sub-Poisson noise, but it uses covering arguments to provide dimension-dependent bounds that depend on the condition number of the variance process. In the sub-Gaussian case, their bounds were shown to be incomparable to the better known log-determinant bounds, but the bounds are inherently restricted to finite-dimensional settings, unlike ours.

Concurrent contributions have also explored dimension-independent self-normalized inequalities. Metelli et al. (2025) developed Bernstein-like concentration inequalities for bounded noises relying on stitching arguments, and Akhavan et al. (2025) leveraged method of mixtures and truncation arguments to also develop a Bernstein-like concentration inequality. Again, these two contributions are restricted to the bounded setting, while our bounds are not.

Finally, the recent preprint of Chugg and Ramdas (2025) developed inequalities via PAC-Bayes; while the bounds are very general and explicit, it is not straightforward to analyze the rate because they depend on the inverse of a tail decay function and some ratio of eigenvalues.

Main contributions and techniques. Our approach is fundamentally different to the aforementioned works, building on tools originally introduced in Pinelis (1994). More concretely, we leverage techniques therein to derive a novel nonnegative supermartingale, which leads to clean concentration inequalities when combined with Ville’s inequality. Importantly, we characterize the concentration inequalities for the self-normalized processes by clearly decoupling the effect of the directions (X_t) and the tail behavior of the noise (ϵ_t) . In particular, our concentration inequalities are applicable as soon as $\mathbb{E}[\exp(\lambda|\epsilon_i|) - \lambda|\epsilon_i| - 1]$ can be controlled, thus naturally adapting to a larger family of light-tailed noises (e.g., sub-exponential or sub-Poisson), not necessarily sub-Gaussian or even bounded. Furthermore, our concentration inequalities are dimension-free and

applicable in any separable Hilbert space. The inequalities are clean and without large constants, and so they are readily applicable to conduct inference, as we elucidate in Section 5.

2. Preliminaries

2.1. Separable Hilbert spaces

Throughout, we will work with random elements in a separable Hilbert space H . We remind the reader that a Hilbert space is a complete inner product space. A *separable* Hilbert space is a Hilbert space that contains a countable, dense subset. Any separable Hilbert space is linearly isometric to $l^2(\mathbb{N}) = \{(x_i)_{i \in \mathbb{N}} : \sum_{i \in \mathbb{N}} x_i^2 < \infty\}$, so we can think of separable Hilbert spaces as potentially infinite sequences whose sum of squares is finite (thus generalizing the usual, finite dimensional Euclidean spaces). Separability is usually assumed given its generality and measure-theoretic convenience (among other reasons, see e.g. [Ledoux and Talagrand \(2013, Chapter 2\)](#)). We highlight that every Hilbert space is a $(2, 1)$ -smooth Banach space ([Pinelis, 1994](#)).

Notation. Given two elements $f, g \in H$, we denote their inner product by either $\langle f, g \rangle$ or $f^T g$, and their outer product by $f g^T := f \langle g, \cdot \rangle$, which is a linear operator from H to itself. Similarly, the identity operator from H to itself is denoted by I .

2.2. Nonnegative supermartingales and Ville's inequality

Nonnegative supermartingales play a central role in deriving anytime valid concentration inequalities due to Ville's inequality ([Ville, 1939](#)). Before presenting the inequality, we introduce some notation. Let $\mathbb{N} = \{0, 1, 2, \dots\}$. A filtration $\mathcal{F} = (\mathcal{F}_t)_{t \in \mathbb{N}}$ is a sequence of σ -algebras such that $\mathcal{F}_t \subseteq \mathcal{F}_{t+1}$, for all t . A stochastic process $M = (M_t)_{t \in \mathbb{N}}$ is a sequence of random variables that are adapted to $(\mathcal{F}_t)_{t \in \mathbb{N}}$, meaning that M_t is \mathcal{F}_t -measurable for all t . M is called *predictable* if M_t is \mathcal{F}_{t-1} -measurable for all t . An integrable stochastic process M is a supermartingale with respect to \mathcal{F} if $\mathbb{E}[M_{t+1} | \mathcal{F}_t] \leq M_t$ for all t , and a martingale if the inequality always holds with equality; inequalities or equalities between random variables are always interpreted to hold almost surely. Throughout, we use the shorthand $\mathbb{E}_t[\cdot]$ for $\mathbb{E}[\cdot | \mathcal{F}_t]$.

Fact 1 (Ville's inequality) *If M is a nonnegative supermartingale (with respect to any filtration \mathcal{F}), then for any $x > 0$,*

$$\mathbb{P}(\exists t \in \mathbb{N} : M_t \geq x) \leq \frac{\mathbb{E}[M_0]}{x}.$$

This result can be seen as a time-uniform version of Markov's inequality, and yields anytime-valid concentration inequalities in a similar way in which Markov's does for fixed t . Indeed, note that if $M_0 = 1$, then selecting $x = 1/\delta$ in Ville's inequality implies that

$$M_t \leq 1/\delta \text{ for all } t \in \mathbb{N} \text{ with probability at least } 1 - \delta,$$

mirroring the form of the probabilistic guarantee stated in (1.2). To obtain the exact expression in (1.2), we will need to carefully rearrange the inequality. The main technical challenge then is to appropriately construct (super)martingales and rearrange the obtained inequalities.

2.3. Vector-valued concentration beyond subGaussianity

The nonnegative supermartingale derived in this contribution makes essential use of the vector-valued results from [Pinelis \(1994\)](#). In contrast to most prominent approaches regarding vector-valued concentration—which exploit either covering arguments, or PAC-Bayes, or mixture martingales—[Pinelis \(1994\)](#) directly worked with the properties of the norm. In particular, we use in this contribution the following result, which is a simplified version of [Pinelis \(1994, Theorem 3.2\)](#).

Fact 2 *Let $f \in H$ be fixed, and $X \in H$ be a random element. It holds that*

$$\mathbb{E} \cosh \|f + \lambda X\| \leq (1 + \tilde{e}_t(\lambda)) \cosh \|f\|,$$

where $\tilde{e}_t(\lambda) := \mathbb{E} [\exp(\lambda \|X\|) - \lambda \|X\| - 1]$.

Note that [Theorem 2](#) allows for working with the moment generating function of $\|X\|$ but without the first order term, which bypasses the main obstacle when using other well-known techniques that relate to Chernoff-bounds. Nonetheless, there is no sense of self-normalization in the original formulation of this inequality, which is the main object under study in this contribution.

3. Problem statement

Let $(X_t)_{t \geq 0}$ be a H -valued stochastic process, where H is a separable Hilbert space, and let $(\epsilon_t)_{t \geq 0}$ be a real-valued stochastic process. We are interested in deriving time-uniform concentration inequalities for $\|(\rho I + V_t)^{-\frac{1}{2}} M_t\|$, where

$$V_t = \sum_{i \leq t} X_i X_i^T, \quad M_t = \sum_{i \leq t} X_i \epsilon_i$$

and $\rho > 0$ is fixed. Throughout, we work under the following assumption.

Assumption 3 *There exists a filtration $\mathcal{F} = (\mathcal{F}_t)_{t \geq 0}$ such that (i) X_t is \mathcal{F}_{t-1} -measurable for all t , (ii) ϵ_t is \mathcal{F}_t -measurable for all t , and (iii) $\mathbb{E}_{t-1} \epsilon_t = 0$, where we recall that $\mathbb{E}_{t-1}[\cdot] = \mathbb{E}[\cdot | \mathcal{F}_{t-1}]$.*

In many applications, \mathcal{F} is naturally taken as the canonical filtration generated by both stochastic processes. That is, $\mathcal{F}_t = \sigma((X_1, \epsilon_1), \dots, (X_t, \epsilon_t))$, with \mathcal{F}_0 being trivial. In this case, (i) implies that X_t is drawn from a distribution that can depend on $(X_1, \epsilon_1, X_2, \epsilon_2, \dots, X_{t-1}, \epsilon_{t-1})$, but cannot depend on the outcomes of (yet to be seen) later rounds. This setting encompasses both independent and identically distributed (iid) draws for X_t , as well as potentially ‘adversarially’ chosen (X_t) , where the adversary has the information only until round $t - 1$. The noises ϵ_t can further depend on X_t , thereby accommodating heteroskedasticity, which corresponds to (ii). However, (iii) imposes a martingale structure on (ϵ_t) : intuitively, we can think of ϵ_t as noise and thus it has zero conditional expectation. As we shall see in later sections, working under [Assumption 3](#) allows for providing concentration inequalities for the self-normalized processes. In particular, the tails of the noises ϵ_t play a key role in such concentration bounds. We focus on two types of light-tailed noises, given their importance and ubiquity in the literature.

Assumption 4 (Bernstein condition) *There exist B and (σ_t) such that (ϵ_t) fulfils*

$$\mathbb{E}_{t-1}|\epsilon_t|^m \leq \frac{1}{2}m!B^{m-2}\sigma_t^2 \quad \forall m \geq 2, \forall t \geq 1,$$

where σ_t is \mathcal{F}_{t-1} -measurable.

Assumption 4 is commonly summarized as: ϵ_t satisfies the Bernstein condition with parameters (σ_t, B) . A sufficient condition for Bernstein's condition to hold is that the random variable is bounded, but it can also be satisfied by several unbounded variables (in fact, it is equivalent to a sub-exponential tail bound condition, see e.g. [Howard et al. \(2020\)](#)), giving it much wider applicability.

Assumption 5 (Bounded noise and variance) *There exist B and (σ_t) such that the stochastic process (ϵ_t) fulfils, for all $t \geq 0$,*

$$|\epsilon_t| \leq B, \quad \mathbb{V}_{t-1}[\epsilon_t] \leq \sigma_t^2,$$

where σ_t is \mathcal{F}_{t-1} -measurable.

Assumption 5 imposes boundedness on the random noises (and hence their variance). Although bounded random variables are also sub-Gaussian, exploiting (the upper bound of) the variance can lead to sharper results.

While we focus our efforts in random noises that fulfill one of the former two assumptions for clarity of exposition, the results presented in this work hold as soon as the conditional expectations $\mathbb{E}_{t-1}[\exp(\lambda|\epsilon_t|) - \lambda|\epsilon_t| - 1]$ can be controlled for (with the former assumptions being specific instances of this case). Thus, the applicability of this work also goes beyond these two examples.

4. Main results

We present in this section the main results of the contribution, namely a supermartingale construction alongside the concentration inequalities that can be derived from it. Contrary to current approaches, our supermartingale construction cleanly decouples the norm representation of the vector-valued process from the concentration of the one-dimensional noises. In order to do that, we elucidate in Section 4.1 the behaviour of the terms $\|G_t\|$, where

$$G_t := \left(\rho I + \sum_{i \leq t} X_i X_i^T \right)^{-\frac{1}{2}} X_t.$$

Section 4.2 combines such G_t terms with the noises in order to provide a supermartingale construction for our central object of study. Such a supermartingale construction composes the main theoretical contribution of this work, from which the remaining results can be derived. The pseudo-variance process $\sum_{i \leq t} \|G_i\|^2$ plays an important role in the supermartingale construction, and we devote Section 4.3 to elucidate its connection to the intrinsic dimension of the Hilbert space. Concentration inequalities are then obtained in Section 4.4, where inequalities for both the fixed time and fully sequential settings are rigorously derived. In particular, Bernstein, Bennett, and empirical Bennett-type concentration inequalities are presented.

4.1. Decoupling the direction from the noise

We highlight the decomposition of the self-normalized process

$$(\rho I + V_t)^{-\frac{1}{2}} M_{t-1} + (\rho I + V_t)^{-\frac{1}{2}} X_t \epsilon_t. \quad (4.1)$$

As it shall be seen in Section 4.2, controlling the tails of the second term in (4.1) suffices to provide a supermartingale construction, which leads to concentration bounds $J_t(\delta)$. We observe that

$$\left\| (\rho I + V_t)^{-\frac{1}{2}} X_t \epsilon_t \right\| = \left\| (\rho I + V_t)^{-\frac{1}{2}} X_t \right\| |\epsilon_t| = \|G_t\| |\epsilon_t|,$$

which decouples the effect of the noise ϵ_t from the vector G_t on which it is projected. In particular, $\|G_t\|$ depends on ρ , thereby concealing the effect of ρ on the bound. It is also key to note that $\|G_t\| \leq 1$. Indeed, we observe that

$$(\rho I + \sum_{i \leq t} X_i X_i^T) \succeq (\rho I + X_t X_t^T),$$

and so, by using the Sherman-Morrison rank-one update formula,

$$\|G_t\|^2 \leq \left\| (\rho I + X_t X_t^T)^{-1/2} X_t \right\|^2 = X_t^T (\rho I + X_t X_t^T)^{-1} X_t = \frac{\|X_t\|^2}{\rho + \|X_t\|^2} \leq 1.$$

4.2. Light-tailed self-normalized process supermartingale constructions

We now provide a nonnegative supermartingale that adapts to different tail behaviours of the random noises through $\mathbb{E}_{t-1} [\exp(\lambda |\epsilon_t|) - \lambda |\epsilon_t| - 1]$. Hence, this construction may be exploited as long as the moment generating function of the absolute value of the noises can be controlled for. As it is usual when deriving concentration inequalities, we will work with the process that is object of study multiplied by an arbitrary value $\lambda > 0$, that is then optimally adjusted or mixed to yield tight inequalities. The next theorem establishes such a nonnegative supermartingale construction.

