

Eventually LIL Regret: Almost Sure $\ln \ln T$ Regret for a sub-Gaussian Mixture on Unbounded Data

Shubhada Agrawal

Indian Institute of Science, Bengaluru, India

SHUBHADA@IISC.AC.IN

Aaditya Ramdas

Carnegie Mellon University, Pittsburgh, USA

ARAMDAS@CMU.EDU

Editors: Matus Telgarsky and Jonathan Ullman

Abstract

We prove that a classic sub-Gaussian mixture proposed by Robbins in a stochastic setting actually satisfies a path-wise (deterministic) regret bound. For every path in a natural “Ville event” \mathcal{E}_α , this regret till time T is bounded by $\ln^2(1/\alpha)/V_T + \ln(1/\alpha) + \ln \ln V_T$ up to universal constants, where V_T is a nonnegative, nondecreasing, cumulative variance process. (The bound reduces to $\ln(1/\alpha) + \ln \ln V_T$ if $V_T \geq \ln(1/\alpha)$.) If the data were stochastic, then one can show that \mathcal{E}_α has probability at least $1 - \alpha$ under a wide class of distributions (eg: sub-Gaussian, symmetric, variance-bounded, etc.). In fact, we show that on the Ville event \mathcal{E}_0 of probability one, the regret on every path in \mathcal{E}_0 is eventually bounded by $\ln \ln V_T$ (up to constants). We explain how this work helps bridge the world of adversarial online learning (which usually deals with regret bounds for bounded data), with game-theoretic statistics (which can handle unbounded data, albeit using stochastic assumptions). In short, conditional regret bounds serve as a bridge between stochastic and adversarial betting.

Keywords: Regret bounds, unbounded data, Robbins mixture, iterated logarithmic regret, LIL.

1. Introduction

Since at least the work of Robbins (Robbins and Siegmund, 1968; Darling and Robbins, 1967a,b), *mixture* supermartingales have been used in conjunction with Ville’s supermartingale inequality (Ville, 1939) in order to deliver time-uniform concentration inequalities and confidence sequences. Since then, extensions of Robbins’ method-of-mixtures have delivered tight bounds in both parametric and nonparametric settings, enabling several modern applications in sequential testing, estimation and bandits (Kaufmann and Koolen, 2021; Howard et al., 2021; Waudby-Smith and Ramdas, 2024). All of these works yield coverage guarantees that are inherently stochastic, and do not appear to deliver any adversarial guarantees.

In contrast, the current paper proves a *deterministic, pathwise* wealth/regret inequality (valid for every data sequence) for a particular sub-Gaussian mixture method proposed by Robbins. What makes this regret bound interesting is that it allows the underlying sequences to have unbounded domain. Due to the unbounded domain, the worst case regret bound is obviously also unbounded. We then show that there is a special “Ville event” \mathcal{E}_α on which the regret is $\ln \ln V_T + \ln^2(1/\alpha)/V_T + \ln(1/\alpha)$ up to constants, where V_T is a nonnegative, nondecreasing, cumulative variance process. What makes \mathcal{E}_α special is that if the data is stochastic (for example, conditionally zero mean and 1-sub-Gaussian), then \mathcal{E}_α holds with probability $1 - \alpha$.

We call this type of bound a *conditional* regret bound, because it implies low regret on a *large* set of paths, rather than all paths, where *large* is quantified by being a high probability set under a suitable set of probability measures. In fact, we show that on \mathcal{E}_0 (a measure one set), every path *eventually* has $\ln \ln V_T$ regret. To summarize at a high level, we prove that, deterministically:

Most paths have $\ln \ln V_T$ regret, and *almost all* paths eventually have $\ln \ln V_T$ regret.

The above guarantee extends existing connections between wealth and regret in the online convex optimization or coin-betting literature, where path-wise regret bounds are standard (e.g., Orabona and Pál (2016); Orabona (2019); Orabona and Jun (2023); Cutkosky and Orabona (2018)), but the data, losses, or gradients in the prior work are generally assumed to be bounded. For unbounded data, it is impossible to prove low regret on all paths since a single extreme observation can blow up regret; our contribution is to prove an $O(\ln \ln V_T)$ regret on a large set of paths. We also point out that getting a $\ln V_T$ conditional regret bound is quite easy; we show how to do that (for a different mixture strategy) as a preliminary step for intuition.

Contributions. We now summarize key contributions of this work.

1. For observations X_1, X_2, \dots , and functions V_t of X_1, \dots, X_t , we show that the classic sub-Gaussian process $(\exp\{\eta \sum_{i=1}^t X_i - \eta^2 V_t/2\})$ for $\eta \in \mathbb{R}$, when mixed over the parameters η weighted using Robbins’ prior, satisfies a deterministic regret bound when compared with $\max_{\eta} \exp\{\eta \sum_i X_i - \eta^2 V_t/2\}$, where the \max_{η} is unrestricted. This bound holds for all data sequences and does not require any assumptions (Eq. (7), Theorem 6).
2. As a consequence, we show that on all paths in an event \mathcal{E}_{α} where the mixture process remains uniformly bounded by $\log(1/\alpha)$ for some $\alpha \in (0, 1)$, the regret till time T is $O(\ln^2(1/\alpha)/V_T + \log(1/\alpha) + \ln \ln V_T)$, for every T (Eq. (8), Theorem 6). We call such bounds that hold on a subset (for example, \mathcal{E}_{α}), rather than on all sequences, *conditional regret bounds*.
3. If in Point 2 above, the data are assumed to be drawn from a distribution (not necessarily iid) such that the mixture processes is a non-negative supermartingale (Definition 2), the “Ville event” \mathcal{E}_{α} has measure at least $1 - \alpha$. Thus, the pathwise regret bound from Point 2 holds for a *large* fraction of paths (Theorem 6). Importantly, the stochasticity is only used to quantify the size of \mathcal{E}_{α} — the regret bounds still hold deterministically on \mathcal{E}_{α} . Our deterministic bounds thus imply the high probability regret bound: $P(R_T = O(\ln \ln V_T + \ln^2(1/\alpha)/V_T + \log(1/\alpha))) \geq 1 - \alpha$ simultaneously for a large class of measures P that ensure W_t is a sub-Gaussian process.
4. Without any assumptions on the data, we show that the conditional regret of $O(\ln \ln V_T)$ holds eventually (as opposed to for all T) on a larger set \mathcal{E}_0 , where the mixture process remains finite, but not necessarily uniformly bounded (Eq. (12), Theorem 8). Interestingly, under the stochasticity assumptions from Point 3, when the mixture process is a non-negative supermartingale, \mathcal{E}_0 has measure 1. While the previous bounds are akin to high probability regret bounds, this type of bound appears to be new in the literature: $P(R_T = O(\ln \ln V_T) \text{ eventually}) = 1$. Moreover, in this case, the mixture process acts as a witness for the (self-normalized) strong law of large numbers (SLLN) as well as an upper law of iterated logarithm (LIL), meaning that either self-normalized SLLN (and upper LIL) holds, or the mixture process gets unbounded (Theorem 8).

Paper outline In Section 2.1, we formalize the setup and introduce the notation. Background material is reviewed in Section 2.2, and related work is discussed in Section 2.3. As a warm-up, Section 3 analyzes logarithmic pathwise conditional regret bounds. Our main results are presented in Section 4, where Theorem 6 establishes an unconditional regret bound valid for all data sequences. While the worst-case regret can grow linearly in T (or V_T), we show that on a certain good event, the

regret is only iterated-logarithmic eventually. Moreover, this good event is a measure 1 set under appropriate stochasticity assumptions. Finally, Section 6 concludes with broader perspectives.

2. Preliminaries: Background and Literature

In this section, we formally introduce the setup, provide the necessary background, and discuss the relevant literature.

2.1. Setup

Let X_1, X_2, \dots be a stream of observations taking values in \mathbb{R} , possibly unbounded. For each $t \in \mathbb{N}$, define the partial sum corresponding to the initial t samples as $S_t := \sum_{i=1}^t X_i$, and let V_t be a function of X_1, \dots, X_t that is nonnegative and increasing (usually strictly) in t . Further, $S_0 = V_0 = 0$ and let $S_t = 0$ whenever $V_t = 0$, and we interpret $0/0 = 0$. For any fixed $\eta \in \mathbb{R}$, let

$$f_t(\eta) := \eta S_t - \eta^2 V_t / 2. \quad (1)$$

Observe that $f_t(\cdot)$ is a concave function. Let

$$\eta_t^* := \begin{cases} S_t/V_t & \text{if } V_t > 0 \\ 0, & \text{otherwise,} \end{cases} \quad \text{and} \quad L_t^* := f_t(\eta_t^*) = \begin{cases} S_t^2/(2V_t), & \text{if } V_t > 0 \\ 0, & \text{otherwise,} \end{cases}$$

denote the unconstrained maximizer for $f_t(\cdot)$, and the corresponding maximum value, respectively.

Next, for any prior π on η , define the mixture (wealth) process

$$W_t := \int_{\mathbb{R}} e^{f_t(\eta)} \pi(\eta) d\eta, \quad (2)$$

with $W_0 = 1$; and $W_t := 1$ for t such that $V_t = 0$. The regret of W_t with respect to the maximum attainable wealth is defined as

$$R_t := L_t^* - \ln W_t.$$

Note that $R_t = 0$ whenever $V_t = 0$, so it is without loss of generality that the reader may consider $V_t > 0$ by default going forward. Finally, for $\alpha \in (0, 1)$, consider the set

$$\mathcal{E}_\alpha = \left\{ \sup_{t \geq 1} \ln W_t \leq \ln \frac{1}{\alpha} \right\}, \quad (3)$$

on which the log-wealth W_t remains bounded. We call \mathcal{E}_α the ‘‘Ville event’’.

Remark 1 *We do not make any assumptions about the data stream. In particular, it is an arbitrary sequence of real numbers. Hence, all the quantities defined above are deterministic. However, whenever the sequence X_i is assumed to be drawn from some probability distribution, all capitalized variables defined above should be interpreted as random variables.*

The main focus of this work is to establish regret bounds for the mixture wealth process W_t , constructed using a mixing distribution π originally introduced in [Robbins and Siegmund \(1968\)](#). We show that this classical stochastic construction can be reinterpreted within an adversarial framework: specifically, W_t admits deterministic, path-wise regret guarantees, which have an interesting nontrivial

interpretation. The worst-case regret is unbounded (linear in t or V_t); but we can prove that one attains low regret on every path in \mathcal{E}_α . What makes \mathcal{E}_α special is that under natural stochastic assumptions (implied by the underlying mixture processes), the event holds with high probability. So we obtain a deterministic low regret bound on every path in a high probability event. In fact, we even show that we achieve a small regret *eventually* on almost every path (i.e. on every path in measure 1 set \mathcal{E}_0). In other words, the set of paths on which regret is unbounded is a measure zero set (under many natural families of distributions specified later).

Our work thus connects worst-case regret analysis with stochastic betting approaches developed in game-theoretic statistics. To place our results in context, we first review the necessary background that will be used throughout. We also refer the reader to Appendix D for a detailed discussion of the betting game considered in this work, including the performance criteria for the betting strategies and the associated notion of regret.

2.2. Background

We begin with notions that characterize the behavior of stochastic processes, such as sub-Gaussian tails, the strong law of large numbers (SLLN), and the law of the iterated logarithm (LIL). We then review the framework of betting and wealth processes, which will serve as useful tools in relating the adversarial and stochastic betting via our conditional regret bounds. After some formal setup, we present a key definition of *sub-Gaussian processes* from de la Pena (2004); Howard et al. (2020).

When dealing with stochastic data, we will work on a filtered measurable space (Ω, \mathcal{F}) , where $\mathcal{F} = (\mathcal{F}_t)_{t \geq 0}$ is a filtration (an increasing sequence of sigma algebras). We will also consider a family Π of probability measures defined on (Ω, \mathcal{F}) .

