

A composite generalization of Ville’s martingale theorem using e-processes

Johannes Ruf* Martin Larsson† Wouter M. Koolen‡
Aaditya Ramdas§

Abstract

We provide a composite version of Ville’s theorem that an event has zero measure if and only if there exists a nonnegative martingale which explodes to infinity when that event occurs. This is a classic result connecting measure-theoretic probability to the sequence-by-sequence game-theoretic probability, recently developed by Shafer and Vovk. Our extension of Ville’s result involves appropriate composite generalizations of nonnegative martingales and measure-zero events: these are respectively provided by “e-processes”, and a new inverse capital outer measure. We then develop a novel line-crossing inequality for sums of random variables which are only required to have a finite first moment, which we use to prove a composite version of the strong law of large numbers (SLLN). This allows us to show that violation of the SLLN is an event of outer measure zero and that our e-process explodes to infinity on every such violating sequence, while this is provably not achievable with a nonnegative (super)martingale.

Keywords: e-process; martingale; measure zero; null set; optional stopping.

MSC2020 subject classifications: Primary 60A10, 60F15, Secondary 60G40.

Submitted to EJP on May 12, 2021, final version accepted on December 13, 2014.

1 Introduction

In his PhD thesis, Ville [26] obtained a martingale characterization of the zero probability events on the filtered probability space generated by an infinite sequence of independent unbiased coin flips. Informally, Ville proved that

an event A has probability $P(A) = 0$ if and only if there exists a nonnegative P-martingale $M = (M_t)_{t \in \mathbb{N}_0}$ which starts at one and diverges to infinity on A .

*Department of Mathematics, London School of Economics. E-mail: j.ruf@lse.ac.uk

†Department of Mathematical Sciences, Carnegie Mellon University. E-mail: larsson@cmu.edu

‡Machine Learning Group, CWI Amsterdam. E-mail: wmkoolen@cwi.nl

§Departments of Statistics and ML, Carnegie Mellon University. E-mail: aramdas@cmu.edu

Here, \mathbb{P} is the probability measure under which the coins are independent and unbiased.

In game-theoretic terms, Ville proved that for any event A having probability zero, one can construct an explicit betting strategy, which starts at unit wealth and never risks more wealth than it has, such that the wealth increases to infinity if that event A occurs. Via this seminal result, applied specifically to the law of the iterated logarithm (LIL), Ville was instrumental in bringing martingales into the forefront of probability theory. Ville's result is also central to the philosophy of probability theory, because it gives an actionable and concrete interpretation to the relatively abstract concept of a measure-zero event. Different measure-zero events — for example, sequences violating the strong law of large numbers (SLLN) and those violating the LIL — obviously have the same probability, but they result in different betting strategies.

Although Ville proved his theorem for binary sequences, his result turns out to be far more general, holding true in general filtered spaces. There is also a version for events of positive measure: an event A has \mathbb{P} -measure at most a if and only if there exists a \mathbb{P} -martingale which starts at one and exceeds any level below $1/a$ if the event A occurs. See Chapter 8.5 in the book by Shafer and Vovk [23] for a detailed treatment and proof.

Ville's result has implications for sequential hypothesis testing: the nonnegative \mathbb{P} -martingale M can be used to construct a sequential test for the null hypothesis that the data sequence was generated according to \mathbb{P} . For any $\alpha \in (0, 1)$, the test which rejects the null the first time that $M_t \geq 1/\alpha$ has type-I error at most α :

$$\mathbb{P} \left(M_t \geq \frac{1}{\alpha} \text{ for some } t \in \mathbb{N} \right) \leq \alpha. \quad (1.1)$$

This inequality, which was proved in the same thesis, is commonly referred to as *Ville's inequality*. Doob, who read Ville's thesis, realized the importance of martingales and credited this particular inequality to Ville; see Doob [4].

We briefly mention how this work connects to modern developments in sequential statistics. The idea of constructing nonnegative test statistics that are martingales under the null hypothesis, but increase under relevant alternative hypotheses, has long (implicitly or explicitly) been a cornerstone of sequential statistics; indeed, the sequential probability ratio test by Wald [27] is based on the same principle (sequential likelihood ratios are easily checked to be nonnegative martingales). Following initial work by Robbins [19], this idea has been extended to many different nonparametric settings in recent years [9, 10]; here, nonnegative *supermartingales* provide a natural and powerful generalization of likelihood ratios. There are yet other problems for which even nonnegative supermartingales do not suffice, and one encounters a new concept called an e-process [7, 17, 18]. We revisit this notion very soon since it plays a fundamental role in the current paper. The aforementioned discoveries underpin a “game-theoretic” approach to statistics, which in particular yields hypothesis tests and confidence sets that can be continuously monitored and adaptively stopped; see the recent survey [16] and references therein for an elaboration on such “safe anytime-valid inference.”

Our contributions. In this paper we develop a *composite* version of Ville's theorem, where the single probability measure \mathbb{P} is replaced by an entire family \mathcal{P} of probability measures (and the space of coin flips is replaced by a general filtered measurable space). In this extended setting, two issues must be addressed:

1. First, one needs an appropriate analog of a zero probability event.
2. Second, one needs an appropriate analog of a martingale.

A natural replacement for the notion of a \mathbb{P} -nullset would be that of a \mathcal{P} -*polar set*, which is an event A such that $\mathbb{P}(A) = 0$ for all $\mathbb{P} \in \mathcal{P}$. However, this turns out to be too

weak. The appropriate replacement relies on a certain outer measure $\nu \equiv \nu_{\mathcal{P}}$ introduced in Section 2. We call this the *inverse capital measure*, because for any set A , the inverse of $\nu(A)$ specifies an asymptotic lower bound on the limiting capital achievable by a gambler (in a particular set of games) starting with an initial capital of one and betting against \mathcal{P} , whenever A occurs.

Next, one may attempt to work with processes that are \mathbb{P} -martingales for every $\mathbb{P} \in \mathcal{P}$. However, this requirement turns out to be too strong, and the appropriate replacement is the notion of an *e-process* [7, 17, 18]. An e-process for a family \mathcal{P} of probability measures, called a \mathcal{P} -e-process for brevity, is a nonnegative process $E = (E_t)_{t \in \mathbb{N}_0}$, adapted to some filtration $(\mathcal{F}_t)_{t \in \mathbb{N}_0}$, such that

$$\mathbb{E}_{\mathbb{P}}[E_{\tau}] \leq 1 \text{ for every } \mathbb{P} \in \mathcal{P} \text{ and every stopping time } \tau, \tag{1.2}$$

where $E_{\infty} = \limsup_{t \rightarrow \infty} E_t$ by convention. (The stopping times above are also defined with respect to $(\mathcal{F}_t)_{t \in \mathbb{N}_0}$.)

This is a much weaker property than being a nonnegative \mathbb{P} -martingale starting at one for every $\mathbb{P} \in \mathcal{P}$. In particular, the martingale property implies the e-process property thanks to the optional stopping theorem. An equivalent definition [17] is: an e-process is any nonnegative process that is upper bounded, for every $\mathbb{P} \in \mathcal{P}$, by a \mathbb{P} -martingale starting at one (a different one for each \mathbb{P}). This style of definition was previously used by Howard et al. [9], but without using the term e-process, which is a very recent christening of this property. As a consequence of this alternative characterization, e-processes also satisfy Ville's inequality (1.1), making their magnitude interpretable and actionable as a measure of evidence.

Using the above two notions — that is, e-processes and the inverse capital measure — we derive a composite version of Ville's theorem. Specifically, we prove that

an event A has inverse capital measure zero if and only if there exists an e-process that grows to infinity whenever A occurs.

We also prove an analogous statement for events with positive inverse capital measure.

E-processes have recently been used to solve a number of problem in sequential statistics, because there exist problems in which there are no nontrivial \mathcal{P} -martingales (other than constants), there are no nontrivial \mathcal{P} -supermartingales (other than decreasing processes), but there do exist nontrivial \mathcal{P} -e-processes which grow to infinity when evaluated on data that is sufficiently incompatible with \mathcal{P} . A particularly interesting and fundamental problem where the preceding observations arise, is to test whether a sequence is exchangeable or not [18]. A game-theoretic interpretation for e-processes can be found in the aforementioned work, and applications to meta-analysis of medical studies were described in Grünwald et al. [7]. A thorough treatment of game-theoretic probability can be found in the book by Shafer and Vovk [24], and a recent case for its use in scientific communication was made in the discussion paper by Shafer [22].

Our composite version of Ville's theorem generalizes his original result. Moreover, unpacking the proof of the theorem we see that it constructs relatively explicit gambling strategies for arbitrary inverse capital measure-zero events that gain infinite wealth when the purportedly rare events actually do occur. In order to demonstrate a nontrivial example of its applicability, we prove a new *composite strong law of large numbers*. Here the inverse capital measure-zero event A is the set of sequences whose empirical averages do not converge to a limit. In other words, it is possible to *test* whether such a composite SLLN holds, with the evidence (e-process) growing to infinity whenever the SLLN fails. Moreover, we prove that there is no nonnegative martingale or

supermartingale that achieves this; the concept of an e-process is necessary for this example.

Proving such a composite strong law is itself nontrivial. Along the way, we develop a new L^1 -variant of the celebrated L^2 -type Dubins-Savage inequality. Such inequalities have been extended to settings with finite p -th moments for $1 < p < 2$ by Khan [11], and also to higher moments [9], but the L^1 -type inequality appears to be new and of independent interest.

To summarize, our main contributions are the following:

- Constructing a composite generalization of Ville's theorem which informally reads as: for every event of *inverse capital* measure equal to $a \in [0, 1]$, we can construct an *e-process* that reaches arbitrarily close to $1/a$ if that event occurs.
- Deriving a new, L^1 -type line-crossing inequality, which is an analog of the L^2 -type Dubins-Savage deviation inequality.
- Developing a new composite strong law of large numbers, for which there is no martingale or supermartingale that can detect its failure, but there exists an explicit e-process that can.

To conclude, it is worth remarking the following. Gambling strategies were called martingales in the 19th century, but Ville switched the term's use to characterize the wealth of the gambler, rather than the gambling strategy itself [14]. In some sense, martingales were thrust into the limelight by Ville's result, giving them a central role in modern "measure-theoretic" probability theory, and also demystifying measure zero events to some extent. In a similar sense, our generalization gives an independent justification for the claim that the e-process is an equally fundamental concept for dealing with composite nulls (as martingales are for point nulls): the e-process arises organically when one tries to generalize Ville's seminal result to the composite setting.