Theorem 6 *Let $(X_t)_{t \geq 1}$ and $(\epsilon_t)_{t \geq 1}$ be Hilbert space valued and real valued processes, respectively, attaining Assumption 3. Let $\lambda > 0$ and recall that $G_t = (\rho I + V_t)^{-\frac{1}{2}} X_t$. Denoting*

$$e_t(\lambda) = \|G_t\|^2 \mathbb{E}_{t-1} [\exp(\lambda |\epsilon_t|) - \lambda |\epsilon_t| - 1],$$

the process

$$S_t = \cosh \left(\lambda \left\| (\rho I + V_t)^{-1/2} M_t \right\| \right) \exp \left(- \sum_{i \leq t} e_i(\lambda) \right) \quad (4.2)$$

is a nonnegative supermartingale.

Proof To prove that S_t is a supermartingale (it is trivially nonnegative), we first observe that

$$\begin{aligned} \left\| \lambda (\rho I + V_t)^{-1/2} M_t \right\| &= \left\| \lambda (\rho I + V_t)^{-1/2} M_{t-1} + (\rho I + V_t)^{-1/2} (\lambda X_t \epsilon_t) \right\| \\ &= \left\| \lambda (\rho I + V_t)^{-1/2} M_{t-1} + \lambda G_t \epsilon_t \right\|. \end{aligned}$$

Plugging the above into Fact 2 establishes that for $\tilde{e}_t(\lambda) := \mathbb{E}_{t-1} [\exp(\lambda \|G_t \epsilon_t\|) - \lambda \|G_t \epsilon_t\| - 1]$, we have

$$\begin{aligned} \mathbb{E}_{t-1} \left[\cosh \left\| \lambda(\rho I + V_t)^{-1/2} M_t \right\| \right] &\leq (1 + \tilde{e}_t(\lambda)) \cosh \left\| \lambda(\rho I + V_t)^{-1/2} M_{t-1} \right\| \\ &\leq \exp(\tilde{e}_t(\lambda)) \cosh \left\| \lambda(\rho I + V_t)^{-1/2} M_{t-1} \right\|. \end{aligned}$$

Since $\|G_t\| \leq 1$ and it is \mathcal{F}_{t-1} -measurable, we see that

$$\begin{aligned} \tilde{e}_t(\lambda) &= \mathbb{E}_{t-1} [\exp(\lambda \|G_t \epsilon_t\|) - \lambda \|G_t \epsilon_t\| - 1] \\ &= \mathbb{E}_{t-1} \left[\sum_{k \geq 2} \frac{(\lambda \|G_t\| |\epsilon_t|)^k}{k!} \right] \\ &\leq \|G_t\|^2 \mathbb{E}_{t-1} \left[\sum_{k \geq 2} \frac{(\lambda |\epsilon_t|)^k}{k!} \right] = e_t(\lambda). \end{aligned}$$

Plugging this back into the earlier expression, and noting that $V_{t-1} \preceq V_t$, we have

$$\begin{aligned} \mathbb{E}_{t-1} \left[\cosh \left\| \lambda(\rho I + V_t)^{-1/2} M_t \right\| \right] &\leq \exp(e_t(\lambda)) \cosh \left\| \lambda(\rho I + V_t)^{-1/2} M_{t-1} \right\| \\ &\leq \exp(e_t(\lambda)) \cosh \left\| \lambda(\rho I + V_{t-1})^{-1/2} M_{t-1} \right\|. \end{aligned}$$

This implies the supermartingale property:

$$\begin{aligned} \mathbb{E}_{t-1}[S_t] &= \mathbb{E}_{t-1} \left[\cosh \left(\left\| \lambda(\rho I + V_t)^{-1/2} M_t \right\| \right) \exp \left(- \sum_{i \leq t} e_i(\lambda) \right) \right] \\ &\leq \exp(e_t(\lambda)) \cosh \left\| \lambda(\rho I + V_{t-1})^{-1/2} M_{t-1} \right\| \exp \left(- \sum_{i \leq t} e_i(\lambda) \right) \\ &= S_{t-1}. \end{aligned}$$

■

Intuitively, the supermartingale is constructed by re-expressing the normalized update so that all geometric complexity is absorbed into $\|G_t\|$, while all randomness is carried solely by the scalar noise $|\epsilon_t|$. In this “whitened” geometry, the vector-valued process behaves like a one-dimensional process whose increments have size at most $|\epsilon_t|$. Applying Pinelis’ inequality to the cosh of the normalized radius yields a predictable exponential correction term, and subtracting these terms creates a supermartingale.

Note that if the noises (ϵ_t) attain the Bernstein condition (Assumption 4) with parameters (σ_t, B) , then (see e.g. [Wainwright \(2019, Equation 2.16\)](#))

$$\mathbb{E}_{t-1} [\exp(\lambda |\epsilon_t|) - \lambda |\epsilon_t| - 1] \leq \frac{\lambda^2}{2(1 - \lambda B)} \sigma_t^2, \quad 0 \leq \lambda < \frac{1}{B}.$$

Otherwise, if the noises (ϵ_t) are bounded by B and their variance is (σ_t^2) (Assumption 5), it follows that

$$\mathbb{E}_{t-1} [\exp(\lambda |\epsilon_t|) - \lambda |\epsilon_t| - 1] \leq \frac{e^{\lambda B} - \lambda B - 1}{B^2} \sigma_t^2.$$

The shrinkage factors $\exp(-\sum_{i \leq t} e_i(\lambda))$ that emerge in our nonnegative supermartingale may appear too large at first glance, given its non-standard form. Shrinkage factors are usually found in the form $\exp(-\psi(\lambda) \sum_{i \leq t} \zeta_i^2)$, where ψ is a CGF-like function, and ζ_i^2 are conditional pseudo-variances; see Howard et al. (2020) for a discussion on this general sub- ψ framework. Despite their different appearances, these factors are analogous. For example, the usual Bennett-inequality includes the shrinkage factor $\exp(-\psi_{P,B}(\lambda) \sum_{i \leq t} \zeta_i^2)$, where $\psi_{P,B}(\lambda) := (\exp(\lambda B) - \lambda B - 1)/B^2$ is the so-called sub-Poisson ψ -function. In contrast, our upper bound for $\mathbb{E}_{t-1}[\exp(\lambda |\epsilon_t|) - \lambda |\epsilon_t| - 1]$ for bounded noises is precisely $\psi_{P,B}(\lambda) \sigma_t^2$. Thus, we recover the sub- $\psi_{P,B}$ shrinkage factor with $\zeta_i^2 = \|G_i\|^2 \sigma_i^2$. Similarly, the Bernstein condition upper bound leads to a shrinkage factor of $\exp(-\psi_{G,B}(\lambda) \sum_{i \leq t} \|G_i\|^2 \sigma_i^2)$, where $\psi_{G,B} = \lambda^2/(2 - 2B\lambda)$ is the usual sub-Gamma ψ -function. Thus, we see that our results fall under the umbrella of sub- ψ shrinkage factors.

Although we generally recover sub- ψ shrinkage factors, we do not do so via the usual sub- ψ definition. That is, we do not assume that $\mathbb{E}_{t-1} \exp(\lambda \epsilon_t) \leq \exp(\psi(\lambda) \zeta_t^2)$. Instead, we (implicitly) exploit an upper bound of the form $\mathbb{E}_{t-1}[\exp(\lambda |\epsilon_t|) - \lambda |\epsilon_t| - 1] \leq \psi(\lambda) \zeta_t^2$. Equivalently, we can rewrite this inequality as $\mathbb{E}_{t-1}[\exp(\lambda |\epsilon_t|) - \lambda |\epsilon_t|] \leq 1 + \psi(\lambda) \zeta_t^2$. Now we note that

$$1 + \psi(\lambda) \zeta_t^2 \leq \exp(\psi(\lambda) \zeta_t^2),$$

and if $\mathbb{E}_{t-1} \epsilon_t = 0$, the inequality

$$\mathbb{E}_{t-1} \exp(\lambda \epsilon_t) = 1 + \mathbb{E}_{t-1} \sum_{k \geq 2} \frac{(\lambda \epsilon_t)^k}{k!} \leq 1 + \mathbb{E}_{t-1} \sum_{k \geq 2} \frac{(\lambda |\epsilon_t|)^k}{k!} \leq \mathbb{E}_{t-1} [\exp(\lambda |\epsilon_t|) - \lambda |\epsilon_t|]$$

holds as well. This implies that our condition on ϵ_t is slightly stronger than the usual sub- ψ condition, and so we cannot present our results under the usual sub- ψ framework. Nonetheless, many of the predominant random variables that are sub- ψ also attain our (only slightly) stronger condition. Clear cases are the aforementioned bounded and Bernstein condition noises achieving our condition, on top of being sub-Poisson and sub-Gamma respectively.

Theorem 6 allows for obtaining concentration inequalities in view of Ville's inequality. The result is formalized in the following proposition, and its proof can be found in Appendix A.1.

Proposition 7 (Light-tailed self-normalized process concentration inequality) *Under the assumptions of Theorem 6, it holds that, with probability $1 - \delta$ and simultaneously for all $t \geq 1$,*

$$\left\| (\rho I + V_t)^{-1/2} M_t \right\| \leq \frac{\sum_{i=1}^t e_i(\lambda) + \log\left(\frac{2}{\delta}\right)}{\lambda}.$$

Proposition 7 immediately yields concentration inequalities for our central object of interest. Nonetheless, the bound provided by Proposition 7 depends on an arbitrary value λ that has to be chosen prior to data collection. Thus, the tightness of the concentration bounds entirely relies on the choice for λ . Furthermore, if the term $\mathbb{E}_{t-1} [\exp(\lambda |\epsilon_t|) - \lambda |\epsilon_t| - 1]$ is constant across t , the sum of $e_i(\lambda)$ scales with $\sum_{i \leq t} \|G_i\|^2 \sigma_i^2$, and the optimal choice of λ relies on an upper bound of this process.

While our focus is on non-asymptotic concentration inequalities, Theorem 6 also leads to the so-called laws of the iterated logarithm (Kolmogoroff, 1929; Stout, 1970). The laws of the iterated logarithm characterize the almost-sure limiting envelope of normalized sums of random variables, describing the oscillatory behavior of stochastic processes between the law of large numbers and the central limit theorem. In particular, the following general upper asymptotic law of the iterated logarithm is immediately obtained from Theorem 6 in combination with Howard et al. (2021, Corollary 1); its proof is presented in Appendix A.2. One can also use Howard et al. (2021, Theorem 1) to derive explicit nonasymptotic laws of the iterated logarithm bounds directly from Theorem 6.

Corollary 8 *Under Assumption 3, and either Assumption 4 or Assumption 5,*

$$\limsup_{t \rightarrow \infty} \frac{\|(\rho I + V_t)^{-1/2} M_t\|}{\sqrt{2 \left(\sum_{i \leq t} \|G_i\|^2 \sigma_i^2 \right) \log \log \left(\sum_{i \leq t} \|G_i\|^2 \sigma_i^2 \right)}} \leq 1 \quad \text{on} \quad \left\{ \sup_t \sum_{i \leq t} \|G_i\|^2 \sigma_i^2 = \infty \right\}.$$

4.3. Upper bounding the pseudo-variance process

While we have stated all our results for arbitrary separable Hilbert spaces, prominent examples of Hilbert spaces in statistical applications are finite-dimensional Euclidean spaces and reproducing kernel Hilbert spaces (RKHS). Given that finite-dimensional Euclidean spaces are instances of RKHS's, our discussion will be framed in terms of the latter.

Often, the quantity $\|G_t\|$ is expected to shrink as t grows, because the matrix V_t accumulates information and enlarges the ellipsoid in which the normalization occurs. As more covariates are observed, the directions in which the data has significant energy get scaled down more aggressively by $(\rho I + V_t)^{-1/2}$, making each X_t appear smaller in the normalized geometry. Intuitively, new observations rarely point in directions that have not already been heavily regularized. This progressive flattening of the geometry is what causes $\|G_t\|$ to typically decrease.