For this paper, the underlying Ω is typically \mathbb{R}^∞ , where we observe a single real-valued observation X_t at each time t so that $\mathcal{F}_t = \sigma(X_1, \dots, X_t)$, with \mathcal{F}_0 being trivial. A stochastic process M is a sequence of random variables adapted to \mathcal{F} , meaning that M_t is \mathcal{F}_t -measurable. For $P \in \Pi$, we say that M is P -integrable if $\mathbb{E}_P[|M_t|] < \infty$ for all t . A P -supermartingale is a P -integrable stochastic process M such that $\mathbb{E}_P[M_t | \mathcal{F}_{t-1}] \leq M_{t-1}$, P -almost surely. Given $\mathcal{P} \subset \Pi$, we say that M is a \mathcal{P} -supermartingale if M is a P -supermartingale for every $P \in \mathcal{P}$. If P or \mathcal{P} are understood from context or not important, we just call M a supermartingale.

Definition 2 (Sub-Gaussian process) *A stochastic process $(S_t)_{t \geq 0}$ is said to be a sub-Gaussian process with a nonnegative and nondecreasing variance proxy $(V_t)_{t \geq 0}$ if $(\exp(\eta S_t - \eta^2 V_t / 2))_{t \geq 0}$ is a (nonnegative) super-martingale for every $\eta \in \mathbb{R}$. Typically, we take $S_0 = V_0 = 0$ by default.*

S_t is often simply a sum process $\sum_{i=1}^t X_i$. We recall in Appendix A that the above condition is satisfied (for different V_t) when \mathcal{P} contains distributions that are (a) centered sub-Gaussian, or (b) symmetric around 0, or (c) have mean zero and finite variance. The condition thus encompasses both light-tailed (case (a)) and heavy-tailed (cases (b) and (c)) settings. Typically, V_t will be strictly positive and strictly increasing. However, we do not assume it because in case (b) above, we have $V_t = \sum_{i \leq t} X_i^2$, and if the distribution of X_1 (for example) is symmetric with a point mass at 0 (eg: it equals 0 with probability half, and a standard Gaussian with probability half), then it is possible for $X_1 = 0$ to occur with positive probability, and then $S_1 = V_1 = 0$. Outside of these sorts of edge cases, V_t will indeed be strictly increasing and positive for $t \geq 1$. Also, typically V_t will increase to infinity as $t \rightarrow \infty$, but we will only assume that when needed to derive limiting statements.

Strong law of large numbers (SLLN) and iterated logarithm (LIL). SLLN and LIL describe the asymptotics of stochastic processes. For example, when X_i are iid from some distribution with mean 0, the SLLN guarantees that $S_t/t \rightarrow 0$ almost surely as $t \rightarrow \infty$, while if they also have finite variance σ^2 , then the LIL refines this as below:

$$\limsup_{t \rightarrow \infty} \frac{|S_t|}{\sigma \sqrt{2t \ln \ln t}} = 1, \quad \text{almost surely.}$$

These results also extend to self-normalized processes, which are especially useful when variances are unknown or data are heavy-tailed (de la Pena et al., 2009); one can derive a self-normalized SLLN and LIL involving S_t/V_t and $S_t/\sqrt{V_t \log \log V_t}$. Our versions of self-normalized SLLN and LIL for sub-Gaussian processes are in Section 4.

Martingales, betting, and wealth processes. There is an equivalence between non-negative (super)martingales and bets in a (conservative) fair game. A wealth process $(W_t)_{t \geq 0}$ is a nonnegative adapted process that satisfies $W_0 = 1$. Such processes can be interpreted as the capital of a gambler who starts with a unit initial wealth ($W_0 = 1$), and at each time t bets a fraction of the current wealth W_{t-1} on the next outcome. The gambler’s wealth at time t is W_t . If the bets are made in a fair (conservative) game, then the wealth process $(W_t)_{t \geq 1}$ is a non-negative (super)martingale with respect to the natural data filtration. Ville’s inequality then yields the following time-uniform guarantee: $P[\sup_t W_t \geq \frac{1}{\alpha}] \leq \alpha$.

Conversely, given a non-negative martingale $(M_t)_{t \geq 1}$ with $M_0 = 1$, it can be viewed as a wealth process generated by a sequence of bets. At time t , the per-round bet is an e-variable (Ramdas and Wang, 2025, §1) that is given by $\frac{M_t}{M_{t-1}} \mathbf{1}(M_{t-1} > 0)$, with $0/0 = 0$. The martingale property ensures that the game is fair. This perspective underlies developments in game-theoretic probability (Shafer and Vovk, 2005, 2019) and game-theoretic statistics (Shafer, 2021; Ramdas et al., 2023; Ramdas and Wang, 2025), where nonnegative martingales are also referred to as wealth processes, and stochastic coin-betting algorithms (Orabona and Pál, 2016). We refer the reader to Appendix D for a detailed discussion on the betting game that we consider in this work.

2.3. Literature Review

In the online learning and information theory literature, Cover (1991); Cover and Ordentlich (2002) prove an $O(\log t)$ deterministic regret bound against all sequences of observations from a bounded domain, say $[0, 1]$, for their mixture strategy. Several later works in online learning literature propose different mixture wealth processes that improve the constant terms in the path-wise regret bound of $O(\log t)$, that holds for any sequence of bounded data.

Cutkosky and Sarlos (2019) use wealth/betting arguments to derive a path-wise regret bound in an adversarial online learning setting. However, they require that the gradients are bounded in a certain dual norm. Koolen and Van Erven (2015) prove a deterministic, pathwise regret bound that is $O(\sqrt{V_t \ln \ln V_t})$ for a certain variance process $V_t \leq t$. However, this result crucially relies on bounded losses.

Orabona and Jun (2023) leverage the wealth-regret connection, but their primary aim is to construct time-uniform confidence sets for the mean of bounded data using the wealth process of certain betting strategies. In their work, they implicitly prove iterated logarithmic conditional regret bounds for bounded data for a wealth process different from W_t . By contrast, we extend the

analysis to unbounded data and also prove regret bounds for W_t that hold for arbitrary sequences of observations.

Agrawal et al. (2021, Lemma F.1) prove a deterministic regret bound of $O(\ln t)$ for unbounded data, that holds for an arbitrary sequence of observations. However, they define regret with respect to a constrained optimizer. In our work, we instead consider regret with respect to the global optimizer and prove a deterministic regret bound for unbounded data, which we show is eventually $O(\ln \ln t)$ on a measure 1 set for stochastic data.

In the coding literature, Rissanen (2002) showed that if the parameter space of an exponential family is restricted to a compact subset of the interior of the parameter space with non-empty interior, then the minimax regret is $O(\log t)$. We refer the reader to Grünwald (2007) for more details. We do not make any such assumptions, and instead, prove a conditional regret bound of $O(\ln \ln t)$, that eventually holds on a measure 1 set (instead of against all sequences).

There are a few works in the literature that derive inequalities for martingales (concentration/bound on moments, etc.) from certain deterministic counterparts. For example, Acciaio et al. (2013) derive Doob’s L^p maximal inequalities (bounds on p^{th} moments of a martingale) as a consequence of elementary deterministic inequalities that hold for all sequences. Beiglböck and Nutz (2014) show that a general martingale inequality can be reduced to certain deterministic inequalities in a small number of variables. Beiglböck and Siorpaes (2015) provide a proof of the Burkholder-Davis-Gundy inequalities as a consequence of elementary deterministic inequalities. Gushchin (2014) derives Doob’s maximal inequality on the probability of a martingale exceeding a level, from deterministic counterparts. These works prove probabilistic martingale inequalities via deterministic statements. Despite also connecting pathwise statements and martingale inequalities, our contributions are orthogonal: we prove deterministic regret bounds for a particular mixture method on unbounded data.

Clerico et al. (2025) develop a method to construct time-uniform confidence sets for parameters of a generalized linear model, via reduction to the online learning framework. They prove certain path-wise data-dependent regret bounds that hold with probability 1, and use them to construct confidence sequences. We establish deterministic regret bounds for unbounded data, and show that under stochastic assumptions, these bounds improve to $O(\ln \ln t)$ on a set of measure 1. We are not aware of any work that proves such rates for unbounded observations, on a set of measure 1.

Finally, Rakhlin and Sridharan (2017) highlight the tight connections between martingale concentration inequalities and regret analysis. They show an equivalence between tail inequalities on the supremum of a collection of martingales and certain path-wise regret bounds. However, our contributions appear orthogonal to theirs; there is no mention in their work of Ville’s inequality, mixture wealth processes, etc. Nevertheless, our work can be considered within the same theme of connecting martingales and regret.

Shafer and Vovk (Shafer and Vovk, 2005, Chapter 5) prove a game-theoretic version of the LIL that holds path-wise and applies to arbitrary forecasting protocols. However, their framework implicitly restricts the sequence of outcomes to be bounded, with possibly time-dependent bounds (see, Shafer and Vovk (2005, Theorems 5.1, 5.2)). Our results are complementary: we handle unbounded data via self-normalization, obtain explicit regret bounds valid for all t , and construct the supermartingale that certifies our version of the SLLN and LIL.

All the above-cited works span different sub-fields. A complete picture that relates these is not entirely clear to us, and forming such a complete picture is a direction for future work.

3. Warm-up: Logarithmic Conditional Regret

To get some intuition, we begin with a simple setting when the prior π on η is a Gaussian distribution with mean 0 and variance σ_0^2 , for some $\sigma_0 > 0$. Then, the mixture (wealth) process from (2) becomes

$$W_t = \frac{1}{\sqrt{1 + \sigma_0^2 V_t}} e^{\frac{\sigma_0^2 S_t^2}{2(1 + \sigma_0^2 V_t)}}, \quad (4)$$

and as earlier, $W_t = 1$ for t such that $V_t = 0$. Theorem 3 below shows that the mixture W_t incurs a regret of $O(\ln V_t + \ln(1/\alpha))$ on the set \mathcal{E}_α , defined in (3).

Theorem 3 For all $t \geq 0$,

$$R_t = \frac{1}{2} \ln(1 + \sigma_0^2 V_t) + \frac{S_t^2}{2V_t(1 + \sigma_0^2 V_t)}.$$

(Recall $0/0=0$.) Further, for any $v_0 > 0$, the following holds on \mathcal{E}_α :

$$\forall t \geq 1, \quad \text{either } V_t \leq v_0 \quad \text{or} \quad R_t \leq \left(\frac{1}{2} + \frac{1}{2\sigma_0^2 v_0}\right) \ln(1 + \sigma_0^2 V_t) + \frac{\ln(1/\alpha)}{\sigma_0^2 v_0}. \quad (5)$$

Moreover, if the data are drawn from any distribution P such that $(S_t)_{t \geq 1}$ is a sub-Gaussian process with variance proxy $(V_t)_{t \geq 1}$, then $P[\mathcal{E}_\alpha] \geq 1 - \alpha$.

The proof of the above theorem proceeds by first establishing an expression for the unconditional regret, which corresponds to the first expression in the theorem statement. Next, on the set \mathcal{E}_α , $\ln W_t$ is bounded. Using this bound in the expression for R_t after appropriate rearrangements, we get a conditional bound on R_t on the set \mathcal{E}_α , in terms of V_t and S_t . The expression in (5) then follows by bounding certain occurrences of V_t by v_0 . The proof is straightforward, and we present it in Appendix B for completeness.

Corollary 4 For the special case when X_i are drawn iid from a standard Gaussian distribution, S_t is a sub-Gaussian process with variance proxy $V_t = t$ (see Appendix A). Assuming $\sigma_0 = 1$ for simplicity, in this case, $R_t = \frac{1}{2} \ln(1 + t) + \frac{S_t^2}{2t(1+t)}$ path-wise. Moreover, on \mathcal{E}_α , which occurs with probability at least $1 - \alpha$, we have $R_t \leq \ln(1 + t) + \ln(1/\alpha)$, for all $t \geq 1$.