2 The inverse capital measure

Consider a filtered measurable space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \in \mathbb{N}})$ where $\mathcal{F} = \sigma(\bigcup_{t \in \mathbb{N}} \mathcal{F}_t)$, and let \mathcal{T} denote the set of all stopping times with respect to this filtration. Fix an arbitrary set \mathcal{P} of probability measures on \mathcal{F} and define a $[0, 1]$ -valued set function ν by

$$\nu(A) = \inf_{\tau \in \mathcal{T}: A \subseteq \{\tau < \infty\}} \sup_{\mathbb{P} \in \mathcal{P}} \mathbb{P}(\tau < \infty), \tag{2.1}$$

for any subset $A \subseteq \Omega$. We call this the *inverse capital* measure in anticipation of its later role in the composite generalization of Ville's theorem. Note that ν depends both on the filtration $(\mathcal{F}_t)_{t \in \mathbb{N}}$ and the set \mathcal{P} .

To interpret $\nu(A)$ one should think of A as a *failure event*. An example that we study in detail below is the event that a certain adapted process $(\bar{X}_t)_{t \in \mathbb{N}}$ fails to converge as t tends to ∞ . One then looks for stopping times τ that detect failure in finite time, meaning that $A \subseteq \{\tau < \infty\}$. Among all such stopping times one searches for those that are *efficient*, uniformly across \mathcal{P} , in the sense that $\sup_{\mathbb{P} \in \mathcal{P}} \mathbb{P}(\tau < \infty)$ is minimal. This can be thought of as minimizing the probability of false alarms.

There is also an interpretation in terms of type-I error control for sequential tests. Indeed, the definition of $\nu(A)$ says that for any $\alpha > \nu(A)$ there exists a level- α sequential test τ for the null hypothesis \mathcal{P} , with the additional property that the null is guaranteed to be rejected if A occurs. Here a stopping time τ is identified with the test that rejects the null when τ occurs, and the level- α property means that the null, if true, will be rejected with probability at most α ; that is, $\sup_{\mathbb{P} \in \mathcal{P}} \mathbb{P}(\tau < \infty) \leq \alpha$. A further interpretation in terms of e-processes arises from Theorem 3.7 below.

There is another natural outer measure, which we call the *maximum likelihood (outer) measure*, namely

$$\mu(A) = \sup_{\mathbb{P} \in \mathcal{P}} \mathbb{P}(A). \tag{2.2}$$

Note that $\mu(A) \leq \nu(A)$ for every $A \in \mathcal{F}$. The inequality can be strict, a fact which is demonstrated and discussed further in Subsection 2.2 below.

2.1 The inverse capital measure is an outer measure

The following lemma gathers some properties of ν .

Lemma 2.1. *The set function ν satisfies the following properties.*

1. *Monotonicity: if $A \subseteq B$ then $\nu(A) \leq \nu(B)$.*
2. *Countable subadditivity: for any sequence of sets $A_n \subseteq \Omega$, $n \in \mathbb{N}$, one has*

$$\nu\left(\bigcup_{n \in \mathbb{N}} A_n\right) \leq \sum_{n \in \mathbb{N}} \nu(A_n).$$

3. *Singletons: if $\mathcal{P} = \{\mathbb{P}\}$ consists of a single element, then $\nu(A) = \mathbb{P}(A)$ for all $A \in \mathcal{F}$.*

In particular, since clearly $\nu(\emptyset) = 0$, the set function ν is an outer measure on Ω .

Proof. Property 1 is immediate from the definition of ν . Let us next prove the remaining two assertions.

Property 2: Consider sets $A_n \subseteq \Omega$, $n \in \mathbb{N}$, and fix any $\varepsilon > 0$. For each n , let τ_n be a stopping time such that $A_n \subseteq \{\tau_n < \infty\}$ and $\sup_{\mathbb{P} \in \mathcal{P}} \mathbb{P}(\tau_n < \infty) \leq \nu(A_n) + \varepsilon 2^{-n}$. Then $\tau = \inf_{n \in \mathbb{N}} \tau_n$ is a stopping time which satisfies $\bigcup_{n \in \mathbb{N}} A_n \subseteq \{\tau < \infty\}$ and, for each $\mathbb{P} \in \mathcal{P}$,

$$\mathbb{P}(\tau < \infty) = \mathbb{P}\left(\bigcup_{n \in \mathbb{N}} \{\tau_n < \infty\}\right) \leq \sum_{n \in \mathbb{N}} \mathbb{P}(\tau_n < \infty) \leq \varepsilon + \sum_{n \in \mathbb{N}} \nu(A_n).$$

Taking supremum over $\mathbb{P} \in \mathcal{P}$ and using that $\varepsilon > 0$ was arbitrary, the result follows.

Property 3: Let $A \in \mathcal{F}$, as assumed. The definition of ν gives $\mathbb{P}(A) \leq \nu(A)$. For the reverse inequality, fix any $\varepsilon > 0$. Since \mathcal{F} is generated by $\bigcup_{t \in \mathbb{N}} \mathcal{F}_t$, the standard Carathéodory construction [1, Section 3] that extends \mathbb{P} to \mathcal{F} shows that we can cover A by a union $B = \bigcup_{t \in \mathbb{N}} A_t$ of disjoint events $A_t \in \mathcal{F}_t$ such that $B \supseteq A$ and $\mathbb{P}(B) \leq \mathbb{P}(A) + \varepsilon$. This ensures that the random variable $\tau = \sum_{t \in \mathbb{N}} t \mathbf{1}_{A_t} + \infty \mathbf{1}_{B^c}$ (meaning that τ equals infinity on B^c) is a stopping time with $\{\tau < \infty\} = B$. Thus by definition, $\nu(A) \leq \mathbb{P}(\tau < \infty) \leq \mathbb{P}(A) + \varepsilon$. Since $\varepsilon > 0$ was arbitrary, we conclude that $\nu(A) \leq \mathbb{P}(A)$. \square

2.2 The inverse capital measure versus the maximum likelihood measure

We begin by noting that the maximum likelihood measure μ defined in (2.2) also satisfies properties 1–3 of Lemma 2.1. As mentioned before, it follows from the definitions that $\mu(A) \leq \nu(A)$ for every $A \in \mathcal{F}$. In the singleton case $\mathcal{P} = \{\mathbb{P}\}$ the inequality is actually an equality since $\mu(A) = \nu(A) = \mathbb{P}(A)$ by Lemma 2.1(3). This implies, in particular, that the maximum likelihood measure can be expressed as

$$\mu(A) = \sup_{\mathbb{P} \in \mathcal{P}} \inf_{\tau \in \mathcal{T}: A \subseteq \{\tau < \infty\}} \mathbb{P}(\tau < \infty),$$

which provides a direct comparison to (2.1). Moreover, for general \mathcal{P} , we have $\nu(A) = \mu(A)$ for any “finite-time” event A of the form $A = \{\tau < \infty\}$. However, the inequality can be strict in general. In fact, we show in Example 2.5 that it is possible to have $\mu(A) = 0$ but $\nu(A) = 1$. In particular, the collection of ν -nullsets can be strictly smaller than the collection of \mathcal{P} -polar sets, which we recall are sets A such that $\mathbb{P}(A) = 0$ for all $\mathbb{P} \in \mathcal{P}$. A

key message of Theorem 3.7, and indeed of this paper, is that in the context of sequential testing and composite versions of Ville's results, ν turns out to be the more relevant object, and not μ .

Example 2.2. Consider the sample space $\Omega = \{0, 1\}^{\mathbb{N}}$, the space of all binary sequences. Let $X = (X_t)_{t \in \mathbb{N}}$ denote the canonical process and set $\mathcal{F}_t = \sigma(X_s, s \leq t)$, $t \in \mathbb{N}$. Let \mathbb{P}_n assign probability $1/2$ to the sequence $0^{\mathbb{N}}$ (all zeros) and to the sequence $0^n \cdot 1 \cdot 0^{\mathbb{N}}$ (a single one occurs after n zeros). We consider $\mathcal{P} = \{\mathbb{P}_n\}_{n \in \mathbb{N}_0}$. Let $A = \{0^{\mathbb{N}}\}$. Clearly $\mathbb{P}_n(A) = 1/2$ for all $n \in \mathbb{N}_0$, and hence $\mu(A) = 1/2$.

However, we have $\nu(A) = 1$. To see why, observe that any stopping time τ that stops on A must stop after some finite number of zeros, say m . But then for $n \geq m$, we have $\mathbb{P}_n(\tau < \infty) = 1$. Since this holds for any τ such that $A \subseteq \{\tau < \infty\}$, we conclude $\nu(A) = 1$.

Curiously, on the complementary event A^c (there is a one somewhere), $\nu(A^c) = 1/2$. This claim is witnessed by the stopping time that stops when it encounters a one.

Remark 2.3. In the singleton case of $\mathcal{P} = \{\mathbb{P}\}$, which is what Ville considered, the outer measure ν of any measurable set A equals its probability $\mathbb{P}(A)$, which does not depend on the filtration. In the general composite case, however, the underlying filtration matters in the definition of ν because the filtration determines the set of stopping times. In particular, the set of ν nullsets depends on the filtration. This can be seen by considering the extreme case where the filtration $(\mathcal{F}_t)_{t \in \mathbb{N}_0}$ is replaced by the constant filtration $(\mathcal{G}_t)_{t \in \mathbb{N}_0}$ where $\mathcal{G}_t = \mathcal{F}$ for all t . In this case ν becomes equal to the maximum likelihood measure μ in (2.2). In general, a larger filtration leads to smaller values of ν (but, of course, it cannot decrease below μ) and to a larger family of increasing e-processes. In this way, more information can lead to more powerful tests.

One may ask for conditions under which ν and μ are equal on \mathcal{F} . An abstract condition of this kind can be extracted directly from the definition of ν . Indeed, the definition yields

$$\mu(A) \leq \nu(A) \leq \mu(A) + \inf_{\tau \in \mathcal{T}: A \subseteq \{\tau < \infty\}} \mu(\{\tau < \infty\} \setminus A).$$

Thus ν and μ coincide on \mathcal{F} provided that any $A \in \mathcal{F}$ admits a *uniform outer approximation* by sets of the form $\{\tau < \infty\}$, $\tau \in \mathcal{T}$, in the sense that there is a non-decreasing sequence of stopping times τ_n such that $A \subset \{\tau_n < \infty\}$ for all n and $\sup_{\mathbb{P} \in \mathcal{P}} \mathbb{P}(\{\tau_n < \infty\} \setminus A) \downarrow 0$ as $n \rightarrow \infty$.