More specifically, note that $\|G_i\| \leq 1$ and $x \leq 2 \log(1 + x)$ for $x \in [0, 1]$, so

$$\begin{aligned} \sum_{i \leq t} \|G_i\|^2 &\leq 2 \sum_{i \leq t} \log(1 + \|G_i\|^2) = 2 \sum_{i \leq t} \log \left(1 + \left\| (\rho I + V_t)^{-1/2} X_t \right\|^2 \right) \\ &\leq 2 \sum_{i \leq t} \log \left(1 + \left\| (\rho I + V_{t-1})^{-1/2} X_t \right\|^2 \right) = 2 \log \det(I + \rho^{-1} V_t), \end{aligned}$$

where the last inequality follows from $V_{t-1} \preceq V_t$, and the last equality from the elliptical potential lemma (Abbasi-Yadkori et al., 2011, proof of Lemma 11). For bounded (X_t) , one may further consider the upper bound

$$\frac{1}{4} \sum_{i \leq t} \|G_i\|^2 \leq \sup_{V_t} \frac{1}{2} \log \det(I + \rho^{-1} V_t) =: \gamma_t(\rho), \quad (4.3)$$

with $\gamma_t(\rho)$ being the maximal information gain, a concept that relates to the intrinsic dimension of the RKHS and that has been widely exploited in sequential decision-making problems (such as bandits). Consequently, upper bounds for the RHS of (4.3) have already been established for bounded (X_t) , e.g. in Vakili et al. (2021, Corollary 1); these have no explicit dimension dependence and are thus called ‘‘dimension-free’’ (so they are applicable in infinite dimensional settings, where they will

capture some notion of effective dimensionality). For Mercer kernels, such bounds depend on their Mercer eigenvalue decay. Of special interest are the scenarios where most of the eigenvalues (λ_m) are 0 (finite dimensional Euclidean spaces), or decay exponentially or polynomially fast (generally RBF and Laplace kernels, respectively).

4.4. Concentration inequalities

Proposition 7 establishes a very general concentration inequality, whose tightness depends on the form of $e_i(\lambda)$ and the choice of λ . In particular, the previously introduced $e_i(\lambda)$ will yield Bernstein-type and Bennett-type inequalities. Furthermore, the choice of λ is also motivated by the nature of the concentration inequalities, with different choices for the fixed sample size and the fully sequential settings. We provide specific instances of such inequalities. We start with a fixed sample size Bernstein-type inequality. Its proof, which optimizes for λ in Proposition 7 for the Bernstein-specific $e_i(\lambda)$, is deferred to Appendix A.3.

Theorem 9 (Bernstein-type concentration inequality) *Fix $n > 0$, and let (X_1, \dots, X_n) and $(\epsilon_1, \dots, \epsilon_n)$ fulfill Assumption 3. Let ϵ_i attain the Bernstein condition with parameters (σ_i, B) (Assumption 4). If $\sum_{i \leq n} \sigma_i^2 \|G_i\|^2 \leq C_n^2$ almost surely, where $C_n \geq 0$ is deterministic, then*

$$\sup_{t \leq n} \left\| (\rho I + V_t)^{-1/2} M_t \right\| \leq B \log \left(\frac{2}{\delta} \right) + C_n \sqrt{2 \log \left(\frac{2}{\delta} \right)}$$

with probability $1 - \delta$. Equivalently, the following holds for all $r \geq 0$:

$$\mathbb{P} \left(\sup_{t \leq n} \left\| (\rho I + V_t)^{-1/2} M_t \right\| \geq r \right) \leq 2 \exp \left(- \frac{r^2}{2(C_n^2 + Br)} \right).$$

If the noises (ϵ_i) are bounded and (an upper bound on) their variance is known, a Bennett-type inequality will be tighter than a Bernstein-type inequality. The proof of the following result, which can be derived similarly to that of the Bernstein-type inequality, is given in Appendix A.4.

Theorem 10 (Bennett-type concentration inequality) *Fix $n > 0$, and let (X_1, \dots, X_n) and $(\epsilon_1, \dots, \epsilon_n)$ attain Assumption 3. Let $|\epsilon_i| \leq B$ and $\mathbb{V}_{i-1}[\epsilon_i] \leq \sigma_i^2$ almost surely (Assumption 5). If $\sum_{i \leq n} \sigma_i^2 \|G_i\|^2 \leq C_n^2$ almost surely, where $C_n \geq 0$ is deterministic, then*

$$\sup_{t \leq n} \left\| (\rho I + V_t)^{-1/2} M_t \right\| \leq \frac{C_n^2}{B} h^{-1} \left(\frac{B^2}{C_n^2} \log \left(\frac{2}{\delta} \right) \right)$$

with probability at least $1 - \delta$, where $h(u) = (1 + u) \log(1 + u) - u$. Equivalently, the following holds for all $r \geq 0$:

$$\mathbb{P} \left(\sup_{t \leq n} \left\| (\rho I + V_t)^{-1/2} M_t \right\| \geq r \right) \leq 2 \exp \left(- \frac{C_n^2}{B^2} h \left(\frac{Br}{C_n^2} \right) \right).$$

We highlight that both Theorem 9 and Theorem 10 exploit deterministic upper bounds C_n , which require a fixed sample size. These concentration inequalities usually suffice when used to conduct theoretical analyses. However, we should expect these inequalities to be conservative in

practice due to their use of a deterministic upper bound C_n . Furthermore, they do not easily adapt to the fully sequential setting, where sample sizes are random stopping times. For bounded random noises, and at an expense of a logarithmic term, these two shortcomings can be addressed by means of a mixture method argument.

Theorem 11 (Mixed Bennett-type concentration inequality) *Let (X_1, X_2, \dots) and $(\epsilon_1, \epsilon_2, \dots)$ attain Assumption 3, and let ϵ_i attain Assumption 5, i.e., $|\epsilon_i| \leq B$ and $\mathbb{V}_{i-1}[\epsilon_i] \leq \sigma_i^2$. Denote*

$$s_t = \|(\rho I + V_t)^{-\frac{1}{2}} M_t\|, \quad \nu_t = \sum_{i \leq t} \sigma_i^2 \|G_i\|^2.$$

For $\theta > 0$, it holds that

$$\mathbb{P} \left(\sup_t \frac{\left(\frac{\theta}{B^2}\right)^{\frac{\theta}{B^2}}}{\Gamma\left(\frac{\theta}{B^2}\right) \bar{\gamma}\left(\frac{\theta}{B^2}, \frac{\theta}{B^2}\right)} \frac{\Gamma\left(\frac{Bs_t + \nu_t + \theta}{B^2}\right) \bar{\gamma}\left(\frac{Bs_t + \nu_t + \theta}{B^2}, \frac{\nu_t + \theta}{B^2}\right)}{\left(\frac{\nu_t + \theta}{B^2}\right)^{\frac{Bs_t + \nu_t + \theta}{B^2}}} \exp\left(\frac{\nu_t}{B^2}\right) < \frac{2}{\delta} \right) \geq 1 - \delta, \quad (4.4)$$

where $\bar{\gamma}(a, x) := (\int_x^\infty u^{a-1} e^{-u} du) / \Gamma(a)$ is the regularized upper incomplete gamma function.

Its proof is deferred to Appendix A.5, and combines Theorem 6 with the Gamma-Poisson mixture argument from Howard et al. (2021, Proposition 10). Theorem 11 does not offer a closed-form confidence interval, but it can be easily obtained by root finding (Howard et al., 2021, Appendix D). The hyperparameter θ can be adjusted by following the considerations from Howard et al. (2021, Section 3.5). While the probabilistic expression is rather difficult to parse, the width of the confidence interval can be studied using the tools from Howard et al. (2021). In particular, Howard et al. (2021, Proposition 2) establishes that the width of a mixed confidence interval scales roughly (and up to constants) as $\sqrt{\nu_n \log(\nu_n)}$. If $\nu_n \approx C_n^2$, the width of the mixed confidence interval is effectively inflated by a logarithmic factor (up to constants) with respect to the confidence interval given by Theorem 9. Nonetheless, if ν_n is substantially smaller than C_n^2 , the mixed confidence interval will be tighter than the one provided by Theorem 9. Theorem 11 is expected to generally work better in practice, given that the upper bound C_n scales with the worst case scenario. Furthermore, the mixed confidence interval is better suited for the fully sequential setting: at time $t \ll n$, the bound is of order $\sqrt{\nu_t \log(\nu_t)} \ll \nu_n \leq C_n^2$.

Theorem 11 still requires knowledge of the variance of the noises, which is unreasonable in practice. We can make it empirical by coupling it with a concentration inequality for the variance of the noises, at the expense of a small inflation of the logarithmic factor, which stems from applying the union bound to the two concentration inequalities. The result is exhibited in the following theorem, and its proof can be found in Appendix A.6.

Theorem 12 (Mixed empirical Bennett-type concentration bound) *Let both (X_1, X_2, \dots) and $(\epsilon_1, \epsilon_2, \dots)$ attain Assumption 3, and let ϵ_i attain Assumption 5 with constant variance, i.e., $|\epsilon_i| \leq B$ and $\mathbb{V}_{i-1}[\epsilon_i] = \sigma^2$. Denote*

$$s_t = \|(\rho I + V_t)^{-\frac{1}{2}} M_t\|, \quad \hat{\nu}_t = \hat{\sigma}_{u,t,\delta_1}^2 \sum_{i \leq t} \|G_i\|^2,$$

where $\hat{\sigma}_{u,t,\delta_1}$ is a $1 - \delta_1$ upper confidence bound for σ^2 , i.e., $\mathbb{P}(\hat{\sigma}_{u,t,\delta_1}^2 < \sigma^2 \text{ for all } t) \geq 1 - \delta_1$. For $\theta > 0$, the time-uniform concentration inequality (4.4) holds with δ_2 replacing δ in the left hand side, and $\delta_1 + \delta_2$ replacing δ in the right hand side.

An upper confidence sequence for σ^2 can be obtained using the inequalities from [Martinez-Taboada and Ramdas \(2025, Corollary 4.3\)](#), which are sharper than previous results from [Audibert et al. \(2009\)](#) and [Maurer and Pontil \(2009\)](#), and allow for non-constant expectation of the outcomes.

4.5. Comparison to existing work

We present here an extended comparison of our contributions to existing work. For simplicity, we assume for now that the conditional variances of the noises are constant, i.e. $\sigma_i = \sigma$, that the vectors X_i are deterministic such that $\|X_i\| \leq L$ for all i , and that $\rho \geq 1 \vee L^2$. We focus on comparing our Bernstein-type concentration inequality (Theorem 9) to related works.

In particular, note that the dominating term of the confidence interval given by Theorem 9 is

$$C_n \sqrt{2 \log(2/\delta)} = \sigma \sqrt{\sum_{i \leq n} \|G_i\|^2} \sqrt{2 \log(2/\delta)}.$$

In contrast, [Abbasi-Yadkori et al. \(2011, Theorem 1\)](#) establishes an interval of

$$R \sqrt{\log \frac{\det(I + \rho^{-1} V_n)}{\delta^2}},$$

where R is the sub-Gaussian parameter. [Abbasi-Yadkori et al. \(2011, Lemma 11\)](#) proves that for $G'_i := (\rho I + V_{i-1})^{-\frac{1}{2}} X_i$, it follows that

$$\log \det(I + \rho^{-1} V_n) \leq \sum_{i \leq n} \|G'_i\|^2 \leq 2 \log \det(I + \rho^{-1} V_n),$$

where the second inequality holds if $\rho \geq 1 \vee L^2$. Since

$$\frac{1}{2} \|G'_i\|^2 \leq \frac{\rho}{\rho + L^2} \|G'_i\|^2 \leq \|G_i\|^2 \leq \|G'_i\|^2,$$

we finally obtain that

$$\frac{1}{2} \sum_{i \leq n} \|G_i\|^2 \leq \log \det(I + \rho^{-1} V_n) \leq 2 \sum_{i \leq n} \|G_i\|^2.$$

Thus, we see that our process $\sum_{i \leq n} \|G_i\|^2$ and Abbasi-Yadkori's $\log \det(I + \rho^{-1} V_n)$ are within a multiplicative factor of 2 from each other. This also elucidates why Abbasi-Yadkori's bound is sharper when the variance and sub-Gaussian parameter are close (i.e., $R \approx \sigma$), given our slightly bigger constants.