In fact, it follows from the following theorem that $O(\ln V_t)$ regret for the Gaussian-mixture wealth holds eventually on every path on a set of measure 1.

Theorem 5 Let $\sigma_0 = 1$ for simplicity. Consider the event

$$\mathcal{E}_0 := \left\{ V_t \uparrow \infty; \limsup_t W_t < \infty \right\}.$$

We have that on \mathcal{E}_0 ,

$$R_t \leq \ln(1 + V_t) + o(1) \text{ eventually,}$$

i.e. on every path in \mathcal{E}_0 , and for every constant d , there exists an n such that for all $t \geq n$, $R_t \leq \ln(1 + V_t) + d$. In particular, dividing by V_t and taking the limit as $t \rightarrow \infty$, we get that $S_t/V_t \rightarrow 0$ path-wise on \mathcal{E}_0 . Finally, if the data are drawn from a distribution P such that $(S_t)_{t \geq 1}$ is sub-Gaussian with variance proxy $(V_t)_{t \geq 1}$, then W_t is a nonnegative super-martingale, immediately implying that \mathcal{E}_0 is a measure 1 event, thus implying the (self-normalized) SLLN. Contrapositively, if there is a path on which $S_t/V_t \rightarrow 0$ does not hold, then $\limsup_t W_t = \infty$ on that path.

Ville's martingale theorem (Ville, 1939) states that fixing a measure P , for every event A such that $P(A) = 0$, there exists a nonnegative P -martingale M^P that increases to ∞ on A . But this M^P is potentially different for every P . Here, we have a single super-martingale $Z = (W_t)_{t \geq 1}$ that serves as a *witness* for the SLLN for an entire class of distributions \mathcal{P} (eg, zero mean, 1-sub-Gaussian distributions). In general, martingales and even super-martingales may not suffice as witnesses for the strong law for composite \mathcal{P} , or more generally for a composite generalization of Ville's theorem; for example, Ruf et al. (2023) showed that in general one must employ e-processes instead (see Ruf et al. (2023) for a definition). A general characterization of when super-martingales suffice is open; the above result provides an interesting example.

Moreover, Theorem 5 provides an explicit witness super-martingale of the violation of the strong law, meaning that if the strong law is violated on some path, $\limsup_t W_t = \infty$ on that path.

Proof Since

$$W_t := \int_{\mathbb{R}} e^{f_t(\eta)} \pi(\eta) d\eta = \frac{e^{\frac{S_t^2}{2(1+V_t)}}}{\sqrt{1+V_t}},$$

we see that \mathcal{E}_0 implies that $S_t^2/(1+V_t) < \sqrt{V_t}$ eventually (where the right hand side can be picked anything asymptotically larger than $\ln(1+V_t)$). Since

$$R_t = L_t^* - \ln W_t = \frac{1}{2} \ln(1+V_t) + \frac{S_t^2}{2V_t(1+V_t)},$$

we then see that \mathcal{E}_0 implies $R_t \leq \ln(1+V_t) + o(1)$. ■

4. Main Result: Iterated Logarithmic Conditional Regret

Recall $S_t, V_t, f_t, \eta_t^*, L_t^*$, and \mathcal{E}_α from Section 2.1. In this section, we establish a regret bound for Robbins' mixture strategy (Robbins and Siegmund, 1968), conditioned on the set \mathcal{E}_α from (3) on which the log of the mixture wealth remains bounded. To achieve this, as in Section 3, we first prove a path-wise regret bound that holds for all data sequences X_1, X_2, \dots . In general, this bound can be arbitrarily large (at least linear in t). However, we show that on the set \mathcal{E}_α , where the mixture process W_t is restricted to be uniformly bounded, the regret is also small. In fact, for the mixture wealth using Robbins' heavy-near-zero prior, we show that the bound on regret till time t , for every t , is $O(\ln \ln V_t)$.

To formalize, for f_t defined in (1), in this section, we consider the mixture wealth process W_t given by

$$W_t = \frac{1}{Z_0} \int_{-1}^1 \frac{e^{f_t(\eta)}}{|\eta| (\ln \frac{c}{|\eta|}) (\ln \ln \frac{c}{|\eta|})^2} d\eta, \quad (6)$$

which is obtained by using the following mixture π in (2): for $c \geq 6.6e$ and $\eta \in [-1, 1]$,

$$\pi(\eta) = \frac{1}{Z_0} \frac{\mathbf{1}_{|\eta| \leq 1}}{|\eta| (\ln \frac{c}{|\eta|}) (\ln \ln \frac{c}{|\eta|})^2}.$$

It equals 0 outside $[-1, 1]$. Here, Z_0 is the normalization constant

$$Z_0 = \int_{-1}^1 \frac{d\eta}{|\eta| (\ln \frac{c}{|\eta|}) (\ln \ln \frac{c}{|\eta|})^2} = \frac{2}{\ln \ln c}.$$

Theorem 6 For all $t \geq 1$ and $\rho \in (0, \frac{1}{4})$,

$$R_t \leq \begin{cases} \frac{1}{2} + \ln \frac{2}{\ln \ln c} + \ln \ln(c\sqrt{1+V_t}) + 2 \ln \ln \ln(c\sqrt{1+V_t}), & \text{if } \frac{|S_t|}{V_t} \leq \frac{1}{\sqrt{1+V_t}} \\ \frac{1}{2} + \ln \frac{2}{\ln \ln c} + \ln \left(\frac{|S_t|}{\sqrt{V_t}} \sqrt{1 + \frac{1}{V_t}} \right) \\ \quad + \ln \ln(c\sqrt{1+V_t}) + 2 \ln \ln \ln(c\sqrt{1+V_t}), & \text{if } \frac{1}{\sqrt{1+V_t}} < \frac{|S_t|}{V_t} \leq 1 \\ \frac{V_t}{2} \left(\frac{|S_t|}{V_t} - 1 + \rho \right)^2 - \ln \left(\frac{\rho \ln \ln c}{\ln(c/1-\rho)(\ln \ln(c/1-\rho))^2} \right), & \text{if } 1 < \frac{|S_t|}{V_t}. \end{cases} \quad (7)$$

Moreover, on \mathcal{E}_α , there exist a universal constant $C > 0$ such that

$$\forall t \geq 1, \quad R_t \leq C \left(1 + \frac{1 + \ln^2(1/\alpha)}{V_t} + \ln(1/\alpha) + \ln \ln \left(c\sqrt{1+V_t} \right) \right). \quad (8)$$

Furthermore, if the data are drawn from any distribution P such that $(S_t)_{t \geq 1}$ is a sub-Gaussian process with variance proxy $(V_t)_{t \geq 1}$, then $P[\mathcal{E}_\alpha] \geq 1 - \alpha$.

Observe that the bound in (7) holds along *all* data sequences. The three cases in this bound capture the behavior of the unrestricted optimizer for the log-wealth. This bound, especially when $|S_t|/V_t > 1$ (third branch), can be arbitrarily large. To get some intuition, let us consider the special case of $S_t = \sum_i X_i$ and $V_t = t$. The third branch then corresponds to $|S_t|/t > 1$. If the data are not stochastic or are stochastic but have a mean greater than 1, then $|S_t|/t$ will likely be greater than 1. Hence, the third branch in (7) is likely to occur, and the bound on the regret will be linear in t (at least). On the other hand, if the data are iid from a 1-sub-Gaussian distribution (for example, $N(0, 1)$), then both SLLN and LIL hold. In particular, $S_t/t \rightarrow 0$ almost surely, and $S_t/\sqrt{t} \approx \sqrt{\ln \ln t}$ infinitely often. These facts suggest that the third branch of the regret bound can occur only finitely many times, and only leads to at most a constant regret (possibly a function of α). In this setting, since the first branch eventually applies, it yields $O(\ln \ln t)$ regret. Moreover, by LIL, the second branch occurs infinitely often; however, the third term within this branch is $O(\ln \ln \ln t)$ (by LIL), which again leads to regret of $O(\ln \ln t)$ in this case as well.

Note that in the above discussion, we assumed the data to be stochastic only to develop some intuition about the bound in (7). However, it does not require any stochasticity assumption, and instead holds path-wise. Next, in (8) we show that on the set where the data is well-behaved in the sense that the wealth process W_t is uniformly bounded from above, the path wise regret is small, and is in fact, $O(\ln \ln V_t)$ at time t . Even the path-wise bound in (8) does not require any stochasticity assumption. Stochasticity is only used in the last part of the theorem, which shows that if the data is stochastic and satisfies certain conditions, the set \mathcal{E}_α where (8) holds is indeed a large set (of probability at least $1 - \alpha$).

Finally, observe from (8) that if $V_t \geq \log(1/\alpha)$, then on \mathcal{E}_α , $R_t = O(\log(1/\alpha) + \ln \ln(1 + V_t))$.

Confidence sequences (CS) from conditional regret bound. Next, under the stochasticity assumption, we observe that our regret bound in (7) immediately implies a $(1 - \alpha)$ -CS (a sequence of confidence intervals with a time-uniform coverage guarantee of $1 - \alpha$) for the mean of the unknown data-generating distribution. To see this, let B_t denote the right hand side of (7), and recall that $R_t = S_t^2/(2V_t) - \log W_t$. Then, from (7), the following holds path wise on \mathcal{E}_α :

$$\forall t \geq 1, \quad \frac{S_t^2}{2V_t} - \log \frac{1}{\alpha} \leq R_t \leq B_t.$$

If the data is drawn from a distribution P such that $(S_t)_{t \geq 1}$ is a sub-Gaussian process with variance proxy $(V_t)_{t \geq 1}$, then the above further implies

$$P \left[\forall t \geq 1, \frac{S_t^2}{2V_t} \leq B_t + \ln \frac{1}{\alpha} \right] \geq 1 - \alpha. \quad (9)$$

Finally, in the special case that the data are conditionally σ_t -sub-Gaussian with mean μ_t (for some predictable μ_t, σ_t), we substitute $S_t = \sum_{i \leq t} (X_i - \mu_i)$ and $V_t = \sum_{i \leq t} \sigma_i^2$ to get that

$$P \left[\forall t \geq 1, \bar{\mu}_t \in \hat{\mu}_t \pm \left\{ \sqrt{\frac{2V_t(B_t + \ln \frac{1}{\alpha})}{t^2}} \right\} \right] \geq 1 - \alpha, \quad (10)$$

where $\bar{\mu}_t = \sum_{i \leq t} \mu_i/t$ and $\hat{\mu}_t = \sum_{i \leq t} X_i/t$. To see that the shape of this CS is optimal, we refer the reader back to Eq. (8). In fact, later in Theorem 8, we show that $B_t = \ln \ln V_t(1 + o(1))$ eventually on \mathcal{E}_α . Thus, the CS in (9), obtained using regret bound in (7), recovers the exact LIL-CS.

Conditional regret bounds from confidence sequences (CS). We now explain how, in martingale-based CS constructions, one can also read off a conditional regret bound similar to those in (8). To this end, consider again the special case where the data is drawn from a distribution P such that $(S_t)_{t \geq 1}$ is a sub-Gaussian process with variance proxy $(V_t)_{t \geq 1}$. Moreover, for some predictable μ_t , let $S_t = \sum_{i \leq t} (X_i - \mu_i)$, and suppose we have the following CS with $\bar{\mu}_t = \sum_{i \leq t} \mu_i/t$, $\hat{\mu}_t = \sum_{i \leq t} X_i/t$, a positive constant c_1 , and $\alpha \in (0, 1)$:

$$P \left[\forall t \geq 1, |\bar{\mu}_t - \hat{\mu}_t| \leq \sqrt{\frac{2V_t(c_1 \ln \ln(e + V_t) + \ln \frac{1}{\alpha})}{t^2}} \right] \geq 1 - \alpha. \quad (11)$$

One frequently used approach for arriving at inequalities of the form (11) is to threshold a non-negative (mixture) super-martingale (or equivalently, a wealth process), say W_t , at a level $\frac{1}{\alpha}$, followed by an application of Ville's inequality (Darling and Robbins, 1967a; Robbins and Siegmund, 1970).