Remark 2.4. For certain choices of \mathcal{P} , the set function μ is *submodular*, meaning that

$$\mu(A \cup B) + \mu(A \cap B) \leq \mu(A) + \mu(B)$$

for all $A, B \in \mathcal{F}$. In this case, submodularity is inherited by ν . Indeed, suppose $A \subset \{\tau < \infty\}$ and $B \subset \{\sigma < \infty\}$ for some $\tau, \sigma \in \mathcal{T}$. Then, because $\{\tau < \infty\} \cup \{\sigma < \infty\} = \{\tau \wedge \sigma < \infty\}$ and $\{\tau < \infty\} \cap \{\sigma < \infty\} = \{\tau \vee \sigma < \infty\}$, the definition of ν and submodularity of μ yield

$$\begin{aligned} \nu(A \cup B) + \nu(A \cap B) &\leq \mu(\tau \wedge \sigma < \infty) + \mu(\tau \vee \sigma < \infty) \\ &\leq \mu(\tau < \infty) + \mu(\sigma < \infty). \end{aligned}$$

Taking the infimum of the right-hand side over all such stopping time τ, σ yields submodularity of ν , namely

$$\nu(A \cup B) + \nu(A \cap B) \leq \nu(A) + \nu(B).$$

However, μ is not always submodular. For example, consider the four-point space $\Omega = \{a, b, c, d\}$ along with $\mathcal{P} = \{\mathbb{P}_1, \mathbb{P}_2\}$ where

$$\mathbb{P}_1 = \left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}, 0\right), \quad \mathbb{P}_2 = \left(0, \frac{2}{3}, 0, \frac{1}{3}\right).$$

In this case we have for $A = \{a, b\}$ and $B = \{b, c\}$ that $\mu(A \cup B) + \mu(A \cap B) = \frac{5}{3}$ but $\mu(A) + \mu(B) = \frac{4}{3}$. This violates submodularity. As a consequence, since ν may coincide with μ (for instance if the filtration is constant), ν may fail to be submodular as well. Submodularity of ν , when it holds, implies that the Choquet integral $\int f d\nu$ with respect to ν of any bounded \mathcal{F} -measurable real-valued function f on Ω admits the representation

$$\int f d\nu = \max_{\mathbb{P} \in \mathcal{P}^*} \mathbb{E}_{\mathbb{P}}[f],$$

where \mathcal{P}^* is the set of all finitely additive probability measures \mathbb{P} on \mathcal{F} such that $\mathbb{P}(A) \leq \nu(A)$ for all $A \in \mathcal{F}$. The Choquet integral is then both subadditive and comonotonically additive. See [6] for more details, in particular Theorem 4.88 and Corollary 4.89. It is an interesting open question what can be said about the relation between \mathcal{P}^* , \mathcal{P} , and the filtration $(\mathcal{F}_t)_{t \in \mathbb{N}_0}$. We would like to thank an anonymous referee for bringing the question of submodularity to our attention.

2.3 The law of large numbers for truncated Cauchy distributions

We now present a more sophisticated example that is of relevance later in the paper. A key added feature compared to Example 2.2 is that we now consider i.i.d. sequences.

Example 2.5. Consider the sample space $\Omega = \mathbb{R}^{\mathbb{N}}$, the space of all infinite real-valued sequences. Let $X = (X_t)_{t \in \mathbb{N}}$ denote the canonical process and set $\mathcal{F}_t = \sigma(X_s, s \leq t)$, $t \in \mathbb{N}$. We will consider a collection \mathcal{P} of laws under which X is i.i.d. with Cauchy distributions that are truncated at higher and higher levels. This will lead to a situation where $\mu(A) = 0$, $\nu(A) = 1$, and $\nu(A^c) = 1$ hold simultaneously for the same event $A = A_{\text{div}}$ given by

$$A_{\text{div}} = \left\{ \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{s=1}^t X_s \text{ does not exist} \right\}. \tag{2.3}$$

Specifically, let \mathbb{Q} denote the law under which X is i.i.d. Cauchy (centered at the origin). For each $a \in (0, \infty)$ we define the truncation of X at magnitude a by

$$X_t^a = X_t \mathbf{1}_{\{|X_t| < a\}} + a \mathbf{1}_{\{X_t \geq a\}} - a \mathbf{1}_{\{X_t \leq -a\}}, \quad t \in \mathbb{N}.$$

Then, under \mathbb{Q} , the process $X^a = (X_t^a)_{t \in \mathbb{N}}$ is i.i.d. with a truncated Cauchy distribution. More precisely, the marginal distribution coincides with the Cauchy on $(-a, a)$ and collects the remaining mass symmetrically at $\pm a$. We let \mathbb{P}_a denote the law of $(X_t^a)_{t \in \mathbb{N}}$ under \mathbb{Q} , and consider the family

$$\mathcal{P} = \{\mathbb{P}_a : a \in (0, \infty)\}.$$

Proposition 2.6. *For the class \mathcal{P} of truncated Cauchy distributions defined above, we have*

$$\mu(A) = 0, \text{ but } \nu(A_{\text{div}}) = 1 = \nu(A_{\text{div}}^c). \tag{2.4}$$

Proof. For each $a \in (0, \infty)$ we have $\mathbb{E}_{\mathbb{P}_a}[|X_1|] = \mathbb{E}_{\mathbb{Q}}[|X_1^a|] \leq a$ and $\mathbb{E}_{\mathbb{P}_a}[X_1] = \mathbb{E}_{\mathbb{Q}}[X_1^a] = 0$. Thus the strong law of large numbers yields

$$\mathbb{P}_a(A_{\text{div}}) = 0,$$

for any $a > 0$. Thus, $\sup_{\mathbb{P} \in \mathcal{P}} \mathbb{P}(A) = \mu(A) = 0$, as claimed.

In order to reason about ν , consider any stopping time τ such that $A_{\text{div}} \subseteq \{\tau < \infty\}$. We write $\tau = \tau(X)$ to emphasize that τ is a function of the trajectory $X = (X_t)_{t \in \mathbb{N}}$. We have

$$\begin{aligned} \mathbb{P}_a(\tau(X) < \infty) &\geq \mathbb{P}_a(\tau(X) < \infty \text{ and } |X_t| < a \text{ for all } t \leq \tau(X)) \\ &= \mathbb{Q}(\tau(X^a) < \infty \text{ and } |X_t^a| < a \text{ for all } t \leq \tau(X^a)) \\ &= \mathbb{Q}(\tau(X) < \infty \text{ and } |X_t| < a \text{ for all } t \leq \tau(X)), \end{aligned} \tag{2.5}$$

where in the last equality we used that $X_t^a = X_t$ whenever either value belongs to $(-a, a)$, and that $\tau(X)$ only depends on the values of X_t for $t \leq \tau(X)$. Next, since X is i.i.d. Cauchy under \mathbb{Q} , and hence the sample averages almost surely fail to converge, we have $\mathbb{Q}(\tau(X) < \infty) \geq \mathbb{Q}(A_{\text{div}}) = 1$. Thus

$$\begin{aligned} & \mathbb{Q}(\tau(X) < \infty \text{ and } |X_t| < a \text{ for all } t \leq \tau(X)) \\ &= \mathbb{Q}(|X_t| < a \text{ for all } t \leq \tau(X)) \\ &\geq \mathbb{Q}(|X_t| < a \text{ for all } t \leq T) - \mathbb{Q}(\tau(X) > T) \end{aligned} \tag{2.6}$$

for any $T \in \mathbb{N}$. Fix any $\varepsilon > 0$ and choose T large enough that $\mathbb{Q}(\tau(X) > T) \leq \varepsilon$, which is possible because $\mathbb{Q}(\tau(X) < \infty) = 1$. Combining (2.5) and (2.6) we then deduce that

$$\mathbb{P}_a(\tau(X) < \infty) \geq \mathbb{Q}(|X_t| < a \text{ for all } t \leq T) - \varepsilon.$$

Sending a to ∞ we find that $\sup_{a \in (0, \infty)} \mathbb{P}_a(\tau(X) < \infty) \geq 1 - \varepsilon$. Since τ was an arbitrary stopping time with $A_{\text{div}} \subseteq \{\tau < \infty\}$ it follows that $\nu(A_{\text{div}}) \geq 1 - \varepsilon$. Since this holds for every $\varepsilon > 0$, we obtain the first equality in (2.4).

Despite the fact that (2.4) holds, we also have $\nu(A_{\text{conv}}) = 1$ where

$$A_{\text{conv}} = A_{\text{div}}^c = \left\{ \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{s=1}^t X_s \text{ exists} \right\}.$$

Indeed, $\nu(A_{\text{conv}}) \geq \mu(A_{\text{conv}}) = \sup_{\mathbb{P} \in \mathcal{P}} \mathbb{P}(A_{\text{conv}}) = 1$ by the strong law of large numbers. □

Note that (2.4) is not a contradiction, since ν is an outer measure, and not a probability measure.

Much intuition from the above example will be necessary when we prove a composite version of the strong law of large numbers in Section 4, where we impose a uniform integrability condition on the class, in order for A_{div} to be a ν -nullset (as one may expect), as opposed to a full outer-measure set in the above example.

3 A composite generalization of Ville's theorem

3.1 Statement and proof of the composite version

We now have the necessary technical instruments in place to present our first main result, which is a composite version of Ville's theorem in terms of e-processes (recall their defining property in (1.2)).

Theorem 3.1. *Let $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \in \mathbb{N}_0})$ be a filtered measurable space with $\mathcal{F} = \sigma(\bigcup_{t \in \mathbb{N}_0} \mathcal{F}_t)$, let \mathcal{P} be a family of probability measures on \mathcal{F} , let ν be the inverse capital outer measure from (2.1), and let $A \subseteq \Omega$ be arbitrary. Then,*

$$\nu(A) = 0 \text{ if and only if there exists a } \mathcal{P}\text{-e-process } (E_t)_{t \in \mathbb{N}_0} \text{ such that } \lim_{t \rightarrow \infty} E_t = \infty \text{ on } A.$$

In the context of the above theorem, any \mathcal{P} -e-process that explodes to infinity on A will be called a "witness e-process for A ". The proof below explicitly constructs such a witness e-process, and we give even more concrete witness e-processes in special cases later.