We shall remind the reader that our inequalities are not tailored to sub-Gaussian noises as above, but rather Bernstein condition or bounded noises (where we aim to adapt to the lower variance). Given that other previous efforts have focused on the latter, we will restrict the comparison here to our bounded noise Bernstein-type inequality. First, let us note that our Theorem 9 strictly improves on [Zhou et al. \(2021, Theorem 4.1\)](#), which has an extra logarithmic term in the sample size, on top of looser constants. In the setting of dimension-free Bernstein bounds, we also highlight three recent dimension-free contributions. [Metelli et al. \(2025, Theorem 5.1\)](#) provides a Bernstein-like dimension-free self-normalized concentration inequality, where the square root of

$\log \det(I + \rho^{-1}V_n)$ appears multiplied by a factor of $\sqrt{73 \log \frac{\pi^2(\rho+1)^2}{3\delta}}$, which is clearly larger than our aforementioned factor of $\sqrt{2 \log(2/\delta)}$. [Akhavan et al. \(2025, Theorem 3.2\)](#) also provides a dimension-free Bernstein bound. For a fixed ρ , their inequality scales with $\log \det(I + \rho^{-1}V_n)$, as opposed to its square root (thus being substantially looser for big n).¹ Lastly, [Chugg and Ramdas \(2025\)](#) also developed dimension-free Bernstein bounds. While the bounds are explicit, it is not straightforward to analyze them because they depend on the inverse of a tail decay function and some ratio of eigenvalues.

We do not argue, however, that our inequalities always dominate these recent alternative approaches. Above, we had assumed that the X_i are deterministic so that the sum of $\|G_i\|^2$ was a constant, avoiding the need to assume an upper bound on it. An identical discussion would have also followed if we were seeking deterministic upper bounds for the inequalities, such as the regret analysis for the linear bandit problem exhibited in [Section 5.2](#). Nonetheless, in practical settings, one would like to adapt to the predictable process $\sum_{i \leq n} \|G_i\|^2$. While some of the alternative approaches do so inherently, we need to rely on a mixture argument (as explained in [Section 4.4](#)), and so we suffer from an extra logarithmic factor. While the magnitude of this inflation is expected to be very mild given its logarithmic nature, it implies that our inequalities do not necessarily dominate their competitors.

5. Applications to online linear regression

5.1. Confidence ellipsoids for online linear regression

Ellipsoids naturally appear in the context of linear regression. To be more precise, let us first revisit linear regression in the finite-dimensional case with Gaussian noise (we roughly follow the discussion presented in [Whitehouse et al. \(2025, Section 5.1\)](#) to motivate some applications of our results). That is, let $H = \mathbb{R}^d$, $\mathbf{Y}_t = (Y_1, \dots, Y_t) \in \mathbb{R}^t$, and $\mathbf{X}_t = (X_1, \dots, X_t)^T \in \mathbb{R}^{t \times d}$, such that

$$\mathbf{Y}_t = \mathbf{X}_t \theta^* + \epsilon_{1:t}, \quad \epsilon_{1:t} \sim N(0, \sigma^2 I_t).$$

The least square estimate for $\theta^* \in \mathbb{R}^d$ is given by $\theta_t(0)$, where $\theta_t(\rho) := (\mathbf{X}_t^T \mathbf{X}_t + \rho I)^{-1} \mathbf{X}_t^T \mathbf{Y}_t$. Assuming that $\mathbf{X}_t^T \mathbf{X}_t$ is full rank, it satisfies

$$\left\| (\mathbf{X}_t^T \mathbf{X}_t)^{\frac{1}{2}} (\theta_t(0) - \theta^*) \right\| \sim \chi_d^2.$$

Consequently, in order to conduct inference on θ^* , a confidence set can be taken to be an ellipsoid centered at $\theta_t(0)$ and thresholded at some quantile of χ_d^2 . Nonetheless, such an ellipsoid fails to be a nonasymptotic confidence set if certain parametric assumptions of linear regression are not attained. In contrast, we can consider a more general sequential setting without assuming homoscedastic Gaussian noises, where the samples simply attain [Assumption 3](#), and the Hilbert space H is of arbitrary dimensions. In the sequential setting, confidence sets are often required to hold uniformly over time, and so the problem is conventionally termed “online” linear regression. Online confidence sets can be obtained from our self-normalized inequalities as exhibited in the following corollary; its proof is based on a simple triangle inequality and can be found in [Appendix A.7](#).

1. The dimension-free Bernstein inequality from [Akhavan et al. \(2025, Theorem 3.2\)](#) is a specific instance of their more general [Akhavan et al. \(2025, Theorem 3.1\)](#) after invoking [Faury et al. \(2020, Theorem 1\)](#). While their techniques can certainly lead to other dimension-free concentration inequalities, we restrict the discussion to the inequalities explicitly established therein.

Corollary 13 *Let $Y_t = X_t^T \theta^* + \epsilon_t$, where the random sequences (X_t) and (ϵ_t) fulfill Assumption 3, and $\|\theta^*\| \leq D < \infty$. If $J_t(\delta)$ is a $1 - \delta$ upper confidence bound for $\|(\rho I + V_t)^{-1/2} M_t\|$ obtained from one of Theorem 9, Theorem 10, Theorem 11, or Theorem 12, then*

$$\mathbb{P} \left(\sup_t \left\| (\rho I + V_t)^{1/2} (\theta_t(\rho) - \theta^*) \right\| - J_t(\delta) \leq \rho^{1/2} D \right) \geq 1 - \delta.$$

5.2. Applications to (kernelized) linear bandits

Online linear regression has immediate applications in (kernelized) linear bandits. In the linear bandit problem, a learner repeatedly chooses actions represented by feature vectors and observes noisy rewards that are assumed to depend linearly on an unknown parameter vector. The goal is to balance exploration (learning about the parameter) and exploitation (selecting actions with high expected reward) in order to minimize cumulative regret. The Gaussian Process Upper Confidence Bound (GP-UCB) algorithm achieves this by maintaining confidence ellipsoids around the estimated parameter and selecting the action with the highest optimistic reward estimate (Srinivas et al., 2010; Chowdhury and Gopalan, 2017; Whitehouse et al., 2023).

Mathematically, in each round $t \in [T]$, the learner uses previous observations to select an action $X_t \in \mathcal{X}$, where \mathcal{X} is a bounded subset of H , and then observes feedback $Y_t := \langle X_t, \theta^* \rangle + \epsilon_t$. It is assumed that $\|\theta^*\| \leq D < \infty$. The learner aims to minimize (with high probability) the regret at time T , which is defined as

$$R_T := \sum_{t=1}^T r_t, \quad r_t := \langle x^*, \theta^* \rangle - \langle X_t, \theta^* \rangle,$$

where $x^* := \arg \max_{x \in \mathcal{X}} \langle x, \theta^* \rangle$. Let

$$\Pi_t(h, \eta) := \left\{ f \in H : \left\| (V_t + \rho I)^{1/2} (f - h) \right\| \leq \eta \right\}$$

denote an ellipsoid in H centered at h . Following the optimism principle, the GP-UCB takes

$$(X_t, \tilde{\theta}_t) = \arg \max_{x \in \mathcal{X}, f \in \Pi_{t-1}(\theta_{t-1}(\rho), \eta_{t-1}^{\text{subG}})} \langle f, x \rangle, \quad (5.1)$$

where η_t^{subG} is obtained from a sub-Gaussian concentration inequality (Abbasi-Yadkori et al., 2011). Considering δ as a constant for simplicity, GP-UCB attains the regret bound

$$R_T = \mathcal{O} \left(B \gamma_T(\rho) \sqrt{T} + \sqrt{\rho \gamma_T(\rho) T} \right),$$

for B -bounded random noises (Whitehouse et al., 2023, Theorem 2), given that the bound is proportional to the sub-Gaussian parameter.

We now consider variants of the GP-UCB procedure, where the threshold η_t is obtained in view of any of our novel inequalities. For the fixed Bernstein and Bennett inequalities, the regret bound is of a similar order, but with the bound of the noises replaced by their variance. For the mixed inequalities, we obtain an extra logarithmic factor. We formalize the result in the following corollary; its proof, provided in Appendix A.8, follows standard arguments.

Corollary 14 *Let (X_1, \dots, X_T) and $(\epsilon_1, \dots, \epsilon_T)$ attain Assumption 3, with ϵ_i attaining Assumption 4 or Assumption 5 with constant σ . Consider variants of the GP-UCB algorithm with η_t taken as $J_t(\delta) + \rho^{1/2}D$, where $J_t(\delta)$ is a $1 - \delta$ upper confidence bound for $\|(\rho I + V_t)^{-1/2}M_t\|$ obtained from one of our self-normalized concentration inequalities. If $J_t(\delta)$ is defined following Theorem 9 (if Assumption 4 holds) or Theorem 10 (if Assumption 5 holds), then*

$$R_T = \mathcal{O} \left(\sigma \gamma_T(\rho) \sqrt{T} + \sqrt{\rho \gamma_T(\rho) T} \right).$$

If Assumption 5 holds, and $J_t(\delta)$ is defined following Theorem 11 or Theorem 12, the above guarantee holds up to logarithmic factors.

Under Assumption 4 (sub-exponential noise distributions), our analysis yields regret bounds that go beyond the commonly studied bounded or sub-Gaussian noise regimes. To the best of our knowledge, such guarantees have not appeared previously in the literature. Under Assumption 5, these variance-dependent type bounds are usually referred to as “second-order” regret guarantees (Kirschner and Krause, 2018; Zhang et al., 2021; Xu et al., 2023; Jun and Kim, 2024; Pacchiano, 2025), with the case $H = \mathbb{R}^d$ and constant ρ being predominantly studied in the literature. Given that $\gamma_T(\rho) = \mathcal{O}(d)$ up to logarithmic factors, Corollary 14 yields a regret bound of $\mathcal{O}(d\sigma\sqrt{T} + \sqrt{dT})$ up to logarithmic factors, immediately recovering a regret bound comparable to many existing works (Zhou et al., 2021; Zhou and Gu, 2022; Kim et al., 2022; Zhao et al., 2023) for constant variance.

5.3. Experiments

In order to elucidate the empirical differences between our variance-dependent inequalities and the predominant sub-Gaussian ones, we run an ablation study on the GP-UCB algorithm for the linear bandit problem using the sub-Gaussian inequality from Abbasi-Yadkori et al. (2011). We also evaluate our mixed Bennett inequality and our empirical mixed Bennett inequality. The code can be found at https://github.com/DMartinezT/self_normalized.

More specifically, we consider a bandit experiment, where at each round t an action X_t is taken following the UCB procedure introduced in (5.1). The covariates are RBF kernel embeddings of one-dimensional points, where the kernel length scale set to 0.01. The bound of the kernel is naturally 1. We take $\rho = 0.05$, $\delta = 2\delta_1 = 2\delta_2 = 0.1$, and the mixing hyperparameter $\theta = 1$. The true regression function is given by a weighted sum of 50 embeddings of different points, which are randomly drawn concentrated around two modes (such that it has two local maxima for the sake of visualization). For the empirical mixed Bennett-type inequality, $\hat{\sigma}_{u,t,\delta_1}^2$ is obtained using the upper confidence sequence from Martinez-Taboada and Ramdas (2025) with $\hat{\mu}_i$ and $\bar{\mu}_i$ taken as the evaluation of $\theta_i(\rho)$ on X_i . No effort has been put into optimizing the hyperparameter choices.