As earlier, let $\mathcal{E}_\alpha := \{\sup_t \ln W_t \leq \ln(1/\alpha)\}$. While directly thresholding a wealth process W_t (as in \mathcal{E}_α) would give an implicit CS for mean, to get an explicit CS of the form in (11), these martingale-based approaches typically proceed by deriving a lower bound for $\ln W_t$ that is relatively tight on the set \mathcal{E}_α , and thresholding this lower bound process, instead. In particular, these methods can be seen to proceed by implicitly establishing an inequality like the following: on \mathcal{E}_α ,

$$\ln W_t \geq \frac{S_t^2}{2V_t} - c_1 \ln \ln(e + V_t), \quad \forall t \geq 1.$$

implying

$$\left\{ \sup_t \ln W_t \leq \ln \frac{1}{\alpha} \right\} \subseteq \left\{ \sup_t \left(\frac{S_t^2}{2V_t} - c_1 \ln \ln(e + V_t) \right) \leq \ln \frac{1}{\alpha} \right\},$$

hence (11). But, the above inequality also implies the following on \mathcal{E}_α :

$$R_t = \frac{S_t^2}{2V_t} - \ln W_t \leq c_1 \ln \ln(e + V_t), \quad \forall t \geq 1,$$

yielding a conditional regret bound of $O(\ln \ln(e + V_t))$ on \mathcal{E}_α , in the same spirit as (8) with $C = c_1$. We remark that our direct approach in Theorem 6 yields tighter path-wise regret bounds that hold unconditionally (Eq. (7)), and lead to exact upper LIL, as shown later in Theorem 8.

Proof sketch. The proof of Theorem 6 proceeds in three steps. In Step 1 (Section 4.1.1), we prove the path-wise bound on R_t in (7). In Step 2 of the proof (Section 4.1.2), we show that on the set \mathcal{E}_α where the mixture-wealth process is uniformly bounded across t , the bounds in (7) are $O(\ln(1/\alpha) + \frac{1}{\sqrt{V_t}} + \frac{1}{\sqrt{V_t}} \ln^2(1/\alpha) + \ln \ln V_t)$, for every t . Finally, Step 3 (Section 4.1.3) shows that when the data is stochastic with appropriate assumptions, the conditional regret bound from Step 2 holds for a large number of sequences. In particular, the set \mathcal{E}_α occurs with probability at least $1 - \alpha$. Note that only the third step of the proof requires the stochasticity assumption on data. We present a complete proof of Theorem 6 in Section 4.1.

Remark 7 *Observe that the Ville's event \mathcal{E}_α does not prevent the maximum hindsight wealth $\frac{S_t^2}{2V_t}$ from being unbounded. In fact, it can drift away from the log of the mixture wealth at the rate of $O(\ln \ln V_t)$. Since the mixture wealth is also uniformly bounded on the Ville's event, this means that the maximum wealth can increase to ∞ at $O(\ln \ln V_t)$ rate.*

Theorem 6 above proves a path-wise regret bound (Eq. (7)), as well as a conditional regret bound of $O(\ln \ln V_t)$ (Eq. (8)) that holds for every t on a set \mathcal{E}_α . If the data is stochastic, then \mathcal{E}_α is large. In the following theorem, we show that the $O(\ln \ln V_t)$ conditional regret bound holds eventually on a larger set \mathcal{E}_0 where the process W_t remains finite, but possibly not uniformly bounded. This further implies that on \mathcal{E}_0 , $|S_t|/V_t \rightarrow 0$ and $\limsup_t |S_t|/\sqrt{V_t \ln \ln V_t} \leq \sqrt{2}$. In other words, either both of these aforementioned convergences hold or the process W_t explodes. In particular, if the data are stochastic with appropriate assumptions, the set \mathcal{E}_0 is a set of measure 1, and we conclude that the process W_t acts as a witness for the self-normalized SLLN and (upper) LIL.

Theorem 8 *Consider the event*

$$\mathcal{E}_0 := \left\{ V_t \uparrow \infty; \limsup_t W_t < \infty \right\}.$$

We have that on \mathcal{E}_0 ,

$$R_t \leq \ln \ln V_t (1 + o(1)) \text{ eventually,} \quad (12)$$

i.e. on every path in \mathcal{E}_0 , and for every $\epsilon > 0$, there exists n such that for all $t \geq n$, $R_t \leq (1 + \epsilon) \ln \ln V_t$. In particular, dividing by V_t and taking the limit as $t \rightarrow \infty$, we get that $|S_t|/V_t \rightarrow 0$ path-wise on \mathcal{E}_0 . Next, dividing by $\ln \ln V_t$ and taking limit as $t \rightarrow \infty$, we get

$$\limsup_t \frac{|S_t|}{\sqrt{2V_t \ln \ln V_t}} \leq 1$$

path-wise on \mathcal{E}_0 . As a corollary, if the data are drawn from a distribution such that $(S_t)_{t \geq 1}$ is a sub-Gaussian process with variance $(V_t)_{t \geq 1}$, then W_t is a nonnegative super-martingale, immediately implying that \mathcal{E}_0 is a measure 1 event, thus implying the self-normalized SLLN and LIL.

As in Theorem 5, Theorem 8 demonstrates a deterministic statement of the form: either $R_t = O(\ln \ln V_t)$, or the wealth process W_t explodes. In case of stochastic data, it provides an explicit super-martingale witnessing the violation of the SLLN and LIL: if either of these is violated on some path, $\limsup_t W_t = \infty$ on that path.

Proof Since $R_t = \frac{S_t^2}{2V_t} - \ln W_t$, from (7) we get a lower bound on W_t . Further, using $\ln(x) \leq x - 1$ and completing the squares, we get the following path-wise lower bound on W_t :

$$\ln W_t \geq \begin{cases} \frac{S_t^2}{2V_t} - \frac{1}{2} - \ln \frac{2}{\ln \ln c} - \ln \ln(c\sqrt{1+V_t}) - 2 \ln \ln \ln(c\sqrt{1+V_t}), & \text{if } \frac{|S_t|}{V_t} \leq \frac{1}{\sqrt{1+V_t}} \\ \frac{1}{2} \left(\frac{|S_t|}{\sqrt{V_t}} - \sqrt{1 + \frac{1}{V_t}} \right)^2 - \ln \frac{2}{\ln \ln c} - \frac{1}{2V_t} \\ \quad - \ln \ln(c\sqrt{1+V_t}) - 2 \ln \ln \ln(c\sqrt{1+V_t}), & \text{if } \frac{1}{\sqrt{1+V_t}} < \frac{|S_t|}{V_t} \leq 1 \\ \frac{S_t^2}{2V_t} - \frac{V_t}{2} \left(\frac{|S_t|}{V_t} - 1 + \rho \right)^2 + \ln \left(\frac{\rho \ln \ln c}{\ln(c/1-\rho)(\ln \ln(c/1-\rho))^2} \right), & \text{if } 1 < \frac{|S_t|}{V_t}. \end{cases} \quad (13)$$

We now argue that on \mathcal{E}_0 , the third branch occurs only finitely many times. On the third branch, since $|S_t| > V_t$, we can further lower bound as

$$\ln W_t \geq V_t \left(\frac{1 - \rho^2}{2} \right) + \ln \left(\frac{\rho \ln \ln c}{\left(\ln \frac{c}{1-\rho} \right) \left(\ln \ln \frac{c}{1-\rho} \right)^2} \right).$$

Recall that on \mathcal{E}_0 , $V_t \uparrow \infty$ and $\limsup_t W_t < \infty$. From this and the inequality above it, we see that on \mathcal{E}_0 , if $\limsup_t \frac{|S_t|}{V_t} > 1$ (i.e., the third branch occurs infinitely often), then $\limsup_t \ln W_t = \infty$, which leads to a contradiction.

Thus, we only focus on the first two cases (i.e., t such that $\frac{|S_t|}{V_t} \leq 1$), henceforth. Now, on \mathcal{E}_0 , dividing both sides of (13) by V_t and taking limit as $t \rightarrow \infty$ we get $|S_t|/V_t \rightarrow 0$. Similarly, dividing by $\ln \ln V_t$ instead, and taking the limit as $t \rightarrow \infty$, we get

$$\limsup_t \frac{|S_t|}{\sqrt{2V_t \ln \ln V_t}} \leq 1. \quad (14)$$

Finally, from (13), we also have the following eventually on \mathcal{E}_0 :

$$R_t \leq \begin{cases} \frac{1}{2} + \ln \frac{2}{\ln \ln c} + \ln \ln(c\sqrt{1+V_t}) + 2 \ln \ln \ln(c\sqrt{1+V_t}), & \text{if } \frac{|S_t|}{V_t} \leq \frac{1}{\sqrt{1+V_t}} \\ \frac{|S_t|}{\sqrt{V_t}} \sqrt{1 + \frac{1}{V_t}} - \frac{1}{2} + \ln \frac{2}{\ln \ln c} \\ \quad + \ln \ln(c\sqrt{1+V_t}) + 2 \ln \ln \ln(c\sqrt{1+V_t}), & \text{if } \frac{1}{\sqrt{1+V_t}} < \frac{|S_t|}{V_t} \leq 1. \end{cases}$$

Using (14) in the above bound, on \mathcal{E}_0 , $R_t \leq \ln \ln V_t(1 + o(1))$, eventually. \blacksquare

Shafer and Vovk, in their book (Shafer and Vovk, 2005, Chapter 5), present a game-theoretic version of the LIL that is closely related to our Theorem 8. Their proof sets up a game between three players: the forecaster, the skeptic, and reality, and shows the existence of a betting strategy (a non-negative supermartingale) that makes the skeptic infinitely rich if an event of the form \mathcal{E}_0 fails. Thus, whenever the skeptic's wealth remains finite, the LIL holds. Shafer and Vovk (2005, Theorem 5.2) prove an exact LIL statement, but under the assumption that the observations are bounded. They also prove just the upper LIL in Shafer and Vovk (2005, Theorem 5.1) under slightly weaker assumptions. However, their bounds only hold if the data are in a bounded range, and they do not derive explicit regret bounds. Our results extend in complementary directions: we obtain explicit path-wise regret bounds that hold for all t for unbounded data. Further, we establish an upper LIL for

self-normalized processes. Our analysis exhibits an explicit supermartingale (wealth process) that witnesses our version of the LIL. We refer the reader to Appendix D for a detailed description of the betting game involving the three players.

Remark 9 *Despite the apparent focus of our bounds above on sub-Gaussian processes (Definition 2), note that they also apply when the data is drawn from a mean-zero distribution with only a finite second moment. Such distributions can still be treated as sub-Gaussian, albeit with a different variance proxy (see Point 4 in the examples in Appendix A). Moreover, recall that the finiteness of the second moment is not only sufficient but also necessary for LIL to hold.*

4.1. Proof of Theorem 6

In this section, we give a proof of Theorem 6. It relies on certain auxiliary lemmas and propositions. While we state these results here, their proofs are given in Appendix C.

Recall from Section 2.1 that $R_t := f_t(\eta_t^*) - \ln W_t$, and that for t such that $V_t = 0$, we also have $S_t = 0$, and hence for such t , for all $\eta \in \mathbb{R}$, $f_t(\eta) = 0$, and $\ln W_t = 0$. Thus, the inequalities in the theorem hold trivially in this case. Henceforth, we focus on t so that $V_t > 0$.