Proof. If $\nu(A) = 0$, then for each $n \in \mathbb{N}$ there exists a stopping time τ_n such that $A \subseteq \{\tau_n < \infty\}$ and $\sup_{\mathbb{P} \in \mathcal{P}} \mathbb{P}(\tau_n < \infty) \leq 2^{-n}$. Define

$$E_t = \sum_{n \in \mathbb{N}} \mathbf{1}_{\{\tau_n \leq t\}}, \quad t \in \mathbb{N}_0. \tag{3.1}$$

This process is nondecreasing and satisfies $E_\infty = \lim_{t \rightarrow \infty} E_t = \sum_{n \in \mathbb{N}} \mathbf{1}_{\{\tau_n < \infty\}}$. In particular, $E_\infty = \infty$ on A . Moreover, $(E_t)_{t \in \mathbb{N}_0}$ is an e-process because for any stopping time τ and $\mathbb{P} \in \mathcal{P}$,

$$\mathbb{E}_{\mathbb{P}}[E_\tau] \leq \mathbb{E}_{\mathbb{P}}[E_\infty] = \sum_{n \in \mathbb{N}} \mathbb{P}(\tau_n < \infty) \leq 1.$$

Conversely, suppose there is an e-process $(E_t)_{t \in \mathbb{N}_0}$ such that $\lim_{t \rightarrow \infty} E_t = \infty$ on A . Pick any $a \in (0, \infty)$ and let $\tau = \inf\{t \geq 0: E_t \geq a\}$. Then $A \subseteq \{\tau < \infty\}$, and the definition of an e-process yields

$$\mathbb{P}(\tau < \infty) \leq \frac{1}{a} \mathbb{E}_{\mathbb{P}}[\mathbf{1}_{\{\tau < \infty\}} E_\tau] \leq \frac{1}{a}$$

for all $\mathbb{P} \in \mathcal{P}$. Thus $\nu(A) \leq 1/a$ and hence, since $a \in (0, \infty)$ was arbitrary, $\nu(A) = 0$. \square

Several remarks are in order.

Remark 3.2 (The singleton case). In the singleton case $\mathcal{P} = \{\mathbb{P}\}$, this result does not quite recover the original theorem of Ville, because our e-process does not reduce to being a \mathbb{P} -martingale. This is easily remedied by modifying (3.1) to $E_t = \sum_{n \in \mathbb{N}} \mathbb{P}(\tau_n < \infty \mid \mathcal{F}_t)$, which is a \mathbb{P} -martingale. To be precise, we must choose versions of the conditional probabilities $\mathbb{P}(\tau_n < \infty \mid \mathcal{F}_t)$ that are identically equal to one on $\{\tau_n \leq t\}$ (not just almost surely equal to one, which is automatic), in order to obtain $\lim_{t \rightarrow \infty} E_t = \infty$ everywhere on A (not just almost everywhere). The original proof by Ville, or the version in Chapter 8.5 of [23], carefully accounts for this subtlety.

Remark 3.3 ($\mu(A) = 0$ does not suffice). We clarify here that a version of the theorem with $\nu(A) = 0$ being replaced by $\mu(A) = 0$ is incorrect. Indeed, Section 2.3 demonstrated an explicit event (for truncated Cauchy distributions) with $\mu(A) = 0$ but $\nu(A) = 1$. If $\mu(A) = 0$ had sufficed for producing a witness \mathcal{P} -e-process (or supermartingale, or martingale) for A , then the converse direction of Theorem 3.1 could be invoked to then further conclude that $\nu(A) = 0$, which yields a contradiction.

Remark 3.4 (Nonnegative \mathcal{P} -supermartingales do not suffice). In the general non-singleton case, one cannot expect an e-process to be a martingale under every $\mathbb{P} \in \mathcal{P}$. (We demonstrate an example in Subsection 4.3 for which no \mathcal{P} -supermartingale can increase to infinity, but an e-process can.) However, in the special case that all elements of \mathcal{P} are mutually locally absolutely continuous and \mathcal{P} satisfies a certain geometric condition known as “fork-convexity”, then it is possible to construct the e-process to be a nonnegative \mathcal{P} -supermartingale. This requires technical machinery involving composite Snell envelopes, that we do not develop further in this paper; see [18, Section 3].

Remark 3.5 (Countable unions of events). If one considers a countable number of events $\{A_n\}_{n \in \mathbb{N}}$ such that $\nu(A_n) = 0$ for each of them, Lemma 2.1 implies that their union also satisfies $\nu(\bigcup_n A_n) = 0$. This fact is witnessed by simply taking any mixture of the underlying e-processes. To clarify, if E^n is the e-process from the above theorem which witnesses the fact that $\nu(A_n) = 0$, then for any positive sequence $\{w_n\}_{n \in \mathbb{N}}$ that sums to one, we see that $\sum_{n \in \mathbb{N}} w_n E^n$ is an e-process that witnesses $\nu(\bigcup_n A_n) = 0$. Indeed, $\sum_{n \in \mathbb{N}} w_n E^n$ is infinite if and only if at least one of its constituent e-processes is also infinite. (This remark is a direct parallel of Lemma 3.2 in [23], except the remark is stated more generally in terms of e-processes.)

Remark 3.6 (An unbounded sup suffices). We claim that if there exists a \mathcal{P} -e-process E such that $\sup_{t \geq 0} E_t = \infty$ on an event A , then there also exists a possibly different \mathcal{P} -e-process E' such that $\lim_{t \rightarrow \infty} E'_t = \infty$. The proof is simple. Let $\tau_n = \inf\{t : E_t \geq 2^n\}$ be the first time that E crosses level 2^n . Since E crosses every finite level on A , we know that $A \subseteq \{\tau_n < \infty\}$. Now define $E'_t = \sum_{n=1}^{\infty} \mathbf{1}_{\{\tau_n \leq t\}}$, exactly as in (3.1), and note that on

A , we have $\lim_{t \rightarrow \infty} E'_t = \infty$. Since E' is adapted, nondecreasing, and

$$\mathbb{E}_{\mathbb{P}} \left[\lim_{t \uparrow \infty} E'_t \right] = \mathbb{E}_{\mathbb{P}} \left[\sum_{n=1}^{\infty} \mathbf{1}_{\{\tau_n < \infty\}} \right] = \sum_{n=1}^{\infty} \mathbb{P}(\tau_n < \infty) \leq \sum_{n=1}^{\infty} \frac{1}{2^n} = 1$$

for all $\mathbb{P} \in \mathcal{P}$, E' is indeed an e-process. In summary, if one can exhibit an e-process E that is unbounded on A , then one can conclude that $\nu(A) = 0$. (This remark is a direct parallel of Lemma 3.1 in [23], except the remark is stated more generally in terms of e-processes.)

3.2 SubGaussian supermartingales and e-processes as witnesses

We now provide two simple examples where a nonnegative \mathcal{P} -supermartingale does bear witness to our main theorem, and a third example where a \mathcal{P} -e-process emerges.

We first recall a standard definition: a zero-mean random variable X is called σ -subGaussian, for some known $\sigma > 0$, if $\mathbb{E}[e^{\lambda X}] \leq e^{\lambda^2 \sigma^2 / 2}$ for every $\lambda \in \mathbb{R}$. A centered Gaussian with variance at most σ^2 is σ -subGaussian, but [8] proved that centered bounded random variables are also subGaussian (for an appropriate σ). Thus, the class of subGaussian random variables is very richly nonparametric, containing discrete and continuous distributions, both bounded and unbounded.

Case 1: zero mean. Now, define \mathcal{P}^0 to consist of all distributions over real-valued sequences $(X_t)_{t \in \mathbb{N}}$ such that for every $t \in \mathbb{N}$, X_t is mean-zero and σ -subGaussian when conditioned on X_1, \dots, X_{t-1} . Recalling (2.3), define the event

$$A_{\text{div}}^0 = \left\{ \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{s=1}^t X_s \text{ does not converge to zero} \right\}, \tag{3.2}$$

and note that $\mathbb{P}(A_{\text{div}}^0) = 0$ for every $\mathbb{P} \in \mathcal{P}^0$. We claim that $\nu(A_{\text{div}}^0) = 0$. It is possible to argue this directly, but we will instead do so by exhibiting a nonnegative \mathcal{P}^0 -supermartingale whose lim sup is unbounded on A_{div}^0 . We treat the case $\sigma = 1$ without loss of generality. First note that it is an immediate consequence of conditional subGaussianity that for any $\lambda \in \mathbb{R}$, the process

$$L_t(\lambda) = \exp \left(\lambda \sum_{i=1}^t X_i - \lambda^2 t / 2 \right) \tag{3.3}$$

is a nonnegative \mathcal{P}^0 -supermartingale with respect to the natural filtration generated by the data. By mixing this supermartingale over all $\lambda \in \mathbb{R}$ using a standard Gaussian distribution over λ , [20] derived the "normal mixture" \mathcal{P}^0 -supermartingale

$$M_t = \frac{\exp \left(\frac{t^2 \bar{X}_t^2}{2(t+1)} \right)}{\sqrt{t+1}},$$

where $\bar{X}_t = \sum_{i=1}^t X_i / t$ is the empirical mean.

Now, we claim that $(M_t)_{t \in \mathbb{N}_0}$ is unbounded on A_{div}^0 . Indeed, if \bar{X}_t does not converge to zero, then, along a subsequence, \bar{X}_t^2 remains above a strictly positive constant c . Along this subsequence, we have $M_t \geq \exp(t^2 c^2 / (2(t+1))) / \sqrt{t+1}$, which converges to infinity. By Remark 3.6, it suffices to treat $(M_t)_{t \in \mathbb{N}_0}$ as the witness \mathcal{P}^0 -supermartingale (and thus \mathcal{P}^0 -e-process) to conclude that $\nu(A_{\text{div}}^0) = 0$.