We consider four experimental settings, where the outcomes correspond to the evaluation of the regression function added to different random noise, which follows either a uniform distribution or beta distributions with different parameters. All of them are rescaled to lie in the interval $(-1, 1)$. Figure 1 illustrates the true regression function, the training points, and the estimated mean of the bandit algorithm after 500 rounds. We generally see that the smaller the variance (with respect to the scale of the regression function and the bound of the noise), the better our inequalities perform in comparison to the sub-Gaussian inequalities from Abbasi-Yadkori et al. (2011). In particular, we observe that for comparatively small variance, i.e. plots (II), (III), and (IV), our inequalities

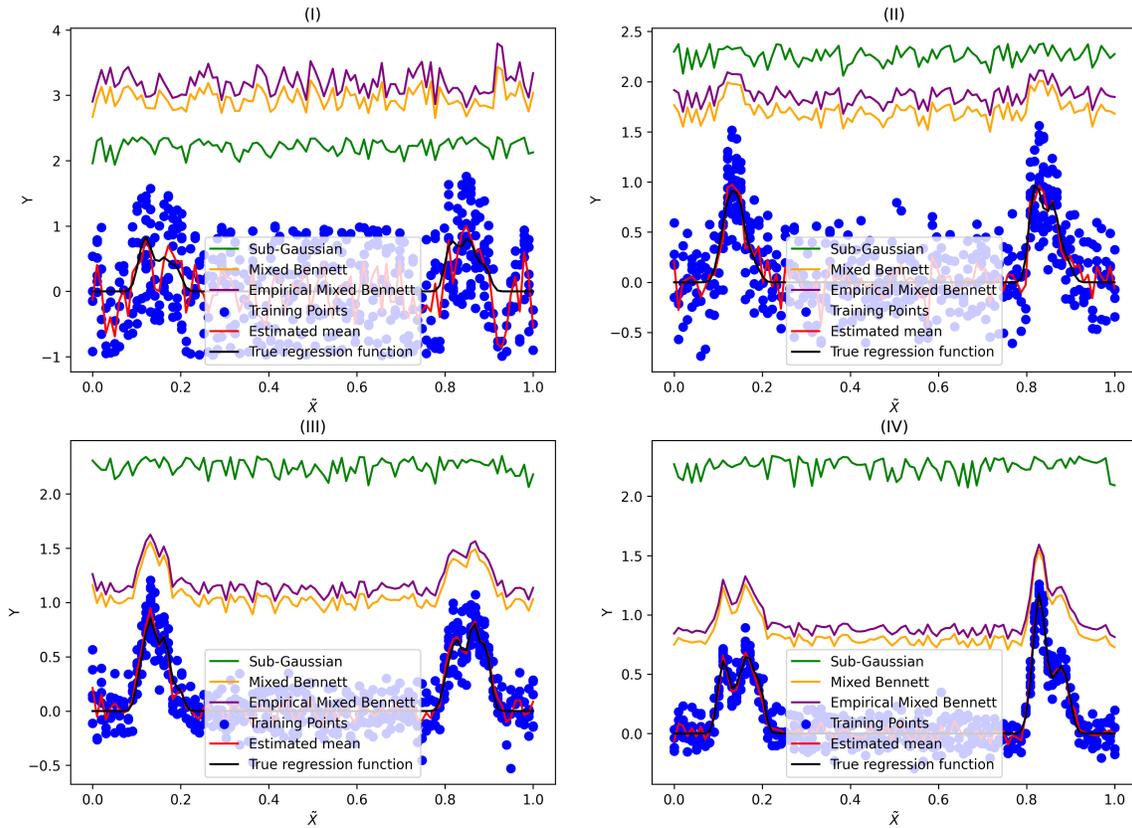


Figure 1: Illustration of the optimistic upper confidence bounds for the regression function after 500 rounds using sub-Gaussian, mixed Bennett, and empirical mixed Bennett inequalities for noises following (I) a rescaled uniform distribution, (II) a rescaled (5, 5)-beta distribution, (III) a rescaled (20, 20)-beta distribution, and (IV) a rescaled (50, 50)-beta distribution. Training points are drawn following a UCB procedure, with the covariates $X_t = k(\cdot, \tilde{X}_t)$ illustrated in the original space (pre-embedded in the RKHS).

lead to sharper bounds of the regression function, elucidating the empirical gains of using variance-aware inequalities. In contrast, our inequalities are empirically outperformed by the sub-Gaussian approach for uniform-distributed noises (comparatively larger variance), where the difference between the variance and bound is not large enough to justify the extra logarithmic term of our mixed inequalities.

6. Conclusion

We have proposed novel concentration inequalities for vector-valued self-normalized processes. These include anytime-valid Bernstein and Bennett type inequalities tailored to a fixed sample size (Proposition 9 and Proposition 10, respectively), a mixed Bennett-type inequality that is well suited for randomly stopped sample sizes, and an empirical mixed Bennett-type inequality that does not

require knowledge of the variance in advance. These inequalities build on the theoretical tools from [Pinelis \(1994\)](#), thus being fundamentally different to previous vector-valued self-normalized inequalities. We have further explored the immediate consequences of our inequalities in the (kernelized) bandit setting, both theoretically in the form of second order regret bounds and empirically.

There are several directions for future work. First, the proposed inequalities can have interesting applications beyond linear bandits. Natural extensions may include reinforcement learning ([Yang et al., 2020](#); [Vakili and Olkhovskaya, 2024](#)), safe Bayesian optimization ([Chowdhury and Gopalan, 2017](#); [Amani et al., 2020](#)), and autoregressive models ([Darolles et al., 2006](#); [Pacurar, 2008](#); [Shao, 2015](#); [Agarwal et al., 2018](#)). Second, our empirical Bennett inequality assumes constant variance for the noise, so we can exploit the upper confidence sequences from [Martinez-Taboada and Ramdas \(2025, Corollary 4.3\)](#). However, there is evident interest in bandit algorithms that adapt to heteroscedastic noises ([Kirschner and Krause, 2018](#); [Kim et al., 2022](#); [Zhao et al., 2023](#)). Extending [Martinez-Taboada and Ramdas \(2025, Corollary 4.3\)](#) to a covariate-dependent inequality would immediately yield heteroscedastic-noise guarantees when in conjunction with [Proposition 11](#). Lastly, our inequalities could potentially be sharpened if the loose bound $V_{t-1} \preceq V_t$ was to be avoided in the supermartingale construction (see proof of [Theorem 6](#)). This limitation may inflate the growth of the inequalities by a logarithmic term (see [Appendix D](#) for a detailed presentation), which implies that sub-Gaussian inequalities are sharper than our mixed inequalities if the variance of the noise is not substantially smaller than the bound of the noise. Finding a refined construction that circumvents this limitation constitutes an important and difficult open direction for future work.

Acknowledgments

DMT acknowledges that the project that gave rise to these results received the support of a fellowship from ‘la Caixa’ Foundation (ID 100010434). The fellowship code is LCF/BQ/EU22/11930075. AR was funded by NSF grant DMS-2310718.

References

- Yasin Abbasi-Yadkori. *Online Learning for Linearly Parametrized Control Problems*. PhD thesis, University of Alberta, 2013.
- Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. *Advances in Neural Information Processing Systems*, 2011.
- Anish Agarwal, Muhammad Jehangir Amjad, Devavrat Shah, and Dennis Shen. Model agnostic time series analysis via matrix estimation. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 2018.
- Arya Akhavan, Amitis Shidani, Alex Ayoub, and David Janz. Bernstein-type dimension-free concentration for self-normalised martingales. *arXiv preprint arXiv:2507.20982*, 2025.
- Sanae Amani, Mahnoosh Alizadeh, and Christos Thrampoulidis. Regret bound for safe Gaussian process bandit optimization. *Learning for Dynamics and Control*, 2020.
- Jean-Yves Audibert, Rémi Munos, and Csaba Szepesvári. Exploration–exploitation tradeoff using variance estimates in multi-armed bandits. *Theoretical Computer Science*, 410(19):1876–1902, 2009.

- Bernard Bercu and Abderrahmen Touati. Exponential inequalities for self-normalized martingales with applications. *The Annals of Applied Probability*, 18(5):1848–1869, 2008.
- Bernard Bercu and Taieb Touati. New insights on concentration inequalities for self-normalized martingales. *Electronic Communications in Probability*, 24:1 – 12, 2019.
- Sayak Ray Chowdhury and Aditya Gopalan. On kernelized multi-armed bandits. *International Conference on Machine Learning*, 2017.
- Ben Chugg and Aaditya Ramdas. A variational approach to dimension-free self-normalized concentration. *arXiv preprint arXiv:2508.06483*, 2025.
- Varsha Dani, Thomas P Hayes, and Sham M Kakade. Stochastic linear optimization under bandit feedback. *Conference on Learning Theory*, 2008.
- DA Darling and Herbert Robbins. Some further remarks on inequalities for sample sums. *Proceedings of the National Academy of Sciences*, 60(4):1175–1182, 1968.
- Donald A Darling and Herbert Robbins. Iterated logarithm inequalities. *Proceedings of the National Academy of Sciences*, 57(5):1188–1192, 1967.
- Serge Darolles, Christian Gouriéroux, and Joann Jasiak. Structural Laplace transform and compound autoregressive models. *Journal of Time Series Analysis*, 27(4):477–503, 2006.
- Victor de la Peña, Michael J Klass, and Tze Leung Lai. Self-normalized processes: exponential inequalities, moment bounds and iterated logarithm laws. *The Annals of Probability*, 32(3):1902 – 1933, 2004.
- Victor de la Peña, Michael J Klass, and Tze Leung Lai. Pseudo-maximization and self-normalized processes. *Probability Surveys*, 4:172 – 192, 2007.
- Victor de la Peña, Michael J Klass, and Tze Leung Lai. Theory and applications of multivariate self-normalized processes. *Stochastic Processes and their Applications*, 119(12):4210–4227, 2009a.
- Victor de la Peña, Tze Leung Lai, and Qi-Man Shao. *Self-normalized Processes: Limit Theory and Statistical Applications*. Springer, 2009b.
- Louis Faury, Marc Abeille, Clément Calauzènes, and Olivier Fercoq. Improved optimistic algorithms for logistic bandits. In *International Conference on Machine Learning*, pages 3052–3060. PMLR, 2020.
- Hamish Flynn and David Reeb. Tighter confidence bounds for sequential kernel regression. In *International Conference on Artificial Intelligence and Statistics*, pages 3844–3852. PMLR, 2025.
- Hamish Flynn, David Reeb, Melih Kandemir, and Jan R Peters. Improved algorithms for stochastic linear bandits using tail bounds for martingale mixtures. *Advances in Neural Information Processing Systems*, 36:45102–45136, 2023.
- Jiafan He, Heyang Zhao, Dongruo Zhou, and Quanquan Gu. Nearly minimax optimal reinforcement learning for linear markov decision processes. In *International Conference on Machine Learning*, pages 12790–12822. PMLR, 2023.

- Steven R Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Time-uniform Chernoff bounds via nonnegative supermartingales. *Probability Surveys*, 17:257–317, 2020.
- Steven R Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Time-uniform, nonparametric, nonasymptotic confidence sequences. *The Annals of Statistics*, 49(2):1055–1080, 2021.
- Kwang-Sung Jun and Jungtaek Kim. Noise-adaptive confidence sets for linear bandits and application to bayesian optimization. *International Conference on Machine Learning*, 2024.
- Yeoneung Kim, Insoon Yang, and Kwang-Sung Jun. Improved regret analysis for variance-adaptive linear bandits and horizon-free linear mixture MDPs. *Advances in Neural Information Processing Systems*, 2022.
- Johannes Kirschner and Andreas Krause. Information directed sampling and bandits with heteroscedastic noise. *Conference on Learning Theory*, 2018.
- A Kolmogoroff. Über das gesetz des iterierten logarithmus. *Mathematische Annalen*, 101(1):126–135, 1929.
- Michel Ledoux and Michel Talagrand. *Probability in Banach Spaces: isoperimetry and processes*. Springer Science & Business Media, 2013.
- Diego Martinez-Taboada and Aaditya Ramdas. Sharp empirical Bernstein bounds for the variance of bounded random variables. *arXiv preprint arXiv:2505.01987*, 2025.
- Diego Martinez-Taboada and Aaditya Ramdas. Empirical Bernstein in smooth Banach spaces. *Annals of Applied Probability*, 2026.
- Andreas Maurer and Massimiliano Pontil. Empirical Bernstein bounds and sample variance penalization. *Conference on Learning Theory*, 2009.
- Alberto Maria Metelli, Simone Drago, and Marco Mussi. Generalized kernelized bandits: Self-normalized Bernstein-like dimension-free inequality and regret bounds. *arXiv preprint arXiv:2508.01681*, 2025.
- Aldo Pacchiano. Second order bounds for contextual bandits with function approximation. *International Conference on Learning Representations*, 2025.
- Maria Pacurar. Autoregressive conditional duration models in finance: a survey of the theoretical and empirical literature. *Journal of Economic Surveys*, 22(4):711–751, 2008.
- Iosif Pinelis. An approach to inequalities for the distributions of infinite-dimensional martingales. In *Probability in Banach Spaces, 8: Proceedings of the Eighth International Conference*, pages 128–134. Springer, 1992.
- Iosif Pinelis. Optimum bounds for the distributions of martingales in Banach spaces. *The Annals of Probability*, pages 1679–1706, 1994.
- Aaditya Ramdas, Johannes Ruf, Martin Larsson, and Wouter Koolen. Admissible anytime-valid sequential inference must rely on nonnegative martingales. *arXiv preprint arXiv:2009.03167*, 2020.