4.1.1. STEP 1: PROVING REGRET BOUND ALONG ALL DATA SEQUENCES

The first two cases in (7) correspond to t such that the unconstrained optimizer η_t^* is in $[-1, 1]$. To get the bound on R_t in (7) in these cases, we show that the mixture wealth W_t , which is an integration of a non-negative function over $[-1, 1]$ weighted by a prior π , is at least the value of the integration over a subset that contains the maximizer η_t^* . This subset is chosen large enough so that it contains a significant fraction of the prior mass. Further, it is also small enough so that the value of the function being integrated remains almost constant (and almost equal to the maximum value) in this range.

Lemma 10 below chooses this subset to be an interval of length r , centered at the optimizer η_t^* .

Lemma 10 *For every $t \geq 1$ such that $|\eta_t^*| \leq 1$, for every $r \in (0, 1)$,*

$$R_t \leq \frac{V_t r^2}{2} - \ln \left(\int_{\substack{\{\eta: |\eta - \eta_t^*| \leq r\} \\ \cap [-1, 1]}} \pi(\eta) d\eta \right).$$

For $r \in (0, 1)$, for all feasible η in the chosen interval $\{\eta : |\eta - \eta_t^*| \leq r\} \cap [-1, 1]$, the function $e^{f_t(\cdot)}$ can be expressed as the maximum value adjusted by a small cost. Lemma 10 bounds this cost to arrive at a bound on R_t in terms of V_t , r , and the mass that the prior π puts on the chosen interval.

Next, Lemma 11 below, lower bounds the mass that the prior π places on the chosen interval. It uses the shape of π , and properties like convexity and radial monotonicity of π , to arrive at the bounds.

Lemma 11 *For $|\eta_t^*| < 1$ and any $r \in (0, 1)$:*

$$\int_{\{\eta: |\eta - \eta_t^*| \leq r\} \cap [-1, 1]} \pi(\eta) d\eta \geq \begin{cases} \frac{\ln \ln c}{2 \ln(c/r) (\ln \ln(c/r))^2}, & \text{if } |\eta_t^*| \leq r, \\ \frac{r \ln \ln c}{2 |\eta_t^*| (\ln(c/r)) (\ln \ln(c/r))^2}, & \text{if } |\eta_t^*| > r. \end{cases}$$

The bound in (7) in the first two cases follows by plugging that from Lemma 11 into Lemma 10, setting $r = \frac{1}{\sqrt{1+V_t}}$, and using $\frac{V_t r^2}{2} \leq \frac{1}{2}$.

Coming to the third branch in (7), we again express the function $e^{f_t(\cdot)}$ in terms of $e^{f_t(\eta_t^*)}$ with the adjustment terms. However, since the optimizer η_t^* lies outside $[-1, 1]$ in this case, we choose the radius r large enough, so that the interval of length r centered at η_t^* still has enough intersection with $[-1, 1]$ (recall, π is only supported on $[-1, 1]$). Since the feasible η in this intersection are far from η_t^* , the cost incurred in relating $e^{f_t(\cdot)}$ to $e^{f_t(\eta_t^*)}$ is relatively high in this case compared to the previous two cases. Lemma 12 below formalizes this and proves the inequality in the third case in (7). First term in the bound corresponds to the cost of relating $e^{f_t(\cdot)}$ with $e^{f_t(\eta_t^*)}$ for $\eta \in [-1, 1] \cap \{\eta : |\eta - \eta_t^*| \leq r\}$. The second term is the mass that the prior π puts on this set.

Lemma 12 *For $t \geq 1$ such that $|\eta_t^*| > 1$, for $\rho \in (0, 1)$,*

$$R_t \leq \frac{V_t}{2} \left(\frac{|S_t|}{V_t} - 1 + \rho \right)^2 - \ln \left(\frac{\rho \ln \ln c}{\left(\ln \frac{c}{1-\rho} \right) \left(\ln \ln \frac{c}{1-\rho} \right)^2} \right).$$

It now remains to show that the bound in (7) is not too large on \mathcal{E}_α . Moreover, \mathcal{E}_α holds along most of the data sequences when the data satisfies appropriate stochasticity assumptions. We show these in the following section.

4.1.2. STEP 2: CONDITIONAL REGRET BOUND ON \mathcal{E}_α

Recall $\mathcal{E}_\alpha = \{\sup_t \ln W_t \leq \ln(1/\alpha)\}$. Observe that to show (8), it suffices to show that the bounds in the middle and last branches of (7) are of the form in (8) on \mathcal{E}_α (it holds for the first branch with $C = 4$).

We first show (8) for the case where $\frac{|S_t|}{V_t} > 1$. As an intermediate step, Proposition 13 below first shows that on \mathcal{E}_α , if V_t is larger than a threshold that depends only on α , say V_α , the third branch cannot occur. Using this observation, it further bounds R_t on the third branch in terms of this threshold V_α , V_t , and a parameter ρ .

Proposition 13 *Let $\rho \in (0, 1/4)$. The following hold on \mathcal{E}_α .*

1. *There exists a threshold*

$$V_\alpha := \frac{2}{1-\rho^2} \left(\ln \frac{1}{\alpha} - \ln \frac{\rho \ln \ln c}{2} + \ln \ln \frac{c}{1-\rho} + 2 \ln \ln \ln \frac{c}{1-\rho} \right) \quad (15)$$

such that for all t such that $V_t \geq V_\alpha$, we have $|\eta_t^| \leq 1$.*

2. *For all $t \geq 1$ such that $|\eta_t^*| \geq 1$ (or equivalently, $\frac{|S_t|}{V_t} \geq 1$),*

$$R_t \leq \frac{(1+\rho)^2 V_\alpha^2}{8V_t} + \frac{\rho V_\alpha}{2} (1+2\rho) - \ln \left(\frac{\rho \ln \ln c}{\left(\ln \frac{c}{1-\rho} \right) \left(\ln \ln \frac{c}{1-\rho} \right)^2} \right). \quad (16)$$

With the above proposition, we get that (8) holds in this case, with

$$C = \frac{3}{(1-\rho)^2} \left(1 \vee \left\{ \ln(\rho \ln \ln c) - \ln \ln \frac{c}{1-\rho} - 2 \ln \ln \ln \frac{c}{1-\rho} - \ln 2 \right\} \right)^2.$$

Now, it only remains to prove (8) for $t \geq 1$ such that $\frac{1}{\sqrt{1+V_t}} < \frac{|S_t|}{V_t} \leq 1$. In this case, since $\ln(x) \leq \frac{x^2}{4}$ for $x \geq 0$, we have from (7) that

$$\begin{aligned} R_t &= \frac{S_t^2}{2V_t} - \ln W_t \\ &\leq \left(\frac{1}{2} + \ln \frac{2}{\ln \ln c} \right) + \ln \ln \left(c\sqrt{1+V_t} \right) + \frac{S_t^2}{4V_t} \left(1 + \frac{1}{V_t} \right) + 2 \ln \ln \ln \left(c\sqrt{1+V_t} \right). \end{aligned} \quad (17)$$

Rearranging the above inequality, using that $|S_t| \leq V_t$, and that on \mathcal{E}_α , $\sup_t \ln W_t \leq \ln \frac{1}{\alpha}$, we get the following on \mathcal{E}_α in this case:

$$\frac{S_t^2}{4V_t} \leq \ln \frac{1}{\alpha} + \left(\frac{1}{2} + \ln \frac{2}{\ln \ln c} \right) + \ln \ln \left(c\sqrt{1+V_t} \right) + \frac{1}{4} + 2 \ln \ln \ln \left(c\sqrt{1+V_t} \right).$$

Substituting the above back in (17), we get that on \mathcal{E}_α , the following holds for all $t \geq 1$ such that $\frac{1}{\sqrt{1+V_t}} < \frac{|S_t|}{V_t} \leq 1$:

$$R_t \leq \ln \frac{1}{\alpha} + 2 \left(\frac{1}{2} + \ln \frac{2}{\ln \ln c} \right) + 2 \ln \ln \left(c\sqrt{1+V_t} \right) + \frac{1}{2} + 4 \ln \ln \ln \left(c\sqrt{1+V_t} \right),$$

proving (8) in this case, with $C = 7$.

4.1.3. STEP 3: STOCHASTIC DATA

Finally, if the data are drawn from a distribution P such that $(S_t)_{t \geq 1}$ is a sub-Gaussian process with variance proxy $(V_t)_{t \geq 1}$, then $(e^{f_t(\eta)})_{t \geq 1}$ is a non-negative super-martingale for every $\eta \in \mathbb{R}$ (Definition 2). Thus, $(W_t)_{t \geq 1}$ is also a non-negative super-martingale. The last statement then follows from Ville's inequality (Ville, 1939). \square

5. Trade-off between regret and growth rate

In this section, we point out an apparently inherent trade-off between the growth rate of mixture martingales and their rate of accumulating regret: asymptotic optimality in one precludes asymptotic optimality in the other.

When the data is stochastic (i.i.d. from a fixed distribution Q), the lower limit $\liminf_{t \rightarrow \infty} \frac{1}{V_t} \ln W_t$ is typically Q -a.s. a constant, and in that case, that constant is termed as the asymptotic growth rate of W_t against Q (Ramdas and Wang, 2025, Chapter 7). It quantifies the rate at which the wealth grows exponentially under Q . In this section, we compare the two wealth processes defined in (4) and (6) at the level of their growth rate and confidence sequence (CS) widths. In contrast to the rest of the paper, this section is purely stochastic. We assume that the data is drawn i.i.d. from a fixed distribution Q that satisfies $\frac{|S_t|}{V_t} \rightarrow m$ as $t \rightarrow \infty$, for some $m \in \mathbb{R}_+$.

For the Gaussian-mixture from Section 3, using the regret bound from (5), we have

$$\liminf_{t \rightarrow \infty} \frac{\ln W_t}{V_t} = \liminf_{t \rightarrow \infty} \left(\frac{S_t^2}{2V_t^2} - \frac{R_t}{V_t} \right) \geq \liminf_{t \rightarrow \infty} \frac{S_t^2}{2V_t^2} = \frac{m^2}{2}.$$

However, as we saw in Section 3, this wealth process has a regret that is $O(\ln V_t)$, and does not achieve $O(\ln \ln V_t)$ regret.

Let us now consider the mixture wealth process defined in (6) using Robbins' prior. From Theorem 6, we see that

$$\limsup_{t \rightarrow \infty} \frac{R_t}{V_t} \leq \begin{cases} \limsup_{t \rightarrow \infty} \frac{1}{2} \left(\frac{|S_t|}{V_t} - 1 + \rho \right)^2, & \text{if } 1 < \frac{|S_t|}{V_t} \\ 0, & \text{otherwise.} \end{cases}$$

Since Robbins' mixture puts a large prior mass at bets close to 0, it has relatively small mass at larger bets. Therefore, when the data is such that the optimal bet $\frac{|S_t|}{V_t}$ is far from 0 (in particular > 1), it lacks sufficient weight to compete with the best-in-hindsight, and hence, suffers large regret in those cases. While the bound above gives a linear *upper bound* on the regret in these cases, we numerically observe that the *regret is indeed linear* in such cases. This linear regret, in fact, affects the growth rate of the wealth process, when m is far from 0. To see this, consider $m > 1$. In this case, the growth rate for the Robbins' mixture wealth is smaller than $m^2/2$ due to the linear regret suffered in this case. It is $m^2/2$ when $m \in [0, 1]$.

Next, it is well-known that for a fixed error probability $\alpha \in (0, 1)$, the CS derived from the Gaussian mixture wealth (4) are tighter early on (for small t) than those obtained using Robbins' mixture. However, from (10), we can see that they are asymptotically looser with an asymptotic width of $O(\ln V_t)$, while the latter has $O(\ln \ln V_t)$ width. The relative tightness of the former (for small t) can be explained from the faster growth rate of the corresponding wealth process, as discussed next.