Case 2: nonpositive mean. Next, consider the event

$$A^{>0} = \left\{ \limsup_{t \rightarrow \infty} \frac{1}{t} \sum_{s=1}^t X_s > 0 \right\},$$

and the set of distributions

$$\mathcal{P}^{\leq 0} = \{ \mathbb{P} : a_t \leq 0 \text{ and } X_t - a_t \text{ is conditionally 1-subGaussian, for all } t \in \mathbb{N} \},$$

where $a_t = \mathbb{E}_{\mathbb{P}}[X_t \mid X_1, \dots, X_{t-1}]$ exists, is finite, and can itself be predictable.

In this case, the process $L_t(\lambda)$ from (3.3) is still a $\mathcal{P}^{\leq 0}$ -supermartingale, but only for $\lambda \geq 0$. Then, mixing over λ using only the positive half of a standard Gaussian, we get the "one-sided normal mixture" $\mathcal{P}^{\leq 0}$ -supermartingale, for which there is no closed form expression [20, 10]. Luckily, a suitable closed-form lower bound is given by

$$N_t = 2 \frac{\exp\left(\frac{2t^2 \bar{X}_t^2}{t+1}\right)}{\sqrt{t+1}} \Phi\left(\frac{2t \bar{X}_t}{\sqrt{t+1}}\right),$$

where Φ is the standard Gaussian cumulative distribution function. Since $(N_t)_{t \in \mathbb{N}_0}$ is not itself a supermartingale, but is upper bounded by one, it is a $\mathcal{P}^{\leq 0}$ -e-process. Now, a similar argument as before proves that $(N_t)_{t \in \mathbb{N}_0}$ is unbounded on $A^{>0}$, that is, $\limsup_{t \rightarrow \infty} N_t = \infty$ on any sequence of observations such that $\limsup_{t \rightarrow \infty} \bar{X}_t > 0$.

Case 3: nonpositive running mean. Finally, consider the class of distributions

$$\mathcal{P}^\dagger = \left\{ \mathbb{P} : \sum_{s \leq t} a_s \leq 0 \text{ and } X_t - a_t \text{ is conditionally 1-subGaussian, for all } t \in \mathbb{N} \right\},$$

where, as before, $a_t = \mathbb{E}_{\mathbb{P}}[X_t \mid X_1, \dots, X_{t-1}]$ exists, is finite, and can itself be predictable.

We have $\mathcal{P}^\dagger \supsetneq \mathcal{P}^{\leq 0}$ since the definition of \mathcal{P}^\dagger allows for a sequence of conditional means like $-1, -1, 2, 0, -2, 2, \dots$, whose individual elements could be positive as long as the cumulative sums are nonpositive. In this case, the process $L_t(\lambda)$ from (3.3) is no longer a \mathcal{P}^\dagger -supermartingale, but it is a \mathcal{P}^\dagger -e-process for all $\lambda \geq 0$. (To see this, simply note that for each $\mathbb{P} \in \mathcal{P}^\dagger$, the process $L_t(\lambda)$ is upper bounded by a different nonnegative \mathbb{P} -supermartingale, namely $\exp(\lambda \sum_{s=1}^t (X_s - a_s) - \lambda^2 t/2)$.) Thus N_t is also a \mathcal{P}^\dagger -e-process, and since we already know that it is unbounded on $A^{>0}$ we can conclude that $\nu(A^{>0}) = 0$ even for \mathcal{P}^\dagger .

We end this subsection by noting that there was nothing particularly special about the above subGaussian example. Howard et al. [9, 10] describe large classes of composite supermartingales and e-processes under different types of assumptions on the underlying random variable (like subexponential, sub-gamma, sub-Poisson, etc.), for which analogous sets of events and distributions could have been defined, and similar arguments could have been made.

3.3 E-process characterization for events with positive measure

One might wonder if there is a version of Theorem 3.1 that is valid for $\nu(A) > 0$. The following result and subsequent discussion and example show what is possible.

Theorem 3.7. *Let $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \in \mathbb{N}_0})$ be a filtered measurable space with $\mathcal{F} = \sigma(\bigcup_{t \in \mathbb{N}_0} \mathcal{F}_t)$, let \mathcal{P} be a family of probability measures on \mathcal{F} , let ν be the inverse capital outer measure from (2.1), and let $A \subseteq \Omega$ be arbitrary. Then, the following claims are true:*

1. For each $\varepsilon > 0$ there exists an e-process $(E_t)_{t \in \mathbb{N}_0}$ such that $\lim_{t \rightarrow \infty} E_t = 1/(\nu(A) + \varepsilon)$ on A .
2. For any $c \in [1, \infty]$, if there exists an e-process $(E_t)_{t \in \mathbb{N}_0}$ such that $\sup_{t \in \mathbb{N}_0} E_t \geq c$ on A , then $\nu(A) \leq 1/c$.

Proof. 1: Let τ_ε be a stopping time with $A \subseteq \{\tau_\varepsilon < \infty\}$ and $\sup_{\mathbb{P} \in \mathcal{P}} \mathbb{P}(\tau_\varepsilon < \infty) \leq \nu(A) + \varepsilon$, and define the one-jump process

$$E_t = \frac{1}{\nu(A) + \varepsilon} \mathbf{1}_{\{\tau_\varepsilon \leq t\}}, \quad t \in \mathbb{N}_0.$$

Then $E_\infty = \mathbf{1}_{\{\tau_\varepsilon < \infty\}}/(\nu(A) + \varepsilon)$, which equals $1/(\nu(A) + \varepsilon)$ on A . Moreover, $(E_t)_{t \in \mathbb{N}_0}$ is an e-process because for any stopping time τ and $\mathbb{P} \in \mathcal{P}$,

$$\mathbb{E}_{\mathbb{P}}[E_\tau] \leq \mathbb{E}_{\mathbb{P}}[E_\infty] = \frac{1}{\nu(A) + \varepsilon} \mathbb{P}(\tau_\varepsilon < \infty) \leq 1.$$

2: This is proved exactly as the converse implication of Theorem 3.1, except that the constant a is chosen from $(0, c)$. One obtains $\nu(A) \leq 1/a$ and hence, since $a \in (0, c)$ was arbitrary, $\nu(A) \leq 1/c$. \square

Remark 3.8. Let \mathcal{P} be the family of i.i.d. truncated Cauchy distributions constructed in Example 2.5. As shown there, $\nu(A_{\text{div}}) = 1$ which, in view of Theorem 3.72, makes it impossible to refute convergence with uniform power for this family \mathcal{P} . On the other hand, we also saw that $\nu(A_{\text{conv}}) = 1$ so refuting divergence is also impossible for \mathcal{P} .

3.4 A counterexample to a natural conjecture

The e-processes constructed in part 1 of Theorem 3.7 depend on ε , and it is natural to ask whether one can find a single e-process $(E_t)_{t \in \mathbb{N}_0}$ that satisfies

$$\sup_{t \in \mathbb{N}_0} E_t \geq \frac{1}{\nu(A)} \text{ on } A \tag{3.4}$$

when $\nu(A) \in (0, 1)$. In the singleton case $\mathcal{P} = \{\mathbb{P}\}$ this can be done, at least if A is measurable and \mathbb{P} -nullsets are ignored, using the Doob martingale $E_t = \mathbb{P}(A \mid \mathcal{F}_t)/\mathbb{P}(A)$. Indeed, Lévy's zero-one law states that $\lim_{t \rightarrow \infty} E_t = \mathbf{1}_A/\mathbb{P}(A)$ almost surely, and this equals $1/\nu(A)$ on A due to Lemma 2.13. However, if \mathcal{P} is not a singleton, it is in general impossible to find a single e-process that satisfies (3.4). The following example shows that this phenomenon can occur even if \mathcal{P} contains just two elements.

Example 3.9. Consider the space $\Omega = \{0, 1\}^{\mathbb{N}}$ of binary sequences (coin flips), let $X = (X_t)_{t \in \mathbb{N}}$ denote the canonical process, and set $\mathcal{F}_t = \sigma(X_s, s \leq t)$, $t \in \mathbb{N}$. We define two probability measures \mathbb{P}_1 and \mathbb{P}_2 as follows. The first coin flip is unbiased,

$$\mathbb{P}_i(X_1 = 1) = \frac{1}{2}, \quad i = 1, 2.$$

Conditionally on X_1 , the remaining coin flips X_2, X_3, \dots are independent under both \mathbb{P}_1 and \mathbb{P}_2 with success probabilities

$$\begin{aligned} \mathbb{P}_1(X_t = 1 \mid X_1 = 1) &= \mathbb{P}_2(X_t = 1 \mid X_1 = 0) = \frac{1}{2}; \\ \mathbb{P}_1(X_t = 1 \mid X_1 = 0) &= \mathbb{P}_2(X_t = 1 \mid X_1 = 1) = 2^{-t} \end{aligned}$$

for $t \geq 2$. Thus under \mathbb{P}_1 , if $X_1 = 1$ the remaining coin flips are i.i.d. Bernoulli- $(1/2)$, while if $X_1 = 0$ they are independent Bernoulli with exponentially decaying (hence summable) success probabilities. Under \mathbb{P}_2 the situation is symmetric.

Consider now the event

$$A = \{X_t = 1 \text{ for infinitely many } t \in \mathbb{N}\}.$$

The Borel–Cantelli lemma yields $\mathbb{P}_i(A) = 1/2$ for $i = 1, 2$. Moreover, if τ_n is the first time that n ones have been observed, then $A \subset \{\tau_n < \infty\}$ and $\mathbb{P}_i(\tau_n < \infty)$ converges to $1/2$ as n tends to infinity, for $i = 1, 2$. It follows that $\nu(A) = 1/2$.