- Aaditya Ramdas, Peter Grünwald, Vladimir Vovk, and Glenn Shafer. Game-theoretic statistics and safe anytime-valid inference. *Statistical Science*, 38(4):576–601, 2023.
- Xiaofeng Shao. Self-normalization for time series: a review of recent developments. *Journal of the American Statistical Association*, 110(512):1797–1817, 2015.
- Niranjan Srinivas, Andreas Krause, Sham M Kakade, and Matthias Seeger. Gaussian process optimization in the bandit setting: no regret and experimental design. *International Conference on Machine Learning*, 2010.
- William F Stout. A martingale analogue of Kolmogorov’s law of the iterated logarithm. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 15(4):279–290, 1970.
- Sattar Vakili and Julia Olkhovskaya. Kernelized reinforcement learning with order optimal regret bounds. *Advances in Neural Information Processing Systems*, 2024.
- Sattar Vakili, Kia Khezeli, and Victor Picheny. On information gain and regret bounds in Gaussian process bandits. *International Conference on Artificial Intelligence and Statistics*, 2021.
- Jean Ville. *Etude critique de la notion de collectif*. Gauthier-Villars Paris, 1939.
- Martin J Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge university press, 2019.
- Justin Whitehouse, Aaditya Ramdas, and Steven Z Wu. On the sublinear regret of GP-UCB. *Advances in Neural Information Processing Systems*, 2023.
- Justin Whitehouse, Ben Chugg, Diego Martinez-Taboada, and Aaditya Ramdas. Mean estimation in Banach spaces under infinite variance and martingale dependence. *arXiv preprint arXiv:2411.11271*, 2024.
- Justin Whitehouse, Zhiwei Steven Wu, and Aaditya Ramdas. Time-uniform self-normalized concentration for vector-valued processes. *The Annals of Applied Probability*, 2025.
- Ruitu Xu, Yifei Min, and Tianhao Wang. Noise-adaptive Thompson sampling for linear contextual bandits. *Advances in Neural Information Processing Systems*, 2023.
- Zhuoran Yang, Chi Jin, Zhaoran Wang, Mengdi Wang, and Michael I. Jordan. On function approximation in reinforcement learning: optimism in the face of large state spaces. *Advances in Neural Information Processing Systems*, 2020.
- Zihan Zhang, Jiaqi Yang, Xiangyang Ji, and Simon S Du. Improved variance-aware confidence sets for linear bandits and linear mixture MDP. *Advances in Neural Information Processing Systems*, 2021.
- Heyang Zhao, Jiafan He, Dongruo Zhou, Tong Zhang, and Quanquan Gu. Variance-dependent regret bounds for linear bandits and reinforcement learning: Adaptivity and computational efficiency. *Conference on Learning Theory*, 2023.
- Dongruo Zhou and Quanquan Gu. Computationally efficient horizon-free reinforcement learning for linear mixture MDPs. *Advances in Neural Information Processing Systems*, 2022.

Dongruo Zhou, Quanquan Gu, and Csaba Szepesvari. Nearly minimax optimal reinforcement learning for linear mixture markov decision processes. *Conference on Learning Theory*, 2021.

Ingvar Ziemann. A vector Bernstein inequality for self-normalized martingales. *Transactions on Machine Learning Research*, 2025.

Appendix A. Proofs

A.1. Proof of Proposition 7

Extend S_t to $t = 0$ with $\lambda_0 = 0$ and $X_0 = 0$. It follows that S_t is a nonnegative supermartingale with $S_0 = 1$. Consequently, Ville's inequality (Fact 1) yields

$$\mathbb{P}\left(\sup_t S_t \geq \frac{1}{\delta}\right) \leq \mathbb{E}[S_0]\delta = \delta.$$

Given that $e^u \leq 2 \cosh(u)$ for all $u \in \mathbb{R}$, it follows from Theorem 6 that

$$\frac{1}{2} \exp\left(\left\|\lambda(\rho I + V_t)^{-1/2} M_t\right\|\right) \exp\left(-\sum_{i=1}^t e_i(\lambda)\right)$$

is dominated by S_t . Thus, with probability $1 - \delta$, and simultaneously for all $t \geq 1$,

$$\frac{1}{2} \exp\left(\left\|\lambda(\rho I + V_t)^{-1/2} M_t\right\|\right) \exp\left(-\sum_{i=1}^t e_i(\lambda)\right) \leq \frac{1}{\delta}.$$

Taking logarithms and dividing both sides by λ , it follows that

$$\left\|(\rho I + V_t)^{-1/2} M_t\right\| \leq \frac{\sum_{i=1}^t e_i(\lambda) + \log\left(\frac{2}{\delta}\right)}{\lambda}.$$

A.2. Proof of Corollary 8

Throughout, we refer to a stochastic process as l_0 -sub- ψ following Howard et al. (2021, Definition 1). Denote

$$\psi_{G,B}(\lambda) := \frac{\lambda^2}{2(1 - \lambda B)} \sigma_t^2, \quad 0 \leq \lambda < \frac{1}{B}; \quad \psi_{P,B}(\lambda) := \frac{e^{\lambda B} - \lambda B - 1}{B^2}, \quad \lambda \geq 0.$$

Based on $e^x \leq 2 \cosh x$, Theorem 6 implies that

$$\tilde{S}_t = \exp\left(\lambda \left\|(\rho I + V_t)^{-1/2} M_t\right\| - \sum_{i \leq t} e_i(\lambda)\right)$$

is dominated by the nonnegative supermartingale $2S_t$. If Assumption 4 holds, then

$$\sum_{i \leq t} e_i(\lambda) \leq \psi_{G,B}(\lambda) \sum_{i \leq t} \|G_i\|^2 \sigma_i^2,$$

and $\|(\rho I + V_t)^{-1/2} M_t\|$ is 2-sub- $\psi_{G,B}$ with variance process $\sum_{i \leq t} \|G_i\|^2 \sigma_i^2$. If Assumption 5 holds,

$$\sum_{i \leq t} e_i(\lambda) \leq \psi_{P,B}(\lambda) \sum_{i \leq t} \|G_i\|^2 \sigma_i^2,$$

and $\|(\rho I + V_t)^{-1/2} M_t\|$ is 2-sub- $\psi_{P,B}$ with variance process $\sum_{i \leq t} \|G_i\|^2 \sigma_i^2$. In either case, Howard et al. (2021, Corollary 1) can be applied to conclude the result, in view of $\psi_{G,B}(\lambda) \approx \lambda^2/2 \approx \psi_{P,B}(\lambda)$ as $\lambda \downarrow 0$.

A.3. Proof of Theorem 9

It follows from Proposition 7 that

$$\|(\rho I + V_t)^{-1/2} M_t\| \leq \frac{\sum_{i=1}^t e_i(\lambda) + \log\left(\frac{2}{\delta}\right)}{\lambda}.$$

simultaneously for all $t \geq 1$ with probability $1 - \delta$. As observed in Section 4.2,

$$e_i(\lambda) \leq \frac{\lambda^2}{2(1-\lambda B)} \|G_i\|^2 \sigma_i^2$$

for $\lambda \in (0, \frac{1}{B})$, and so

$$\begin{aligned} \sup_{t \leq n} \|(\rho I + V_t)^{-1/2} M_t\| &\leq \sup_{t \leq n} \frac{\frac{\lambda^2}{2(1-\lambda B)} \sum_{i=1}^t \sigma_i^2 \|G_i\|^2 + \log\left(\frac{2}{\delta}\right)}{\lambda} \\ &\leq \frac{\frac{\lambda^2}{2(1-\lambda B)} C_n^2 + \log\left(\frac{2}{\delta}\right)}{\lambda}. \end{aligned}$$

Now we optimize over $\lambda \in (0, 1/B)$. Denote $L := \log\left(\frac{2}{\delta}\right)$. Consider the function

$$f(\lambda) := \frac{\frac{\lambda^2}{2(1-\lambda B)} C_n^2 + L}{\lambda} = \frac{C_n^2 \lambda}{2(1-\lambda B)} + \frac{L}{\lambda}, \quad \lambda \in (0, 1/B).$$

Writing f in terms of $x := \lambda^{-1}$ gives

$$f\left(\frac{1}{x}\right) = \frac{C_n^2}{2(x-B)} + Lx, \quad x > B.$$

Differentiating with respect to x yields

$$\frac{d}{dx} \left(\frac{C_n^2}{2(x-B)} + Lx \right) = -\frac{C_n^2}{2(x-B)^2} + L,$$

and equating it to 0 leads to the minimizer

$$x^* = B + \sqrt{\frac{C_n^2}{2L}}.$$

Thus

$$\begin{aligned} \min_{0 < \lambda < 1/B} f(\lambda) &= f\left(\frac{1}{x^*}\right) = Lx^* + \frac{C_n^2}{2(x^* - B)} \\ &= L\left(B + \sqrt{\frac{C_n^2}{2L}}\right) + \frac{C_n^2}{2\sqrt{C_n^2/(2L)}} \\ &= BL + \sqrt{2C_n^2 L}. \end{aligned}$$

We thus conclude that

$$\sup_{t \leq n} \left\| (\rho I + V_t)^{-1/2} M_t \right\| \leq B \log\left(\frac{2}{\delta}\right) + \sqrt{2C_n^2 \log\left(\frac{2}{\delta}\right)}.$$

A.4. Proof of Theorem 10

It follows from Proposition 7 that

$$\left\| (\rho I + V_t)^{-1/2} M_t \right\| \leq \frac{\sum_{i=1}^t e_i(\lambda) + \log\left(\frac{2}{\delta}\right)}{\lambda}. \quad (\text{A.1})$$

simultaneously for all $t \geq 1$ with probability $1 - \delta$. As observed in Section 4.2,

$$e_i(\lambda) \leq \frac{e^{\lambda B} - \lambda B - 1}{B^2} \sigma_i^2 \|G_i\|^2,$$

and so

$$\begin{aligned} \sup_{t \leq n} \left\| (\rho I + V_t)^{-1/2} M_t \right\| &\leq \sup_{t \leq n} \frac{\frac{e^{\lambda B} - \lambda B - 1}{B^2} \sum_{i=1}^t \sigma_i^2 \|G_i\|^2 + \log\left(\frac{2}{\delta}\right)}{\lambda} \\ &\leq \frac{\frac{e^{\lambda B} - \lambda B - 1}{B^2} C_n^2 + \log\left(\frac{2}{\delta}\right)}{\lambda}. \end{aligned}$$

Denote

$$L := \log\left(\frac{2}{\delta}\right), \quad s := B\lambda, \quad \phi(s) := e^s - s - 1, \quad A := \frac{C_n^2}{B^2}.$$

It follows that

$$\sup_{t \leq n} \left\| (\rho I + V_t)^{-1/2} M_t \right\| \leq Bg(s), \quad g(s) := \frac{A\phi(s) + L}{s}.$$

The best bound is obtained by minimizing $g(s)$ over $s > 0$. Differentiating gives

$$g'(s) = \frac{A\phi'(s)s - (A\phi(s) + L)}{s^2}.$$

Setting $g'(s) = 0$ yields $As\phi'(s) = A\phi(s) + L$. Recalling $\phi(s) = e^s - s - 1$ and $\phi'(s) = e^s - 1$, this becomes

$$s(e^s - 1) = e^s - s - 1 + \frac{L}{A}.$$

Let now $u := e^s - 1$, so $s = \log(1 + u)$. Substituting and simplifying leads to

$$h(u) = \frac{B^2}{C_n^2} L,$$

where $h(u) = (1+u)\log(1+u) - u$. Let u^* denote such an optimal value for u , and analogously for s^* . Using that $As^*\phi'(s^*) = A\phi(s^*) + L$, we compute

$$g(s^*) = \frac{A\phi(s^*) + L}{s^*} = \frac{As^*\phi'(s^*)}{s^*} = A\phi'(s^*) = A(e^{s^*} - 1) = Au^*.$$

Therefore the optimized upper bound evaluates to

$$Bg(s^*) = BAu^* = \frac{C_n^2}{B}u^*.$$

Recalling that u^* is the unique solution of $h(u) = \frac{B^2}{C_n^2} \log \frac{2}{\delta}$, we conclude that

$$\sup_{t \leq n} \left\| (\rho I + V_t)^{-1/2} M_t \right\| \leq \frac{C_n^2}{B} h^{-1} \left(\frac{B^2}{C_n^2} \log \left(\frac{2}{\delta} \right) \right)$$

holds with probability at least $1 - \delta$.

A.5. Proof of Theorem 11

The proof of this corollary primarily relies on the combination of Theorem 6 and the Gamma-Poisson mixture argument from Howard et al. (2021, Proposition 10).

Given that $e^u \leq 2 \cosh(u)$ for all $u \in \mathbb{R}$, it follows that

$$\epsilon_t := \exp \left(\left\| \lambda(\rho I + V_t)^{-1/2} M_t \right\| - \sum_{i=1}^t e_i(\lambda) \right) \leq 2S_t,$$

where S_t is the nonnegative supermartingale established in Theorem 6 with $S_0 = 1$. Denoting $\psi_{P,B}(\lambda) := B^{-2}(e^{\lambda B} - \lambda B - 1)$, and in view of

$$e_i(\lambda) \leq \psi_{P,B}(\lambda) \sigma_i^2 \|G_i\|^2,$$

the process ϵ_t falls under Howard et al. (2021, Definition 1) as 2-sub- $\psi_{P,B}$ with variance process $\sum_{i \leq t} \sigma_i^2 \|G_i\|^2$. Thus, Howard et al. (2021, Proposition 10) can be invoked to yield the corollary.