Recall that one way to construct a $(1 - \alpha)$ -CS for mean is to run, in parallel, a collection of one-sided sequential tests of the null hypothesis that the mean is m_0 , for each $m_0 \in \mathbb{R}$. At time t , the $(1 - \alpha)$ -CS is the set of all values m_0 whose corresponding test has not yet stopped (and hence, has not rejected m_0). For simplicity, consider testing the mean of a Gaussian distribution with known unit variance. To test mean being m_0 , we may use either of the mixture wealth processes in (4) or (6), with $S_t = \sum_{i \leq t} (X_i - m_0)$, and $V_t = t$. Both wealth processes are non-negative martingales under the null, and therefore yield valid one-sided sequential tests with the stopping time given by

$$\tau_{m_0} := \min \left\{ t : W_t \geq \frac{1}{\alpha} \right\}.$$

From above, one can deduce that for fixed m_0 , the stopping time (to reject m_0) will be smaller if the wealth W_t grows faster. Since the growth rate for Robbins' mixture wealth from (6) is smaller for m_0 that are "far" from the true mean, the corresponding test rejects m_0 later than that using Gaussian mixture from (4). However, once the "far-off" m_0 are excluded, the Robbins' wealth suffers a smaller regret, and hence has a faster growth rate (in non-asymptotic sense) than the Gaussian mixture. Thus, it rejects the nearby m_0 faster, leading to narrower CS eventually. In fact, we don't see the CS derived using Robbins' mixture wealth process perform poorly in practice. This is because the mean values very far from the true mean are eliminated from both the CSs early on (since we often pick $\alpha = 0.05$), since the signal from them is high, but the effect would be more visible for small α .

Thus, there appears to be a trade-off between the asymptotic growth rate of the mixture wealth process and the regret incurred. To achieve the smaller $O(\ln \ln V_t)$ regret when the mean shift is small, the mixture in (6) pays in its asymptotic growth rate against larger mean shifts. To the best of our knowledge, these observations have not been pointed out in the literature.

6. Conclusions

In this work, we proved deterministic regret bounds for sub-Gaussian mixture processes of the form W_t , which mix $\exp\{\eta S_t - \eta^2 V_t/2\}$ over $\eta \in \mathbb{R}$. Our regret bounds hold for all sequences of data. However, without any assumption on the data sequence, they can be large (for example, linear in t or V_t). We then show that for certain sequences of data, specifically along which the mixture W_t remains uniformly bounded, the regret for the mixture using Robbins' heavy-near-0 prior is $O(\ln \ln V_t)$ for every t . Furthermore, if the data are assumed to be stochastic, then the set of such sequences is large, and hence $O(\ln \ln V_t)$ holds along many sequences of observations. In fact, we show that regret is $O(\ln \ln V_t)$ eventually on a set of measure 1.

We strongly believe that the deterministic regret bounds of the form we prove in this work for sub-Gaussian wealth processes from the previous paragraph, also hold for a broader class of (sub- ψ) wealth processes obtained by mixing $\exp\{\eta S_t - \psi(\eta)V_t\}$ over η for a non-negative, convex function ψ . Moreover, such mixtures are (super)martingales under a much broader (sub- ψ) class of probability measures (see, [Howard et al. \(2020\)](#) for a definition). When the data is assumed to be stochastic from such a distribution, we expect to get similar conditional regret bounds on Ville's events. Our work, thus, gives a path towards writing down (conditional) regret bounds for unbounded data, which may be shown to be small for stochastic settings. We, in fact, conjecture that:

Behind all concentration results that are derived by constructing an appropriate non-negative (super)martingale, followed by an application of Ville's inequality, there exists a conditional regret bound, that is, a bound on the regret of the chosen martingale that is small on the Ville event.

Finally, in game-theoretic statistics, non-negative (super)martingales, or equivalently, wealth processes in a (conservative) fair game, play a central role, together with Ville's inequality. At their core, these are probabilistic objects. In parallel, the field of adversarial online learning is built around deterministic, path-wise regret inequalities. There appear to be strong connections between these two worlds. For example, Ville, in his thesis ([Ville, 1939](#)), proved a striking deterministic statement about martingales: given a probability measure P , for every event A of measure zero, there exists a nonnegative P -martingale M^P that becomes unbounded on A . Ville's martingale theorem thus serves as one bridge between the probabilistic and deterministic viewpoints. Related connections have been explored in subsequent works, ([Rakhlin and Sridharan, 2017](#); [Acciaio et al., 2013](#); [Beiglböck and Nutz, 2014](#); [Beiglböck and Siorpaes, 2015](#); [Gushchin, 2014](#); [Shafer and Vovk, 2005](#); [Clerico et al., 2025](#)), which derive probabilistic results from certain deterministic inequalities. However, a complete and unified understanding of the relationship between the two perspectives is still missing in the literature. We view our contributions as another step toward unifying these two paradigms.

Acknowledgments

SA acknowledges the generous support from the Pratiksha Trust, Bangalore, through the Young Investigator Award, and the DST INSPIRE Faculty Grant IFA24-ENG-389. AR acknowledges funding from NSF grant 2310718 and a Sloan fellowship.

References

- B Acciaio, M Beiglböck, F Penkner, W Schachermayer, and J Temme. A trajectorial interpretation of Doob’s martingale inequalities. *The Annals of Applied Probability*, pages 1494–1505, 2013.
- Shubhada Agrawal, Wouter M Koolen, and Sandeep Juneja. Optimal best-arm identification methods for tail-risk measures. *Advances in Neural Information Processing Systems*, 34:25578–25590, 2021.
- Mathias Beiglböck and Marcel Nutz. Martingale inequalities and deterministic counterparts. *Electronic Journal of Probability*, 19:1–15, 2014.
- Mathias Beiglböck and Pietro Siorpaes. Pathwise versions of the Burkholder-Davis-Gundy inequality. *Bernoulli*, pages 360–373, 2015.
- Eugenio Clerico, Hamish Flynn, and Gergely Neu. Confidence sequences for generalized linear models via regret analysis. *arXiv preprint arXiv:2504.16555*, 2025.
- Thomas M Cover. Universal portfolios. *Mathematical finance*, 1(1):1–29, 1991.
- Thomas M Cover and Erik Ordentlich. Universal portfolios with side information. *IEEE Transactions on Information Theory*, 42(2):348–363, 2002.
- Ashok Cutkosky and Francesco Orabona. Black-box reductions for parameter-free online learning in Banach spaces. In *Conference On Learning Theory*, pages 1493–1529. PMLR, 2018.
- Ashok Cutkosky and Tamas Sarlos. Matrix-free preconditioning in online learning. In *International Conference on Machine Learning*, pages 1455–1464. PMLR, 2019.
- Donald A Darling and Herbert Robbins. Confidence sequences for mean, variance, and median. *Proceedings of the National Academy of Sciences*, 58(1):66–68, 1967a.
- Donald A Darling and Herbert Robbins. Iterated logarithm inequalities. *Proceedings of the National Academy of Sciences*, 57(5):1188–1192, 1967b.
- Victor H de la Pena. Self-normalized processes: Exponential inequalities, moment bounds and iterated logarithm laws. *The Annals of Probability*, 32(3A):1902–1933, 2004.
- Victor H de la Pena, Tze Leung Lai, and Qi-Man Shao. *Self-normalized processes: Limit theory and Statistical Applications*. Springer, 2009.
- Peter D Grünwald. *The minimum description length principle*. MIT press, 2007.
- AA Gushchin. On pathwise counterparts of Doob’s maximal inequalities. *Proceedings of the Steklov Institute of Mathematics*, 287(1):118–121, 2014.
- Steven R Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Time-uniform Chernoff bounds via nonnegative supermartingales. *Probability Surveys*, 2020.
- Steven R Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Time-uniform, nonparametric, nonasymptotic confidence sequences. *The Annals of Statistics*, 49(2):1055–1080, 2021.

- Emilie Kaufmann and Wouter M. Koolen. Mixture martingales revisited with applications to sequential tests and confidence intervals. *Journal of Machine Learning Research*, 22(246):1–44, 2021.
- Wouter M Koolen and Tim Van Erven. Second-order quantile methods for experts and combinatorial games. In *Conference on Learning Theory*, pages 1155–1175. PMLR, 2015.
- Francesco Orabona. A modern introduction to online learning. *arXiv preprint arXiv:1912.13213*, 2019.
- Francesco Orabona and Kwang-Sung Jun. Tight concentrations and confidence sequences from the regret of universal portfolio. *IEEE Transactions on Information Theory*, 2023.
- Francesco Orabona and Dávid Pál. Coin betting and parameter-free online learning. *Advances in Neural Information Processing Systems*, 29, 2016.
- Alexander Rakhlin and Karthik Sridharan. On equivalence of martingale tail bounds and deterministic regret inequalities. In *Conference on Learning Theory*, pages 1704–1722. PMLR, 2017.
- Aaditya Ramdas and Ruodu Wang. *Hypothesis testing with e-values*. Foundations and Trends in Statistics, 2025.
- Aaditya Ramdas, Peter Grünwald, Vladimir Vovk, and Glenn Shafer. Game-theoretic statistics and safe anytime-valid inference. *Statistical Science*, 38(4):576–601, 2023.
- Jorma J Rissanen. Fisher information and stochastic complexity. *IEEE Transactions on Information Theory*, 42(1):40–47, 2002.
- Herbert Robbins and David Siegmund. Iterated logarithm inequalities and related statistical procedures. *Mathematics of the Decision Sciences*, 2:267–279, 1968.
- Herbert Robbins and David Siegmund. Boundary crossing probabilities for the wiener process and sample sums. *The Annals of Mathematical Statistics*, pages 1410–1429, 1970.
- Johannes Ruf, Martin Larsson, Wouter M Koolen, and Aaditya Ramdas. A composite generalization of ville’s martingale theorem using e-processes. *Electronic Journal of Probability*, 28:1–21, 2023.
- Glenn Shafer. Testing by betting: A strategy for statistical and scientific communication. *Journal of the Royal Statistical Society Series A: Statistics in Society*, 184(2):407–431, 2021.
- Glenn Shafer and Vladimir Vovk. *Probability and finance: it’s only a game!*, volume 491. John Wiley & Sons, 2005.
- Glenn Shafer and Vladimir Vovk. *Game-theoretic foundations for probability and finance*. John Wiley & Sons, 2019.
- Jean Ville. *Etude critique de la notion de collectif*. Gauthier-Villars Paris, 1939.
- Ian Waudby-Smith and Aaditya Ramdas. Estimating means of bounded random variables by betting. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 86(1):1–27, 2024.

Appendix A. Sub-Gaussian Process with Variance Proxy

In this section, we recall some examples of sub-Gaussian processes $(S_t)_t$ with the corresponding variance proxy $(V_t)_t$ to demonstrate that Definition 2 is quite general.

1. First, if X_i are drawn iid from the standard Gaussian distribution (with mean 0 and variance 1), then $(S_t)_{t \geq 1}$ is a 1-sub-Gaussian process with variance proxy $V_t = t$. This is easy to see since for all $\eta \in \mathbb{R}$, $\exp \eta S_t - \eta^2 t / 2$ is a non-negative martingale, satisfying the conditions in Definition 2.
2. For $\sigma > 0$, if X_i are drawn iid from a centered (mean zero) σ^2 -sub-Gaussian distribution P (that satisfies $\mathbb{E}_P[\exp(\eta X - \sigma^2 \eta^2 / 2)] \leq 1$), then $(S_t)_{t \geq 1}$ is sub-Gaussian with $V_t = \sigma^2 t$.
3. If X_t are drawn iid from a distribution P that is symmetric around 0, then $(S_t)_{t \geq 1}$ is sub-Gaussian with $V_t = \sum_{i=1}^t X_i^2$. Notably, in this case, P may not even have a finite first moment (for instance, Cauchy distribution).
4. If X_t are drawn iid from a distribution P with mean 0 and finite variance (i.e., $\mathbb{E}_P[X^2] < \infty$), then $(S_t)_{t \geq 1}$ is sub-Gaussian with $V_t = \frac{1}{3} \sum_{i=1}^t (X_i^2 + 2\mathbb{E}[X_i^2])$.