However, we claim that no e-process for $\mathcal{P} = \{\mathbb{P}_1, \mathbb{P}_2\}$ can satisfy (3.4). To show this, consider any e-process $(E_t)_{t \in \mathbb{N}_0}$ and suppose that $\sup_{t \in \mathbb{N}_0} E_t \geq c$ on A for some constant $c \in [1, \infty)$. We must show that $c < 2$. To do so, fix $T \in \mathbb{N}$ and $\varepsilon \in (0, c/2)$ such that $\{E_T \geq \varepsilon\} \neq \emptyset$. Define $B = \{E_T \geq \varepsilon\} \cap \{X_1 = 0\}$ and consider the case $B \neq \emptyset$ (the complementary case $\{E_T \geq \varepsilon\} \cap \{X_1 = 1\} \neq \emptyset$ is argued symmetrically). Denote $\delta = \mathbb{P}_1(B) \in (0, 1/2]$. (Indeed $\delta > 0$ since $B \in \mathcal{F}_T$ and $B \neq \emptyset$ by assumption.) Next, define the stopping time

$$\tau = \inf\{t: E_t \geq c - \delta\varepsilon\} \wedge T_B,$$

where $T_B = T\mathbf{1}_B + \infty\mathbf{1}_{B^c}$ (meaning T_B takes the value infinity on the complement of B). Note that both A and B are subsets of $\{\tau < \infty\}$. Moreover, the Borel–Cantelli lemma yields $A = \{X_1 = 1\}$ up to a \mathbb{P}_1 -nullset, and hence $\mathbb{P}_1(A \cap B) = 0$. We conclude that

$$\mathbb{E}_{\mathbb{P}_1}[E_\tau \mathbf{1}_{\{\tau < \infty\}}] \geq \mathbb{E}_{\mathbb{P}_1}[E_\tau \mathbf{1}_A] + \mathbb{E}_{\mathbb{P}_1}[E_\tau \mathbf{1}_B] \geq \frac{1}{2}(c - \delta\varepsilon) + \delta \min\{c - \delta\varepsilon, \varepsilon\} > \frac{c}{2}.$$

On the other hand, the left-hand side is bounded by one since $(E_t)_{t \in \mathbb{N}_0}$ is an e-process. This shows that $c < 2$, as required.

3.5 Gambling interpretations of the inverse capital measure

Nonnegative martingales have a straightforward game-theoretic interpretation: they are the wealth of a gambler playing a fair game. Nonnegative supermartingales are similar: they are the wealth of a gambler playing a game whose odds are stacked against them (or playing a fair game, and throwing away some money in each round). \mathcal{P} -e-processes admit a slightly more sophisticated game-theoretic interpretation: at the very beginning, the gambler agrees to play (in parallel) a separate fair game against each $P \in \mathcal{P}$ (or one with odds stacked against them), and their wealth at any time is given by their minimum wealth across all these games. For more on this interpretation, see Section 5.4 in Ramdas et al. [18]. Thus it is natural to ask: does ν also admit a game-theoretic interpretation?

In this subsection, we rewrite Theorem 3.7 to bring to the fore its pricing interpretation of ν . The following corollary is a composite analog of Proposition 8.13 in Shafer and Vovk [23]. It shows that $1/\nu(A)$ is the maximum level that can be reached by an e-process started at 1, whenever the event A occurs.

Corollary 3.10. *Let \mathcal{P} be a family of probability measures on Ω and let ν be the associated inverse capital measure (2.1). Then*

$$\nu(A) = \inf \left\{ c \in (0, 1] : \liminf_{t \rightarrow \infty} cE_t \geq \mathbf{1}_A \right\} = \inf \left\{ c \in (0, 1] : \sup_{t \in \mathbb{N}_0} cE_t \geq \mathbf{1}_A \right\}, \quad (3.5)$$

where $(E_t)_{t \in \mathbb{N}_0}$ ranges over the set of all e-processes for \mathcal{P} .

Proof. Let $V_{\lim \inf}$ and V_{\sup} denote the value of the first and second infimum respectively in (3.5). As $\liminf_{t \rightarrow \infty} cE_t \leq \sup_{t \in \mathbb{N}_0} cE_t$, the inf defining V_{\sup} is over a larger set, and hence $V_{\sup} \leq V_{\lim \inf}$. To then prove the claimed equalities (3.5), it suffices to show the sandwich $V_{\lim \inf} \leq \nu(A) \leq V_{\sup}$. In order to prove the first inequality, Theorem 3.7(1) yields an e-process showing that $c = \nu(A) + \epsilon$ is feasible for $V_{\lim \inf}$ for any $\epsilon > 0$,

which implies the first inequality by letting $\epsilon \rightarrow 0$. Moving on to the second inequality, Theorem 3.7(2) shows that for V_{sup} every feasible pair of c and e-process must have $c \geq \nu(A)$, proving the second inequality. This completes the proof. \square

The following rewrite brings out yet another aspect related to betting. For a process $(M_t)_{t \in \mathbb{N}_0}$, thought of as payoffs, we can say it has fair price $\sup_{\rho \in \mathcal{T}} \mathbb{E}_{\mathbb{P}}[M_{\rho}]$ under \mathbb{P} . The following result shows that the inverse capital measure of an event A is also the infimum of all initial capitals that suffice to super-replicate the payoff $\mathbf{1}_A$. Due to this interpretation we could also have plausibly called the inverse capital measure the “initial capital measure”.

Lemma 3.11. *The following dual representation for ν holds:*

$$\nu(A) = \inf_{\tau \in \mathcal{T}: A \subseteq \{\tau < \infty\}} \sup_{\mathbb{P} \in \mathcal{P}} \mathbb{P}(\tau < \infty) = \inf_{\substack{(M_t)_t \geq 0 \text{ s.t.} \\ \sup_{t \in \mathbb{N}_0} M_t \geq \mathbf{1}_A}} \sup_{\mathbb{P} \in \mathcal{P}} \sup_{\rho \in \mathcal{T}} \mathbb{E}_{\mathbb{P}}[M_{\rho}]. \quad (3.6)$$

Proof. The first expression is simply the definition of $\nu(A)$, and is included for convenience of comparison to the second expression. We prove both directions in turn.

- For \geq , fix any feasible τ on the left. Then $M_t = \mathbf{1}_{\tau < \infty}$ is feasible on the right, by monotonicity in time it is best to take $\rho = \infty$, upon which $\mathbb{E}_{\mathbb{P}}[M_{\rho}] = \mathbb{P}(\tau < \infty)$.
- For \leq , fix any feasible $(M_t)_t$ on the right. For any $\epsilon \in (0, 1)$, consider the stopping time $\tau = \inf\{t : M_t \geq 1 - \epsilon\}$, so that $A \subseteq \{\tau < \infty\}$ and hence τ is feasible on the left. Moreover, $\mathbb{P}(\tau < \infty)(1 - \epsilon) \leq \mathbb{E}_{\mathbb{P}}[M_{\tau}] \leq \sup_{\rho} \mathbb{E}_{\mathbb{P}}[M_{\rho}]$. Taking $\epsilon \rightarrow 0$ proves the claim.

This completes the proof. \square

4 Line-crossing inequalities and a composite SLLN

Before we introduce our composite strong law of large numbers (SLLN), we present an important technical tool that may be of independent interest.

4.1 An L^1 -type line-crossing inequality

Recall the following famous inequality due to Dubins and Savage [5]: for any square integrable martingale $(M_t)_{t \in \mathbb{N}}$ with $M_0 = 0$, having predictable quadratic variation $V_t = \sum_{s=1}^t \mathbb{E}[(M_s - M_{s-1})^2 | \mathcal{F}_{s-1}]$, we have for any constants $\gamma, \epsilon > 0$,

$$\mathbb{P}\left(\sup_{t \in \mathbb{N}} \frac{M_t}{\gamma + V_t} \geq \epsilon\right) \leq \frac{1}{\gamma\epsilon^2 + 1}.$$

The following theorem can be understood as an L^1 -type Dubins–Savage inequality. Namely, it gives a probability bound on the likelihood that a random walk starting at zero with i.i.d. centered increments (that are only required to be in L^1), ever crosses a line with positive slope and intercept. To the best of our knowledge, such inequalities have only been proven in the presence of higher moments for the random walk increments; see, for example, Khan [11] for L^p variants for $1 < p \leq 2$, and Howard et al. [9] for other variants under two, three, and infinite moments. It might be of interest to improve the constants in the statement, although this is not needed for our purposes.

Theorem 4.1. *Suppose that $(X_t)_{t \in \mathbb{N}}$ is i.i.d. with $\mathbb{E}[|X_1|] < \infty$ and $\mathbb{E}[X_1] = 0$ and define the tail function*

$$r(K) = \mathbb{E}\left[|X_1| \mathbf{1}_{\{|X_1| > K\}}\right], \quad K \geq 0. \quad (4.1)$$

Then for any positive numbers ε, γ , and any $K \geq 1$ one has

$$\mathbb{P} \left(\sup_{t \in \mathbb{N}} \frac{1}{\gamma + t} \left| \sum_{s=1}^t X_s \right| > \varepsilon + r(K) \right) \leq \frac{8K^2}{\gamma\varepsilon^2} + \left(\frac{16}{\varepsilon^2} + 2 \right) r(K).$$

Proof. We use the triangle inequality to bound

$$\frac{1}{\gamma + t} \left| \sum_{s=1}^t X_s \right| \leq \frac{1}{\gamma + t} \left| \sum_{s=1}^t Y_s \right| + \frac{1}{t} \left| \sum_{s=1}^t Z_s \right| + \frac{1}{t} \left| \sum_{s=1}^t R_s \right|, \tag{4.2}$$

where

$$\begin{aligned} Y_s &= X_s \mathbf{1}_{\{|X_s| \leq K\}} - \mathbb{E}[X_s \mathbf{1}_{\{|X_s| \leq K\}}], \\ Z_s &= X_s \mathbf{1}_{\{K < |X_s| \leq s-1\}} - \mathbb{E}[X_s \mathbf{1}_{\{K < |X_s| \leq s-1\}}], \\ R_s &= X_s \mathbf{1}_{\{|X_s| > K \vee (s-1)\}} - \mathbb{E}[X_s \mathbf{1}_{\{|X_s| > K \vee (s-1)\}}]. \end{aligned}$$

We handle the three terms on the right-hand side of (4.2) one by one.