A.6. Proof of Theorem 12

The corollary is obtained in view of Theorem 11 and a union bound. Let A be the event that

$$\hat{\sigma}_{u,t,\delta_1}^2 < \sigma^2$$

for some $t \geq 0$, and C the event that

$$\frac{\left(\frac{\theta}{B^2}\right)^{\frac{\theta}{B^2}} \Gamma\left(\frac{Bs_t + \nu_t + \theta}{B^2}\right) \bar{\gamma}\left(\frac{Bs_t + \nu_t + \theta}{B^2}, \frac{\nu_t + \theta}{B^2}\right)}{\Gamma\left(\frac{\theta}{B^2}\right) \bar{\gamma}\left(\frac{\theta}{B^2}, \frac{\theta}{B^2}\right) \left(\frac{\nu_t + \theta}{B^2}\right)^{\frac{Bs_t + \nu_t + \theta}{B^2}}} \exp\left(\frac{\nu_t}{B^2}\right) > \frac{2}{\delta_2} \quad (\text{A.2})$$

for some $t \geq 0$, where $\nu_t = \sigma^2 \sum_{i \leq t} \|G_i\|^2$.

We observe that $\mathbb{P}(A) \leq \delta_1$ by assumption, and $\mathbb{P}(C) \leq \delta_2$ by Theorem 11. Thus $P(A \cup C) \leq \delta_1 + \delta_2 \leq \delta$ in view of the union bound. Thus, $P(\bar{A} \cap \bar{C}) = 1 - P(A \cup C) \geq 1 - \delta$.

Denote the LHS of (A.2) by ϵ_t , and its empirical counterpart

$$\frac{\left(\frac{\theta}{B^2}\right)^{\frac{\theta}{B^2}}}{\Gamma\left(\frac{\theta}{B^2}\right) \bar{\gamma}\left(\frac{\theta}{B^2}, \frac{\theta}{B^2}\right)} \frac{\Gamma\left(\frac{Bs_t + \hat{\nu}_t + \theta}{B^2}\right) \bar{\gamma}\left(\frac{Bs_t + \hat{\nu}_t + \theta}{B^2}, \frac{\hat{\nu}_t + \theta}{B^2}\right)}{\left(\frac{\hat{\nu}_t + \theta}{B^2}\right)^{\frac{Bs_t + \hat{\nu}_t + \theta}{B^2}}} \exp\left(\frac{\hat{\nu}_t}{B^2}\right)$$

by $\hat{\epsilon}_t$. Since both expressions are obtained as mixtures of functions that are decreasing on σ , if \bar{A} holds, then $\hat{\epsilon}_t \leq \epsilon_t$. Furthermore, if \bar{C} holds, then $\sup_t \epsilon_t \leq \frac{2}{\delta_2}$. Together, these imply that $\sup_t \hat{\epsilon}_t \leq \frac{2}{\delta_2}$ with probability at least $1 - \delta$.

A.7. Proof of Corollary 13

It suffices to observe that

$$\begin{aligned} \left\| (V_t + \rho I)^{1/2} (\theta_t(\rho) - \theta^*) \right\| &= \left\| (V_t + \rho I)^{1/2} \left((V_t + \rho I)^{-1} \mathbf{X}_t^T \epsilon_{1:t} - \rho(\rho I + V_t)^{-1} \theta^* \right) \right\| \\ &\leq \left\| (V_t + \rho I)^{-\frac{1}{2}} \mathbf{X}_t^T \epsilon_{1:t} \right\| + \left\| \rho(\rho I + V_t)^{-\frac{1}{2}} \theta^* \right\| \\ &= \left\| (V_t + \rho I)^{-\frac{1}{2}} M_t \right\| + \left\| \rho(\rho I + V_t)^{-\frac{1}{2}} \theta^* \right\| \\ &\leq J_t(\delta) + \rho^{1/2} D. \end{aligned}$$

A.8. Proof of Corollary 14

Define $\eta_0(\delta) = \rho^{1/2} D$, $\theta_0(\rho) = 0$ and consider $V_0 = 0$, so that $\Pi_0(\theta_0(\rho), \eta_0)$ is the ball centered at 0 of radius D containing θ^* . For $t \geq 1$, let $\eta_t(\delta) = J_t(\delta) + \rho^{1/2} D$, where $J_t(\delta)$ is obtained from one of our concentration inequalities. Taking

$$(X_t, \tilde{\theta}_t) = \arg \max_{x \in \mathcal{X}, f \in \Pi_{t-1}(\theta_{t-1}(\rho), \eta_{t-1})} \langle f, x \rangle,$$

the regret can be upper bounded with probability $1 - \delta$ as

$$\begin{aligned} r_t &= \langle \theta^*, x^* \rangle - \langle \theta^*, X_t \rangle \stackrel{(i)}{\leq} \langle \tilde{\theta}_t, X_t \rangle - \langle \theta^*, X_t \rangle \\ &= \langle \tilde{\theta}_t - \theta_{t-1}(\rho), X_t \rangle - \langle \theta_{t-1}(\rho) - \theta^*, X_t \rangle \\ &\stackrel{(ii)}{\leq} \left\| (V_{t-1} + \rho I)^{-1/2} X_t \right\| \left(\left\| (V_{t-1} + \rho I)^{1/2} (\tilde{\theta}_t - \theta_{t-1}(\rho)) \right\| + \left\| (V_{t-1} + \rho I)^{1/2} (\theta_{t-1}(\rho) - \theta^*) \right\| \right) \\ &\stackrel{(iii)}{\leq} 2\eta_{t-1}(\delta) \left\| (V_{t-1} + \rho I)^{-1/2} X_t \right\|, \end{aligned}$$

where (i) follows from the definition of $(X_t, \tilde{\theta}_t)$ together with the fact that we are considering a high probability event where $\theta^* \in \Pi(\theta_t(\rho), \eta_t)$ for every $t \geq 0$, (ii) follows from Cauchy-Schwarz inequality, and (iii) follows from the definition of the ellipsoid $\Pi(\theta_{t-1}(\rho), \eta_{t-1})$ along with the fact

that both $\tilde{\theta}_t$ and θ^* belong to it. Hence,

$$\begin{aligned} R_T &= \sum_{t=1}^T r_t \stackrel{(i)}{\leq} \sqrt{T \sum_{t=1}^T r_t^2} \\ &\stackrel{(ii)}{\leq} \sqrt{T \sum_{t=1}^T 4\eta_{t-1}(\delta)^2 \left\| (V_{t-1} + \rho I)^{-1/2} X_t \right\|^2} \\ &\stackrel{(iii)}{\leq} 4\eta_T(\delta) \sqrt{T\gamma_T(\rho)} \end{aligned}$$

with probability $1 - \delta$, where (i) follows from the Cauchy-Schwarz inequality, (ii) follows from the elliptical potential lemma and (iii) is obtained given that $t \mapsto \eta_t(\delta)$ is non-decreasing.

If $\eta_T(\rho)$ is obtained from Theorem 9, then

$$J_T(\rho) \leq B \log\left(\frac{2}{\delta}\right) + \sigma \sqrt{\gamma_T(\rho)} \sqrt{2 \log\left(\frac{2}{\delta}\right)} = \mathcal{O}\left(\sigma \sqrt{\gamma_T(\rho)}\right).$$

If $\eta_T(\rho)$ is obtained from Theorem 10, it is not closed form. However, it is well known that

$$h^{-1}(y) \leq \sqrt{2y} + \frac{y}{3},$$

and so it follows that $J_T(\rho)$ is upper bounded by

$$C_T \sqrt{2 \log(2/\delta)} + \frac{B}{3} \log(2/\delta) \leq \sigma \sqrt{8\gamma_T(\rho) \log(2/\delta)} + \frac{B}{3} \log(2/\delta) = \mathcal{O}\left(\sigma \sqrt{\gamma_T(\rho)}\right),$$

where in the first inequality we used that $C_T := \sigma \sqrt{4\gamma_T(\rho)}$ is an upper bound on $\sigma \sqrt{\sum_{t \leq T} \|G_i\|^2}$ by definition of $\gamma_T(\rho)$. Consequently,

$$\eta_T(\rho) = J_T(\rho) + \rho^{1/2} D = \mathcal{O}\left(\sigma \sqrt{\gamma_T(\rho)} + \sqrt{\rho}\right),$$

which implies

$$R_T = \mathcal{O}\left(\left(\sigma \sqrt{\gamma_T(\rho)} + \sqrt{\rho}\right) \sqrt{T\gamma_T(\rho)}\right).$$

If $J_T(\rho)$ is obtained from Theorem 11, Howard et al. (2021, Proposition 2) implies that $J_T(\rho)$ is also $\mathcal{O}(\sigma \sqrt{\gamma_T(\rho)})$ up to logarithmic factors, from which the same regret bound (up to logarithmic factors) follows. Lastly, if $J_T(\rho)$ is obtained from Theorem 12, and $\hat{\sigma}_{u,t,\delta_1}$ is $\sigma(1 + o(1))$ with high probability, then the same regret bound holds. The $\sigma(1 + o(1))$ condition holds for the inequalities from Martinez-Taboada and Ramdas (2025), with Martinez-Taboada and Ramdas (2025, Section 4.4) establishing that $\hat{\sigma}_{u,T,\delta_1} \lesssim \sigma + c/\sqrt{T}$ for some constant c .

Appendix B. Extended related work

Self-normalized scalar processes. A prominent line of research concerns self-normalized concentration inequalities developed by (de la Peña et al., 2004, 2007, 2009b), which establish time-uniform guarantees on the behavior of self-normalized scalar processes. These results are obtained via the method of mixtures, a probabilistic technique originally introduced by Robbins (Darling and Robbins, 1967, 1968), which constructs bounds by averaging over a parameterized family of exponential supermartingales. Building on this framework, Bercu and Touati (2008) further explored the self-normalized regime, deriving concentration inequalities accommodating asymmetric and heavy-tailed increment distributions. Later on, they extended this analysis by incorporating both predictable and empirical quadratic variations (Bercu and Touati, 2019).

Sub-Gaussian self-normalized vector processes. Going from one dimension to several or infinite dimensions is far from straightforward. For this reason, most of the advances for self-normalized processes are in the sub-Gaussian case, where mixture methods provide clean concentration inequalities (Abbasi-Yadkori et al., 2011; Chowdhury and Gopalan, 2017; Flynn et al., 2023; Flynn and Reeb, 2025). We highlight the seminal work by Abbasi-Yadkori et al. (2011), which was extended to Hilbert spaces by Abbasi-Yadkori (2013) (see also Whitehouse et al. (2023)). Chowdhury and Gopalan (2017) also provided a related (though inferior) concentration inequality using a ‘double mixture’ technique.

Self-normalized vector processes beyond sub-Gaussianity. de la Peña et al. (2009a) worked out some multivariate inequalities in more general regimes. However, these are not closed form and their theoretical properties hard to study. More recently, Whitehouse et al. (2025) presented tractable self-normalized inequalities for general light-tailed noises; however, their argument relies on a covering argument that is dimension dependent and not generalizable to infinite dimensions. Similarly, Ziemann (2025) provided a self-normalized vector Bernstein inequality that is also dimension dependent and restricted to finite dimensional spaces, via PAC-Bayes arguments. Concurrent contributions explored dimension-independent self-normalized inequalities. Chugg and Ramdas (2025) developed inequalities via PAC-Bayes; while the bounds are explicit, it is not straightforward to analyze the rate because they depend on the inverse of a tail decay function and some ratio of eigenvalues. Metelli et al. (2025, Theorem 5.1) developed Bernstein-like concentration inequalities for bounded noises relying on stitching arguments. Akhavan et al. (2025) leveraged method of mixtures and truncation arguments to also develop a Bernstein-like concentration inequality. Our approach is fundamentally different to these recent works, building on the tools originally introduced in Pinelis (1994). In a different line of research, Zhou et al. (2021, Theorem 4.1) also presented a Bernstein-type self-normalized concentration inequality for vector-valued martingales, whose proof exploits solely univariate concentration inequalities (similarly to Dani et al. (2008, Theorem 5)). In contrast to our Bernstein-type inequality, this results in substantially looser constants and an extra logarithmic term in the sample size. Furthermore, our main theorem is significantly more general, leading to Bennett-type inequalities and generalizing to unbounded random variables. Follow-up work sharpened the inequalities in specific scenarios (Zhou and Gu, 2022; He et al., 2023; Zhao et al., 2023), e.g. weighted linear regression, which are out of the scope of this contribution.