The examples above show that the condition encompasses both light-tailed (cases 1 and 2) and heavy-tailed (cases 3 and 4) settings. Moreover, the condition of X_i being iid can be greatly relaxed. We refer the reader to Howard et al. (2021, Table S3) and the discussions therein for more details.

Appendix B. Proofs from Section 3

Proof [Proof of Theorem 3] We first prove the expression for R_t . To this end, observe from (4), that

$$-\ln W_t = \frac{1}{2} \ln(1 + \sigma_0^2 V_t) - \frac{\sigma_0^2 S_t^2}{2(1 + \sigma_0^2 V_t)}, \quad (18)$$

which equals 0 when $V_t = 0$. Add $L_t^* = S_t^2 / (2V_t)$ (or 0 when $V_t = 0$) to obtain

$$R_t = \frac{1}{2} \ln(1 + \sigma_0^2 V_t) + \frac{S_t^2}{2V_t} - \frac{\sigma_0^2 S_t^2}{2(1 + \sigma_0^2 V_t)} = \frac{1}{2} \ln(1 + \sigma_0^2 V_t) + \frac{S_t^2}{2V_t(1 + \sigma_0^2 V_t)},$$

proving the expression for R_t from the theorem statement.

Next, rearrange (18) to get

$$\frac{S_t^2}{1 + \sigma_0^2 V_t} = \frac{2}{\sigma_0^2} (\ln W_t + \frac{1}{2} \ln(1 + \sigma_0^2 V_t)).$$

Plugging this into the expression for R_t from above, we get

$$R_t = \frac{1}{2} \ln(1 + \sigma_0^2 V_t) + \frac{1}{\sigma_0^2 V_t} (\ln W_t + \frac{1}{2} \ln(1 + \sigma_0^2 V_t)).$$

On \mathcal{E}_α , since $\ln W_t \leq \ln(1/\alpha)$ for all t , we further get

$$R_t \leq \frac{1}{2} \ln(1 + \sigma_0^2 V_t) + \frac{1}{\sigma_0^2 V_t} \left(\ln \frac{1}{\alpha} + \frac{1}{2} \ln(1 + \sigma_0^2 V_t) \right).$$

When $V_t \geq v_0 > 0$, replacing V_t by v_0 in the denominators of the above expression, we get the desired bound.

Finally, if X_i are iid from P , and $(S_t)_{t \geq 1}$ is sub-Gaussian with variance $(V_t)_{t \geq 1}$, then for $\eta \in \mathbb{R}$, the process $(e^{f_t(\eta)})_{t \geq 1}$ is a non-negative super-martingale (Definition 2). Hence, $(W_t)_{t \geq 1}$ is also a non-negative super-martingale. Ville's inequality (Ville, 1939) then implies that $P[\mathcal{E}_\alpha] \geq 1 - \alpha$. ■

Appendix C. Proofs from Section 4

Proof [Proof of Lemma 10] Recall that $\eta_t^* = \frac{S_t}{V_t}$ and $L_t^* = f_t(\eta_t^*) = \frac{S_t^2}{2V_t}$. On algebraic manipulation, we get $f_t(\eta) = L_t^* - \frac{V_t}{2}(\eta - \eta_t^*)^2$, which implies

$$W_t = \int_{-1}^1 e^{f_t(\eta)} \pi(\eta) d\eta = e^{L_t^*} \int_{-1}^1 e^{-\frac{V_t}{2}(\eta - \eta_t^*)^2} \pi(\eta) d\eta \geq e^{L_t^* - \frac{1}{2}V_t r^2} \int_{\substack{\{\eta: |\eta - \eta_t^*| \leq r\} \\ \cap [-1, 1]}} \pi(\eta) d\eta.$$

Upon taking logs and rearranging, we get the desired inequality. ■

Proof [Proof of Lemma 11] For $y \in (0, 1]$ and $c > 6.6e$ chosen as in $\pi(\cdot)$, let

$$\phi(y) := \frac{1}{y \left(\ln \frac{c}{y} \right) \left(\ln \ln \frac{c}{y} \right)^2}.$$

Then, $\pi(\eta) = \frac{\phi(|\eta|)}{Z_0}$, where recall that Z_0 is the normalization constant. Let $I_t := [\eta_t^* - r, \eta_t^* + r]$, and $|I_t \cap [-1, 1]|$ denote the length of the interval $I_t \cap [-1, 1]$.

For the first case ($|\eta_t^*| \leq r < 1$), if $\eta_t^* > 0$, then $[\eta_t^* - r, \eta_t^*] \subset I_t \cap [-1, 1]$, and otherwise, $[\eta_t^*, \eta_t^* + r] \subset I_t \cap [-1, 1]$. Let $s_t := \text{Sign}(\eta_t^*)$, and let $\bar{I}_t := [s_t(|\eta_t^*| - r), s_t|\eta_t^*|]$ denote the interval $[\eta_t^* - r, \eta_t^*]$ when $s_t = 1$, and $[\eta_t^*, \eta_t^* + r]$, otherwise. Then, the following inequalities hold:

$$\begin{aligned} \frac{1}{r} \int_{I_t \cap [-1, 1]} \frac{\phi(|\eta|)}{Z_0} d\eta &\geq \frac{1}{r} \int_{\bar{I}_t} \frac{\phi(|\eta|)}{Z_0} d\eta && (\bar{I}_t \subset I_t \cap [-1, 1]) \\ &\geq \frac{1}{Z_0} \phi \left(\frac{1}{r} \int_{\bar{I}_t} |\eta| d\eta \right) && (\text{Convexity of } \phi(\cdot) \text{ on both sides of } 0) \\ &\stackrel{(a)}{\geq} \frac{1}{Z_0} \phi(r), \end{aligned}$$

where (a) follows by observing that $\phi(\cdot)$ is radially decreasing as well as for all $\eta \in \bar{I}_t$, $|\eta| \leq r$. On rearranging the above inequality, we get the desired inequality in this case.

For the second case ($|\eta_t^*| > r$), consider the following inequalities:

$$\frac{1}{|I_t \cap [-1, 1]|} \int_{I_t \cap [-1, 1]} \phi(|\eta|) d\eta \stackrel{(a)}{\geq} \phi \left(\frac{1}{|I_t \cap [-1, 1]|} \int_{I_t \cap [-1, 1]} |\eta| d\eta \right) \stackrel{(b)}{\geq} \phi(|\eta_t^*|),$$

where in (a), we used Jensen's inequality (the interval $I_t \cap [-1, 1]$ lies entirely in either side of 0 depending on the sign of η_t^* , and $\phi(\cdot)$ is a convex function on $[0, 1]$ as well as on $[-1, 0]$); for (b), observe that $\int_{I_t \cap [-1, 1]} |\eta| d\eta \leq |\eta_t^*|$, and ϕ is radially decreasing.

Furthermore, since $|I_t \cap [-1, 1]| \geq r$, using this, re-arranging the above inequality, and dividing by Z_0 , we get

$$\int_{I_t \cap [-1, 1]} \frac{\phi(|\eta|)}{Z_0} d\eta \geq \frac{r}{Z_0} \phi(|\eta_t^*|) = \frac{r \ln \ln c}{2|\eta_t^*| \left(\ln \frac{c}{|\eta_t^*|} \right) \left(\ln \ln \frac{c}{|\eta_t^*|} \right)^2} \geq \frac{r \ln \ln c}{2|\eta_t^*| \left(\ln \frac{c}{r} \right) \left(\ln \ln \frac{c}{r} \right)^2},$$

where, for the last inequality, we used $r < |\eta_t^*| < 1$. This completes the proof. \blacksquare

Proof [Proof of Lemma 12] First, observe the following:

$$f_t(\eta) = f_t(\eta_t^*) - \frac{V_t}{2}(\eta - \eta_t^*)^2.$$

For $\rho \in (0, 1)$, choose $r := |\eta_t^*| - 1 + \rho$, and set $s_t := \text{Sign}(\eta_t^*) \in \{-1, 1\}$. Then, integrating the exponent of the above at the intersection $\{|\eta - \eta_t^*| \leq r\} \cap [-1, 1]$, we get

$$W_t = \int_{-1}^1 e^{f_t(\eta)} \pi(\eta) d\eta \geq e^{f_t(\eta_t^*) - \frac{V_t r^2}{2}} \int_{\{|\eta - \eta_t^*| \leq r\} \cap [-1, 1]} \pi(\eta) d\eta.$$

Observe that $\{|\eta - s_t| \leq \rho\} \subset \{|\eta - \eta_t^*| \leq r\} \cap [-1, 1]$, and on that subset $|\eta| \in [1 - \rho, 1]$, implying

$$\int_{\{|\eta - \eta_t^*| \leq r\} \cap [-1, 1]} \pi(\eta) d\eta \geq \int_{|\eta - s_t| \leq \rho} \pi(\eta) d\eta \stackrel{(a)}{\geq} \frac{2\rho}{Z_0 \left(\ln \frac{c}{1-\rho} \right) \left(\ln \ln \frac{c}{1-\rho} \right)^2},$$

where in (a), we used that $\frac{1}{|\eta|} \geq 1$, and that

$$\frac{1}{\left(\ln \frac{c}{|\eta|} \right) \left(\ln \ln \frac{c}{|\eta|} \right)^2} \geq \frac{1}{\left(\ln \frac{c}{1-\rho} \right) \left(\ln \ln \frac{c}{1-\rho} \right)^2}.$$

Proof [Proof of Proposition 13] We first show that (15) holds. Towards this, suppose at time $t > 0$, $|\eta_t^*| > 1$. Then $|S_t| > V_t$. Let $s_t := \text{sign}(S_t) \in \{-1, +1\}$. Then, for every $\eta \in [s_t(1 - \rho), s_t]$ (to be interpreted as either $[-1, -1 + \rho]$ or $[1 - \rho, 1]$, depending on sign of s_t), we have $\eta S_t \geq (1 - \rho)|S_t|$ and $\eta S_t - \frac{1}{2}\eta^2 V_t$ is increasing in η for $\eta > 0$ and decreasing otherwise. Hence,

$$f_t(\eta) \geq (1 - \rho)|S_t| - \frac{V_t}{2}(1 - \rho)^2 \geq (1 - \rho)V_t - \frac{V_t}{2}(1 - \rho)^2 = \frac{V_t}{2}(1 - \rho^2).$$

Therefore,

$$\ln W_t = \ln \int_{-1}^1 e^{f_t(\eta)} \pi(\eta) d\eta \geq \ln \left(\int_{1-\rho}^1 \pi(\eta) d\eta \right) + \frac{V_t}{2}(1 - \rho^2).$$

Now, on \mathcal{E}_α , since $\sup_t \ln W_t \leq \ln(1/\alpha)$, from the above inequality, we must also have

$$\begin{aligned} \frac{V_t}{2}(1 - \rho^2) &\leq \ln \frac{1}{\alpha} - \ln \left(\int_{1-\rho}^1 \pi(\eta) d\eta \right) \\ &\stackrel{(a)}{\leq} \ln \frac{1}{\alpha} - \ln \frac{\rho}{Z_0} + \ln \ln \frac{c}{1-\rho} + 2 \ln \ln \ln \frac{c}{1-\rho}, \end{aligned}$$

which implies, for such a t , $V_t \leq V_\alpha$, with V_α defined in (15). In (a), we used that

$$\ln \int_{1-\rho}^1 \pi(\eta) d\eta \geq \ln \frac{\rho}{Z_0} - \ln \ln \frac{c}{1-\rho} - 2 \ln \ln \ln \frac{c}{1-\rho}, \quad (19)$$

which follows by observing that for $\eta \in [1 - \rho, 1]$,

$$\eta < 1, \quad \text{and} \quad \ln \frac{c}{|\eta|} \left(\ln \ln \frac{c}{\eta} \right)^2 \leq \ln \frac{c}{1-\rho} \left(\ln \ln \frac{c}{1-\rho} \right)^2,$$

and hence, $\frac{1}{\pi(\eta)} \leq Z_0 (\ln \frac{c}{1-\rho}) (\ln \ln \frac{c}{1-\rho})^2$.