Step 1: Since the Y_s are centered i.i.d. with variance bounded by K^2 , the classical Dubins–Savage line-crossing inequality [5] implies that

$$\mathbb{P} \left(\exists t \in \mathbb{N} \text{ with } \sum_{s=1}^t Y_s > \beta + \alpha K^2 t \right) \leq \frac{1}{1 + \alpha\beta} \leq \frac{1}{\alpha\beta} \tag{4.3}$$

for any positive numbers α, β . We apply this with $\alpha = \varepsilon/K^2$ and $\beta = \varepsilon\gamma$ to get the following tail bound for the first term on the right-hand side of (4.2):

$$\mathbb{P} \left(\sup_{t \in \mathbb{N}} \frac{1}{\gamma + t} \left| \sum_{s=1}^t Y_s \right| > \varepsilon \right) \leq 2 \mathbb{P} \left(\exists t \in \mathbb{N} \text{ with } \sum_{s=1}^t Y_s > \varepsilon\gamma + \varepsilon t \right) \leq \frac{2K^2}{\gamma\varepsilon^2}. \tag{4.4}$$

Step 2: Summation by parts yields the identity

$$\frac{1}{t} \sum_{s=1}^t Z_s = \sum_{s=1}^t \frac{Z_s}{s} - \frac{1}{t} \sum_{j=2}^t \sum_{s=1}^{j-1} \frac{Z_s}{s}. \tag{4.5}$$

Writing $S = \sup_{t \in \mathbb{N}} \left| \sum_{s=1}^t Z_s/s \right|$ we deduce $\left| \sum_{s=1}^t Z_s \right|/t \leq S + \sum_{j=1}^t S/t = 2S$ for all $t \in \mathbb{N}$. Noting that $\sum_{s=1}^t Z_s/s$ is a martingale with independent increments, we may use Kolmogorov’s maximal inequality and monotone convergence to get

$$\mathbb{P} \left(\sup_{t \in \mathbb{N}} \frac{1}{t} \left| \sum_{s=1}^t Z_s \right| > \varepsilon \right) \leq \mathbb{P} \left(\sup_{t \in \mathbb{N}} \left| \sum_{s=1}^t \frac{Z_s}{s} \right| > \frac{\varepsilon}{2} \right) \leq \frac{4}{\varepsilon^2} \mathbb{E} \left[\sum_{s=1}^{\infty} \frac{Z_s^2}{s^2} \right]. \tag{4.6}$$

Since $\mathbb{E}[Z_s^2] \leq \mathbb{E}[X_1^2 \mathbf{1}_{\{K < |X_1| \leq s-1\}}]$, the expectation on the right-hand side can be bounded as follows:

$$\mathbb{E} \left[\sum_{s=1}^{\infty} \frac{Z_s^2}{s^2} \right] \leq \mathbb{E} \left[X_1^2 \mathbf{1}_{\{|X_1| > K\}} \sum_{s=1}^{\infty} \mathbf{1}_{\{|X_1| \leq s-1\}} \frac{1}{s^2} \right] \leq \mathbb{E}[|X_1| \mathbf{1}_{\{|X_1| > K\}}],$$

using in the last step the inequality $\sum_{s=1}^{\infty} \mathbf{1}_{\{x \leq s-1\}} 1/s^2 \leq \int_x^{\infty} 1/s^2 ds = 1/x, x \in (0, \infty)$. This establishes the tail bound

$$\mathbb{P} \left(\sup_{t \in \mathbb{N}} \frac{1}{t} \left| \sum_{i=1}^t Z_s \right| > \varepsilon \right) \leq \frac{4}{\varepsilon^2} r(K), \tag{4.7}$$

where $r(K)$ is the tail function (4.1).

Step 3: Write $R_s = X_s \mathbf{1}_{\{|X_s| > K \vee (s-1)\}} - c_s$ with $c_s = \mathbb{E}[X_s \mathbf{1}_{\{|X_s| > K \vee (s-1)\}}]$. On the event where $|X_s| \mathbf{1}_{\{|X_s| > K\}} \leq s - 1$ for all $s \in \mathbb{N}$, one has $R_s = -c_s$ for all $s \in \mathbb{N}$, and hence

$$\frac{1}{t} \left| \sum_{s=1}^t R_s \right| = \frac{1}{t} \left| \sum_{s=1}^t c_s \right| \leq \frac{1}{t} \sum_{s=1}^t \mathbb{E}[|X_s| \mathbf{1}_{\{|X_s| > K\}}] = \mathbb{E}[|X_1| \mathbf{1}_{\{|X_1| > K\}}] = r(K)$$

for all $t \in \mathbb{N}$. Therefore,

$$\mathbb{P}(|X_s| \mathbf{1}_{\{|X_s| > K\}} \leq s - 1 \text{ for all } s \in \mathbb{N}) \leq \mathbb{P}\left(\sup_{t \in \mathbb{N}} \frac{1}{t} \left| \sum_{s=1}^t R_s \right| \leq r(K)\right).$$

We now take complements, apply the union bound, and use that $\sum_{n \in \mathbb{N}} \mathbb{P}(\xi > n - 1) \leq \mathbb{P}(\xi > 0) + \mathbb{E}[\xi]$ for any nonnegative random variable ξ to get

$$\begin{aligned} \mathbb{P}\left(\sup_{t \in \mathbb{N}} \frac{1}{t} \left| \sum_{s=1}^t R_s \right| > r(K)\right) &\leq \sum_{s \in \mathbb{N}} \mathbb{P}(|X_1| \mathbf{1}_{\{|X_1| > K\}} > s - 1) \\ &\leq \mathbb{P}(|X_1| > K) + \mathbb{E}[|X_1| \mathbf{1}_{\{|X_1| > K\}}] \\ &\leq 2\mathbb{E}[|X_1| \mathbf{1}_{\{|X_1| > K\}}] = 2r(K). \end{aligned} \tag{4.8}$$

Here the third inequality follows because $K \geq 1$.

Step 4: Combining (4.2), (4.4), (4.7), and (4.8) yields

$$\mathbb{P}\left(\sup_{t \in \mathbb{N}} \frac{1}{\gamma + t} \left| \sum_{s=1}^t X_s \right| > 2\varepsilon + r(K)\right) \leq \frac{2K^2}{\gamma \varepsilon^2} + \left(\frac{4}{\varepsilon^2} + 2\right) r(K).$$

Replacing ε by $\varepsilon/2$ completes the proof. □

By choosing $K = \gamma^{1/3}$ one obtains the following corollary. In particular, this makes it clear that the left-hand side converges to 0 as γ tends to ∞ , and that this occurs at a rate that only depends on the tail function $r(K)$. Recall that $r(K)$ decreases to 0 as K tends to ∞ .

Corollary 4.2. *In the setting of Theorem 4.1 one has*

$$\mathbb{P}\left(\sup_{t \in \mathbb{N}} \frac{1}{\gamma + t} \left| \sum_{s=1}^t X_s \right| > 2\varepsilon\right) \leq \frac{8}{\varepsilon^2} \gamma^{-1/3} + \left(\frac{16}{\varepsilon^2} + 2\right) r(\gamma^{1/3}) \tag{4.9}$$

for any positive $\varepsilon > 0$ and any large enough $\gamma \geq 0$ such that $r(\gamma^{1/3}) \leq \varepsilon$. Since

$$\frac{1}{k} \left| \sum_{s=1}^k X_s \right| \leq 2 \sup_{t \in \mathbb{N}} \frac{1}{k + t} \left| \sum_{s=1}^t X_s \right|,$$

we further conclude that

$$\mathbb{P}\left(\frac{1}{k} \left| \sum_{s=1}^k X_s \right| > \varepsilon\right) \leq \frac{128}{\varepsilon^2} k^{-1/3} + \left(\frac{256}{\varepsilon^2} + 2\right) r(k^{1/3}) \tag{4.10}$$

for any positive ε and any $k \in \mathbb{N}$ such that $r(k^{1/3}) \leq \varepsilon$.

4.2 A composite strong law of large numbers

The line-crossing inequality of Theorem 4.1 implies a composite strong law of large numbers. Specifically, if ν is induced by a family \mathcal{P} of i.i.d. laws that satisfy a certain uniform integrability condition, then the event that the time averages of $(X_t)_{t \in \mathbb{N}}$ diverge,

$$A_{\text{div}} = \left\{ \text{the limit } \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{i=1}^t X_i \text{ does not exist} \right\},$$

will be a ν -nullset.

Theorem 4.3. *Let \mathcal{P} be a family of probability measures \mathbb{P} under which $(X_t)_{t \in \mathbb{N}}$ is i.i.d. with $\mathbb{E}_{\mathbb{P}}[|X_1|] < \infty$. Assume also that \mathcal{P} satisfies the centered uniform integrability condition*

$$\lim_{K \rightarrow \infty} \sup_{\mathbb{P} \in \mathcal{P}} \mathbb{E}_{\mathbb{P}} \left[|X_1 - \mathbb{E}_{\mathbb{P}}[X_1]| \mathbf{1}_{\{|X_1 - \mathbb{E}_{\mathbb{P}}[X_1]| > K\}} \right] = 0. \tag{4.11}$$

Then $\nu(A_{\text{div}}) = 0$.

Let us contrast this composite SLLN result with classical SLLNs. The latter assume some condition (e.g. i.i.d. with finite variance) and prove that A_{div} has measure zero. That is, they show that for any set of distributions \mathcal{P} satisfying the assumed condition, the maximum likelihood measure is zero, that is $\mu(A_{\text{div}}) = 0$. This in turn implies that $\mathbb{P}(A_{\text{div}}) = 0$ for every $\mathbb{P} \in \mathcal{P}$, providing a strong argument for why A_{div} will not happen. Our composite SLLN instead strengthens this to $\nu(A_{\text{div}}) = 0$. This implies that there is a \mathcal{P} -e-process $(E_t)_{t \in \mathbb{N}_0}$ that explodes when A_{div} occurs. Equivalently, for every confidence level $\delta \in (0, 1)$, there is a stopping time τ that detects (by stopping at some finite time) that A_{div} is occurring, with a type-I error bounded by δ for all $\mathbb{P} \in \mathcal{P}$. In words, $\mu(A_{\text{div}}) = 0$ means that convergence occurs, while $\nu(A_{\text{div}}) = 0$ means that we can in addition detect divergence in finite time.

An example of \mathcal{P} that satisfies condition (4.11) is a location family created by translating a distribution with finite mean. To elaborate, let $\phi(x)$ be a density such that $\int |x| \phi(x) dx < \infty$. Abusing notation to use $\phi(x)$ to also denote the distribution with that density, define $\mathcal{P}^\phi = \{\phi(x - a)\}_{a \in \mathbb{R}}$ as distributions yielding i.i.d. draws from some shifted version of ϕ . Then \mathcal{P}^ϕ satisfies (4.11). In fact, more can be said. For any given set of i.i.d. product distributions \mathcal{P} , let \mathcal{P}^0 be the family of distributions obtained by centering each of the distributions in \mathcal{P} , so that every $\mathbb{P} \in \mathcal{P}^0$ yields zero-mean i.i.d. draws. Then \mathcal{P} satisfies (4.11) if and only if \mathcal{P}^0 satisfies the more standard uniform integrability condition, which we know (from the de la Vallée-Poussin theorem) holds if and only if there exists a nonnegative increasing function G such that $\lim_{x \rightarrow \infty} G(x)/x = \infty$ and $\sup_{\mathbb{P} \in \mathcal{P}^0} \mathbb{E}[G(|X_1|)] < \infty$.