Light-tailed vector-valued concentration inequalities. In the context of sums of random vectors, Pinelis (1992, 1994) introduced a martingale based approach tailored to light-tailed random vectors, which led to generalizations of well-known concentration inequalities (such as Hoeffd-

ing and Bernstein inequalities) that hold uniformly over time in smooth Banach spaces. In his framework, any dependence on dimensionality is effectively substituted by a geometric property of the underlying Banach space, i.e. its smoothness parameter (which equals one in Hilbert spaces). Recent vector-valued concentration bounds, such as sharp vector-valued empirical Bernstein inequalities (Martinez-Taboada and Ramdas, 2026, 2025) or heavy-tailed vector-valued concentration inequalities (Whitehouse et al., 2024) build on the theoretical tools introduced by Pinelis. However, these were not self-normalized, and our contribution pushes on this trajectory by generalizing the Pinelis framework to self-normalized processes.

Time-uniform Chernoff bounds. The use of nonnegative supermartingale techniques to derive concentration inequalities has gained significant traction to provide probabilistic guarantees for streams of data that are continuously monitored and adaptively stopped, with Ville’s inequality (Ville, 1939) being the theoretical pillar of this line of research. The results presented in this work fall within the broader umbrella of time-uniform concentration, aligning with the anytime-valid Chernoff-style bounds exhibited in Howard et al. (2020, 2021).

Appendix C. Further experiments

In Section 5.3, we presented an empirical comparison between our inequalities and those of Abbasi-Yadkori et al. (2011). We extend these experiments to additionally include the inequalities proposed in Akhavan et al. (2025, Theorem 3.2) and Metelli et al. (2025, Theorem 5.1). All hyperparameter choices are kept identical to those used in Section 5.3. Figure 2 reports the resulting confidence bounds corresponding to the different concentration inequalities. We observe that our inequalities, as well as those of Abbasi-Yadkori et al. (2011), yield substantially tighter bounds in practice than these more recent alternatives. This behavior is consistent with the theoretical considerations discussed in Section 4.5.

Appendix D. Limitations of our work

We discuss in this section the limitations of our contribution. For simplicity, let us focus on bounded random noises (Assumption 5) such that the conditional standard deviation σ_t is constant and equal to σ . In such a setting, our Bernstein-type inequality from Theorem 9 establishes a confidence interval with radius

$$B \log \left(\frac{2}{\delta} \right) + C_n \sqrt{2 \log \left(\frac{2}{\delta} \right)}.$$

In view of $C_n^2 \leq \sigma^2 \sum_{i \leq n} \|G_i\|^2 \leq 4\sigma^2 \gamma_n(\rho)$, the dominating term of the above expression can be upper bounded by

$$2\sigma \sqrt{2\gamma_n(\rho) \log \left(\frac{2}{\delta} \right)}.$$

Furthermore, the sub-Gaussian concentration inequality from Abbasi-Yadkori et al. (2011) yields a confidence interval with a radius that can be upper bounded by

$$\frac{B}{2} \sqrt{2 \log \left(\frac{1}{\delta} \right) + 2\gamma_n(\rho)}.$$

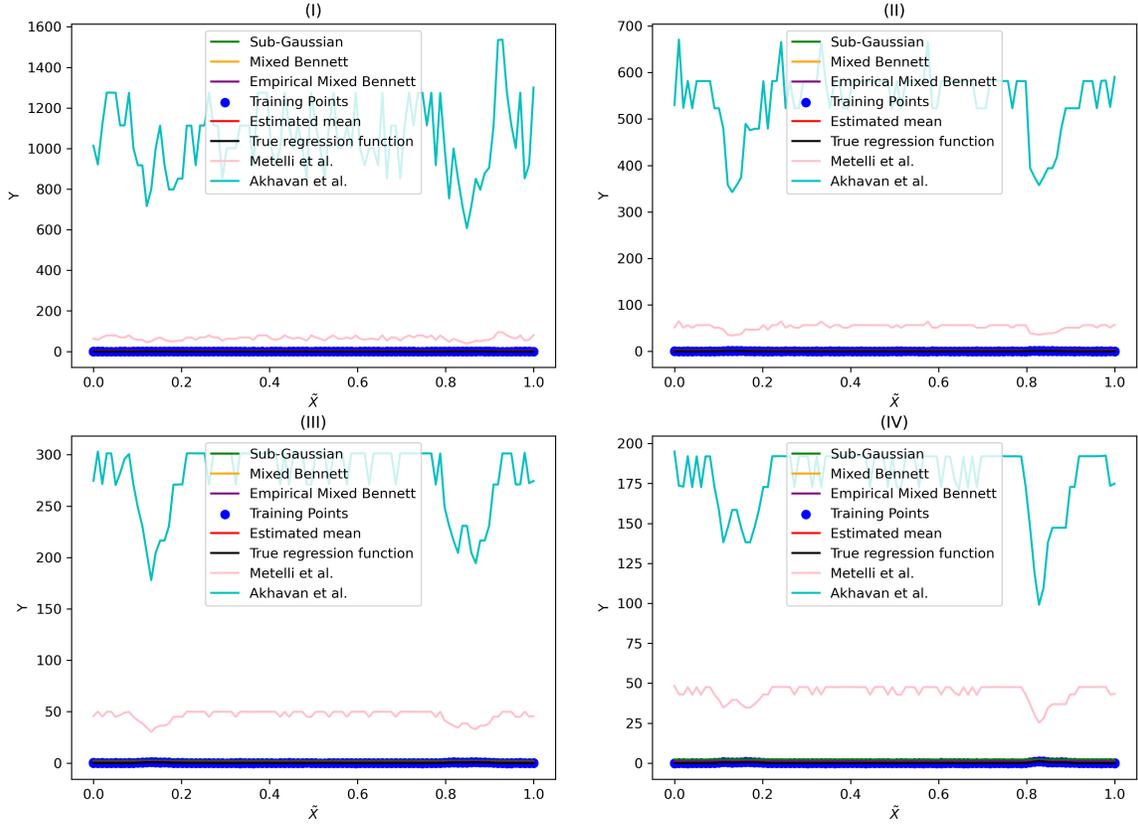


Figure 2: Illustration of the optimistic upper confidence bounds for the regression function after 500 rounds using sub-Gaussian, mixed Bennett, empirical mixed Bennett inequalities, [Akhavan et al. \(2025, Theorem 3.2\)](#), and [Metelli et al. \(2025, Theorem 5.1\)](#) for noises following (I) a rescaled uniform distribution, (II) a rescaled (5, 5)-beta distribution, (III) a rescaled (20, 20)-beta distribution, and (IV) a rescaled (50, 50)-beta distribution. Training points are drawn following a UCB procedure, with the covariates $X_t = k(\cdot, \tilde{X}_t)$ illustrated in the original space (pre-embedded in the RKHS).

Assuming that $\sigma \approx B$ for ease of comparison, note that both radii scale as $\mathcal{O}(\sqrt{\gamma_n(\rho)})$ with n (this is precisely the term that is usually considered in the regret analyses for the bandit problem).

However, our inequality was obtained after optimizing for λ for a given n , while the inequality from [Abbasi-Yadkori et al. \(2011\)](#) uses a mixture argument of the analogous hyperparameter (which is a vector in their case, and they mix following a standard multivariate Gaussian distribution). It is well understood that mixing generally leads to confidence intervals that are inflated by a logarithmic factor of the pseudo-variance process in comparison to tightly optimized inequalities, see e.g. [Howard et al. \(2021, Section 3\)](#) for a discussion. This begs the question of whether our inequalities could be improved by a logarithmic factor.

In order to address this question, let us consider the reduction of our problem to one dimension. We can think of the one-dimensional problem as the multivariate setting where all the directions

are $X_t = r_t e_1$. For simplicity, we assume that these vectors have constant unit norms ($r_t = 1$), so $X_t = e_1$ for all t . In this case,

$$\left\| (\rho I + V_t)^{-1/2} M_t \right\| = \frac{\left| \sum_{i \leq t} \epsilon_i \right|}{\sqrt{\rho + t}}, \quad \|G_t\| = \frac{1}{\sqrt{\rho + t}}.$$

Observing that $\sum_{i \leq n} \|G_i\|^2 = \sum_{i \leq n} \frac{1}{\rho + i} \asymp \log n$, we find that our self-normalized Bernstein inequality (Theorem 9) provides confidence radii scaling as $\mathcal{O}(\sqrt{\log n})$. By contrast, applying the classical univariate Bernstein inequality to $\left| \sum_{i \leq t} \epsilon_i \right|$ yields radii of order \sqrt{n} ; after division by $\sqrt{\rho + n}$, these become $\mathcal{O}(1)$. Thus, our inequalities are loose by a logarithmic factor, at least in this scenario. This extra logarithmic factor can be directly recognized in the supermartingale construction, which in this one dimensional setting reduces to

$$S_t = \cosh \left(\lambda \frac{\left| \sum_{i \leq t} \epsilon_i \right|}{\sqrt{\rho + t}} \right) \exp \left(-\sigma^2 \psi_{P,B}(\lambda) \sum_{i \leq t} \frac{1}{\rho + i} \right),$$

where $\psi_{P,B}(\lambda) = \frac{e^{\lambda B} - \lambda B - 1}{B^2}$. The term $\sum_{i \leq t} 1/(\rho + i)$ being $\mathcal{O}(\log(t))$, as opposed to $\mathcal{O}(1)$, is what causes the looseness in the final concentration inequality.

One may wonder whether this looseness stems from the overall approach (i.e., looking for a nonnegative supermartingale that is already self-normalized), or rather a weak technical analysis of it. We shall argue for the former. In order to see this, let us consider the simpler (only involving exp, not cosh) one-dimensional supermartingale construction

$$S_t^+ = \exp \left(\lambda \frac{\sum_{i \leq t} \epsilon_i}{\sqrt{\rho + t}} - \sigma^2 \psi_{P,B}(\lambda) \sum_{i \leq t} \frac{1}{\rho + i} \right).$$

Such a nonnegative supermartingale is of the form

$$S_t(\psi, V_t) = \exp \left(\lambda \frac{\sum_{i \leq t} \epsilon_i}{\sqrt{\rho + t}} - \psi(\lambda) V_t \right)$$

with $\psi(\lambda) = \psi_{P,B}(\lambda)$, and $V_t = \sigma^2 \sum_{i \leq t} \frac{1}{\rho + i}$. We would like to find ψ and V_t , such that ψ is *CGF-like*² and V_t is nonnegative and $\mathcal{O}(1)$. If there do not exist such ψ and V_t , then we can conclude that the limitations of our inequalities stem indeed from the approach itself. To see that this is the case, observe that

$$\frac{S_t(\psi, V_t)}{S_{t-1}(\psi, V_{t-1})} = \exp \left(\lambda \left(\frac{1}{\sqrt{\rho + t}} - \frac{1}{\sqrt{\rho + t + 1}} \right) \sum_{i \leq t-1} \epsilon_i + \lambda \frac{\epsilon_t}{\sqrt{\rho + t}} - \psi(\lambda) (V_t - V_{t-1}) \right).$$

Note that

$$\left(\frac{1}{\sqrt{\rho + t}} - \frac{1}{\sqrt{\rho + t + 1}} \right) \sum_{i \leq t-1} \epsilon_i$$

2. A real valued function ψ with domain $[0, \lambda_{\max})$ is called *CGF-like* if it is strictly convex and twice continuously differentiable with $\psi(0) = \psi'(0_+) = 0$ and $\sup_{\lambda \in [0, \lambda_{\max})} \psi(\lambda) = \infty$. See [Howard et al. \(2020, Section 2.1\)](#) for a detailed presentation.

can be arbitrarily small (on an event with non-zero probability), and hence in order to obtain the supertingale condition

$$\frac{\mathbb{E}_{t-1} S_t(\psi, V_t)}{S_{t-1}(\psi, V_{t-1})} \leq 1,$$

we ought to have

$$\mathbb{E}_{t-1} \exp \left(\lambda \frac{\epsilon_t}{\sqrt{\rho + t}} - \psi(\lambda)(V_t - V_{t-1}) \right) \lesssim 1.$$

If ϵ_t is sub- ψ with pseudo-variance σ^2 , i.e.

$$\mathbb{E}_{t-1} \exp (\lambda \epsilon_t - \psi(\lambda) \sigma^2) \leq 1,$$

then $\epsilon_t / \sqrt{\rho + t}$ is sub- ψ with pseudo-variance $\sigma^2 / (\rho + t)$. Consequently, $V_t - V_{t-1}$ can be taken as $\sigma^2 / (\rho + t)$. However, this implies that V_t is $\omega(1)$. Hence, the limitations of our inequalities seem to stem from the approach itself, not than a loose technical analysis. Investigating whether these inequalities can be further refined through an alternative approach remains an open direction of research.