For (16), let $\alpha_t := |\eta_t^*| - 1$. As above, for $\eta \in [s_t(1 - \rho), s_t]$ we have

$$f_t(\eta) \geq (1 - \rho)|S_t| - \frac{V_t}{2} = \frac{V_t}{2} [(1 - 2\rho) + 2(1 - \rho)\alpha_t].$$

Then,

$$\begin{aligned} \ln W_t &\geq \ln \left(\int_{1-\rho}^1 \pi(\eta) d\eta \right) + \frac{V_t}{2} ((1 - 2\rho) + 2(1 - \rho)\alpha_t) \\ &\geq \ln \left(\int_{1-\rho}^1 \pi(\eta) d\eta \right) + V_t (1 - \rho) \alpha_t. \end{aligned}$$

On \mathcal{E}_α , since $\sup_t \ln W_t \leq \ln(1/\alpha)$; using this and (19) in the above inequality, and rearranging,

$$\alpha_t \leq \frac{(1 + \rho)}{2V_t} V_\alpha.$$

Further, from the previous part, we also have that $V_t \leq V_\alpha$, and recall that Lemma 12 gives a bound on R_t in this case. Thus, for t such that $|S_t| \geq V_t$, the first term in the bound in Lemma 12 is bounded as below:

$$\begin{aligned} \frac{V_t}{2}(\alpha_t + \rho)^2 &= \frac{V_t \alpha_t^2}{2} + \frac{V_t \rho^2}{2} + V_t \alpha_t \rho \\ &\leq \frac{(1 + \rho)^2 V_\alpha^2}{8V_t} + \frac{V_\alpha \rho^2}{2} + \frac{\rho(1 + \rho)V_\alpha}{2} \\ &= \frac{(1 + \rho)^2 V_\alpha^2}{8V_t} + \frac{\rho V_\alpha}{2} (1 + 2\rho). \end{aligned}$$

By subtracting the constant prior mass term from Lemma 12, we get the inequality in (16). ■

Appendix D. The Sub-Gaussian Betting Game

In this section, we present the betting game that we consider in this work. To build intuitions, we begin by presenting the game in the simplest setting of Gaussian observations and extend it to the more general sub-Gaussian setting towards the end.

Gaussian data. Consider a game between two risk-neutral players: a buyer (named Skeptic) and a seller (named Forecaster) of bets. In this game, Forecaster believes that the distribution of an unseen outcome is $N(0, 1)$ (we will generalize this later). Skeptic is skeptical about Forecaster’s belief and instead believes that it is a mean-shifted Gaussian, i.e., $N(\cdot, 1)$, with non-zero mean.

For every $\eta \in \mathbb{R}$, Forecaster sells a bet η such that every dollar placed on η (before observing the data) pays back $e^{\eta X - \eta^2/2}$ dollars, when the data X is revealed. Why are these the only bets offered? Because it is well understood that likelihood ratios are the only admissible bets against a point null (Ramdas and Wang, 2025, Chapter 3), and $e^{\eta X - \eta^2/2}$ is just the likelihood ratio between $N(\eta, 1)$ and $N(0, 1)$, which is the log-optimal bet if Skeptic knew the alternative to be $N(\eta, 1)$. Note that any such bet is fair from Forecaster’s viewpoint, because the expected value of $e^{\eta X - \eta^2/2}$ (the amount Forecaster pays up per dollar of investment in η) equals one if $X \sim N(0, 1)$.

Starting with a unit wealth (i.e., $W_0 = 1$), Skeptic bets against Forecaster, hoping to get rich. The betting game between Forecaster and Skeptic can be described as below.

Let $\mathcal{P}(\mathbb{R})$ denote the collection of all probability measures with support in \mathbb{R} . For $t = 1, 2, \dots$:

- Forecaster makes the aforementioned bets available.
- For some $\pi_t \in \mathcal{P}(\mathbb{R})$, Skeptic invests a fraction $\pi_t(\eta)$ of the current wealth, i.e., $W_{t-1}\pi_t(\eta)$ dollars, in the bet η , for every $\eta \in \mathbb{R}$.
- Reality reveals the data X_t .
- Forecaster pays back $e^{\eta X_t - \eta^2/2}$ dollars to Skeptic for every dollar invested in the bet η .
- The wealth of Skeptic at time t thus becomes $W_t = W_{t-1} \int \pi_t(\eta) e^{\eta X_t - \eta^2/2} d\eta$.

The goal of Skeptic is to choose a betting strategy $(\pi_t)_{t \geq 1}$ whose wealth is close to that of the best fixed betting strategy in hindsight. (A slightly weaker goal for Skeptic is the following: when the true mean is nonzero, they would like to earn as much wealth as an oracle that knows the true mean; of course, when the true mean is zero, neither can earn wealth systematically.)

Formally, the performance of a given betting strategy can be evaluated by considering the difference between its log-wealth and that generated by using the best-in-hindsight strategy. This difference, termed as regret till time t , is given by

$$R_t := \ln \left(\max_{\eta \in \mathbb{R}} e^{\eta \sum_{i=1}^t X_i - \eta^2 t/2} \right) - \ln W_t = \max_{\eta \in \mathbb{R}} \left(\eta \sum_{i=1}^t X_i - \eta^2 t/2 \right) - \ln W_t.$$

Clearly, $\eta_t^* = S_t/t$ optimizes the wealth in hindsight above, so that the first term just equals $S_t^2/(2t)$. The question we study is whether there is a betting strategy that can make R_t above grow only as $\log \log t$ (since achieving a $\log t$ regret is straightforward, see Section 3).

In this work, we will use a Bayesian betting strategy that starts with a “prior” $\pi \in \mathcal{P}(\mathbb{R})$, and at round t , bets using a “posterior” given by Bayes’ rule as:

$$\pi_t(\eta) = \frac{\pi(\eta) \exp(\eta \sum_{i \leq t-1} X_i - \frac{\eta^2}{2}(t-1))}{\int_{\mathbb{R}} \exp(\eta \sum_{i \leq t-1} X_i - \frac{\eta^2}{2}(t-1)) \pi(\eta) d\eta}.$$

Prior and posterior are put in quotes for two reasons. First, there is nothing Bayesian assumed in the problem setup, and Bayes rule is simply being used as a way of describing the strategy. One could have used “mixture” or “working prior” instead. Second, the likelihood term is only a valid likelihood in this Gaussian setup, but will no longer be one in the more general sub-Gaussian process setup described later. We claim that in this case, the expression for W_t simplifies to

$$W_t = \int_{\mathbb{R}} e^{\eta \sum_{i=1}^t X_i - \eta^2 t / 2} \pi(\eta) d\eta.$$

This is directly checked by noticing that the above definition of W_t results in the normalizing constant $\pi_t(\eta)$ equaling W_{t-1} , and thus W_t does indeed satisfy the relationship

$$W_t = W_{t-1} \int \pi_t(\eta) e^{\eta X_t - \eta^2 / 2} d\eta.$$

In our paper, we found it easier to directly work with the expression for the wealth process W_t above than to explicitly specify and manipulate the per-round bets π_t , which are left implicit.

Remark 14 *As mentioned before, the η -bet that Forecaster makes available to Skeptic in the above game is actually the likelihood ratio between $N(\eta, 1)$ and $N(0, 1)$, and Skeptic bets on the shifts η in the mean. If the Forecaster instead believes that the distribution of the unseen outcome is $N(m_0, 1)$, for some $m_0 \neq 0$, and the Skeptic believes that it is a mean-shifted Gaussian, that is, $\{N(m, 1) : m \neq m_0\}$, then in the above game, the bets that the Forecaster makes available at each time will be $\{e^{\eta(X-m_0)-\eta^2/2} : \eta \in \mathbb{R}\}$. Again, $e^{\eta(X-m_0)-\eta^2/2}$ is a likelihood ratio between $N(m_0 - \eta, 1)$ and $N(m_0, 1)$. Thus, Forecaster’s belief that the distribution is 0-mean Gaussian is without loss of generality, because one can simply redefine $X' = X - m_0$ and proceed. This also holds if the Forecaster changes their prediction to $N(m_t, 1)$ in each round t : such time-varying forecasts reduce to zero mean forecasts by recentering. Finally, the assumption of the data being unit variance is also without loss of generality. Had the Forecaster claimed $N(0, \sigma)$ instead, we could have redefined the game using $X' = X/\sigma$ (or more generally $X'_t = (X_t - \mu_t)/\sigma_t$). As long as Skeptic is only challenging the mean forecast (which is what we study in this paper), recentering and rescaling do not qualitatively change the game.*

Sub-Gaussian process. Recall the functions S_t and V_t introduced in Section 2.1. We now consider Forecaster who believes that the process S_t is sub-Gaussian with a variance proxy V_t , and Skeptic who is skeptical about its mean. The protocol of the game remains the same as earlier, except that a unit dollar invested in the bet η offered by Forecaster at time t now pays back

$$e^{\eta(S_t - S_{t-1}) - \frac{\eta^2}{2}(V_t - V_{t-1})}.$$

With this, the wealth process of Skeptic becomes

$$W_t = W_{t-1} \int_{\mathbb{R}} \pi_t(\eta) e^{\eta(S_t - S_{t-1}) - \frac{\eta^2}{2}(V_t - V_{t-1})} d\eta.$$

We can similarly define the regret of Skeptic's wealth or strategy. We focus on the performance of the class of strategies whose wealth at time t takes the form

$$W_t = \int_{\mathbb{R}} e^{\eta S_t - \frac{\eta^2}{2} V_t} \pi(\eta) d\eta.$$

The sub-Gaussian process (Definition 2) captures a large class of distributions. Thus, the game described in the sub-Gaussian process setting is very general. In particular, the game in the Gaussian data setting considered earlier, can be recovered by setting $S_t = \sum_{i=1}^t X_i$ and $V_t = t$. In fact, the games described above can be generalized even further, as discussed next.

Forecaster's belief about the nature are captured by the bets he makes available. It can even change across rounds. For example, consider Forecaster who believes that the data is $N(0, 1)$ at $t = 1$, $N(m_0, \sigma^2)$ at $t = 2$ (conditional on X_1), and a distribution that is symmetric around 0 at $t = 3$ (conditional on X_1, X_2), and so on (in fact, the forecast at time 3 can depend on the first two outcomes: for example, if X_1 is positive, X_3 can be forecast to be symmetric, and if X_1 is negative, X_3 can be forecast to any mean-zero and unit variance distribution). Recall from Appendix A that each of these forecasts can be modeled as sub-Gaussian processes.

More generally, Forecaster does not have to specify their beliefs as sets of distributions, but simply specify the available bets. Since we work with sub-Gaussian bets, they only need to specify how they think S_t and V_t will change at round t .

In our game, Skeptic believes that the data is a mean-shifted version of Forecaster's belief. In the time-varying example considered in the previous paragraph, a unit dollar invested in bet η gives back $e^{\eta X_1 - \eta^2/2}$ dollars when the data X_1 is revealed at $t = 1$. At $t = 2$, it returns $e^{\eta(X_2 - m_0) - \eta^2 \sigma^2/2}$ dollars, at $t = 3$, it pays off $e^{\eta X_3 - X_3^2 \eta^2/2}$, etc. The game protocol remains the same as earlier.

Our results deliver low regret in this much more general setting (than the Gaussian setting specified earlier).