An example of \mathcal{P} that fails condition (4.11) is the set of all truncated zero-centered Cauchy distributions encountered in Example 2.5. This example, and especially Proposition 2.6, reminds us that if uniform integrability conditions are dropped entirely, then $\nu(A_{\text{div}})$ will not be zero, but could even be equal one. Thus, condition (4.11) is not only sufficient for concluding $\nu(A_{\text{div}}) = 0$, but also that some condition like (4.11) is necessary to restrict \mathcal{P} . However, we leave the exact characterization of the necessary condition for future work.

Proof of Theorem 4.3. We use the notation $\bar{X}_t = \sum_{s=1}^t X_s/t$. For each $n \in \mathbb{N}$, define the event

$$A_n = \left\{ \limsup_{t \rightarrow \infty} \bar{X}_t - \liminf_{t \rightarrow \infty} \bar{X}_t > \frac{2}{n} \right\}. \tag{4.12}$$

Then $A_{\text{div}} \subseteq \bigcup_{n \in \mathbb{N}} A_n$, so we wish to bound $\nu(A_n)$. To this end, observe that for any $k, n \in \mathbb{N}$ one has $A_n \subseteq \{\tau_{k,n} < \infty\}$, where

$$\tau_{k,n} = \inf \left\{ t \geq 0 : |\bar{X}_{k+t} - \bar{X}_k| > \frac{1}{n} \right\}. \tag{4.13}$$

(Indeed, if $\tau_{k,n} = \infty$, \bar{X}_{k+t} stays within $1/n$ of \bar{X}_k for all $t \geq 0$. But then so do both $\limsup_{t \rightarrow \infty} \bar{X}_t$ and $\liminf_{t \rightarrow \infty} \bar{X}_t$. This shows that $\{\tau_{k,n} = \infty\} \subseteq A_n^c$.) We conclude that

$$\nu(A_n) \leq \sup_{\mathbb{P} \in \mathcal{P}} \mathbb{P}(\tau_{k,n} < \infty).$$

It is shown below that the right-hand side tends to 0 as k tends to ∞ , for each $n \in \mathbb{N}$. Thus $\nu(A_n) = 0$ for all $n \in \mathbb{N}$, and hence $\nu(A_{\text{div}}) \leq \sum_{n \in \mathbb{N}} \nu(A_n) = 0$ as required.

It remains to show, for any fixed $n \in \mathbb{N}$, that $\sup_{\mathbb{P} \in \mathcal{P}} \mathbb{P}(\tau_{k,n} < \infty)$ tends to 0 as k tends to ∞ . Observe that for any $k \in \mathbb{N}$ and $c \in \mathbb{R}$,

$$|\bar{X}_{k+t} - \bar{X}_k| \leq \frac{1}{k+t} \left| \sum_{i=1}^t (X_{k+i} - c) \right| + 2|\bar{X}_k - c|. \tag{4.14}$$

Thus for any $\mathbb{P} \in \mathcal{P}$ and $k \in \mathbb{N}$ we may take $c = \mathbb{E}_{\mathbb{P}}[X_1]$ to obtain

$$\mathbb{P}(\tau_{k,n} < \infty) \leq \mathbb{P} \left(\sup_{t \in \mathbb{N}} \frac{1}{k+t} \left| \sum_{i=1}^t (X_{k+i} - \mathbb{E}_{\mathbb{P}}[X_1]) \right| \geq \frac{1}{2n} \right) + \mathbb{P} \left(|\bar{X}_k - \mathbb{E}_{\mathbb{P}}[X_1]| \geq \frac{1}{4n} \right).$$

The line-crossing inequality (4.9) and the inequality (4.10) in Corollary 4.2 now show that for k large enough the right-hand side is bounded by

$$C \left(k^{-1/3} + \mathbb{E}_{\mathbb{P}} \left[|X_1 - \mathbb{E}_{\mathbb{P}}[X_1]| \mathbf{1}_{\{|X_1 - \mathbb{E}_{\mathbb{P}}[X_1]| > k^{1/3}\}} \right] \right)$$

for a constant C which, for instance, can be taken to be $C = 4352n^2 + 4$. Thanks to the assumption (4.11), it follows that $\sup_{\mathbb{P} \in \mathcal{P}} \mathbb{P}(\tau_{k,n} < \infty)$ tends to 0 as k tends to ∞ . This completes the proof of the theorem. \square

Remark 4.4. With \mathcal{P} as in Theorem 4.3, this result together with the composite Ville's theorem, Theorem 3.1, shows that there exists an e-process for \mathcal{P} which converges to infinity on A_{div} . An inspection of the proofs shows that this e-process can be chosen as

$$E_t = \sum_{n \in \mathbb{N}} \mathbf{1}_{\{\tau_{k_n,n} \leq t\}}, \quad t \in \mathbb{N}_0,$$

where $\tau_{k_n,n}$ is given by (4.13) with $k = k_n$, and $k_n \gg n^6 2^{3n}$ is chosen large enough that

$$(4352n^2 + 4) \left(k_n^{-1/3} + \sup_{\mathbb{P} \in \mathcal{P}} \mathbb{E}_{\mathbb{P}} \left[|X_1 - \mathbb{E}_{\mathbb{P}}[X_1]| \mathbf{1}_{\{|X_1 - \mathbb{E}_{\mathbb{P}}[X_1]| > k_n^{1/3}\}} \right] \right) \leq 2^{-n}.$$

This yields a game-theoretic composite SLLN for unbounded observations without requiring any moment beyond the first. It may be contrasted with other game-theoretic SLLNs, for example by [23] and [13], which are pointwise claims and hold for bounded observations.

Remark 4.5. Chung [2] proved a uniform strong law of large numbers that can be used instead of our line-crossing inequality to prove Theorem 4.3. Here, we summarize the statement from Shorack's textbook [25, Theorem 10.12.1], which also has a self-contained proof: suppose $(X_t)_{t \in \mathbb{N}}$ is i.i.d. for each $\mathbb{P} \in \mathcal{P}$, where \mathcal{P} satisfies the uniform integrability condition

$$\lim_{K \rightarrow \infty} \sup_{\mathbb{P} \in \mathcal{P}} \mathbb{E}_{\mathbb{P}}[|X_1| \mathbf{1}_{\{|X_1| \geq K\}}] = 0. \tag{4.15}$$

Then, for any $\varepsilon > 0$,

$$\lim_{n \rightarrow \infty} \sup_{\mathbb{P} \in \mathcal{P}} \mathbb{P} \left(\sup_{m \geq n} |\bar{X}_m - \mathbb{E}_{\mathbb{P}}[X_1]| \geq \varepsilon \right) = 0. \tag{4.16}$$

Interestingly, our slightly weaker assumption (4.11) is already sufficient to yield Chung's conclusion (4.16): given a family \mathcal{P} satisfying (4.11), define $\mathcal{P}' = \{\mathbb{P} - \mathbb{E}_{\mathbb{P}}[X_1] : \mathbb{P} \in \mathcal{P}\}$, where we write " $\mathbb{P} - \mathbb{E}_{\mathbb{P}}[X_1]$ " as shorthand for the law under \mathbb{P} of the centered sequence $(X_t - \mathbb{E}_{\mathbb{P}}[X_1])_{t \in \mathbb{N}}$. Then \mathcal{P}' satisfies (4.15), and applying Chung's result to \mathcal{P}' yields (4.16), first with \mathcal{P}' in place of \mathcal{P} , but then also for \mathcal{P} itself. Once this has been done, a proof of Theorem 4.3 can be obtained as follows. Up until (4.14) the proof is unchanged, but then, instead of (4.14), one uses the triangle inequality to get

$$|\bar{X}_{k+t} - \bar{X}_k| \leq |\bar{X}_{k+t} - c| + |\bar{X}_k - c|$$

for any constant c . For each $\mathbb{P} \in \mathcal{P}$ one then takes $c = \mathbb{E}_{\mathbb{P}}[X_1]$ to get

$$\mathbb{P}(\tau_{k,n} < \infty) \leq 2 \mathbb{P} \left(\sup_{t \in \mathbb{N}} |\bar{X}_{k+t} - \mathbb{E}_{\mathbb{P}}[X_1]| \geq \frac{1}{2n} \right).$$

Applying (4.16) now yields $\lim_{k \rightarrow \infty} \sup_{\mathbb{P} \in \mathcal{P}} \mathbb{P}(\tau_{k,n} < \infty) = 0$, as required.

4.3 The insufficiency of test (super)martingales for the composite strong law

Consider any \mathcal{P} that satisfies the conditions of Theorem 4.3: it contains distributions having a finite mean and satisfying the uniform integrability condition (4.11). The composite strong law states that $\nu(A_{\text{div}}) = 0$, meaning that the "erratic" sequences — those on which the empirical averages do not converge to a limit — form a ν -nullset. Theorem 3.7 in turn constructs an explicit e-process $(E_t)_{t \in \mathbb{N}_0}$ — a nonnegative process such that $\mathbb{E}_{\mathbb{P}}[E_\tau] \leq 1$ for any $\mathbb{P} \in \mathcal{P}$ and stopping time τ — that grows to infinity on any erratic sequence. Further, $(\mathbf{1}_{\{E_t > 1/\alpha\}})_{t \in \mathbb{N}_0}$ yields a level- α sequential test for the null hypothesis that the empirical average of the observations converges to some limit.

One important question is whether an e-process was necessary for this purpose, or if a nonnegative martingale (with respect to \mathcal{P}) beginning at one would have sufficed. In Ville's original result, \mathcal{P} was a singleton, and hence such martingales do exist and were explicitly constructed. However, in our setting, the question is whether a composite \mathcal{P} -martingale exists, which would satisfy the martingale property for every $\mathbb{P} \in \mathcal{P}$.

The answer is an emphatic no: for any sufficiently rich \mathcal{P} , there is no nonnegative \mathcal{P} -martingale that can grow to infinity on ν -nullsets. In fact, it is possible that the only nonnegative \mathcal{P} -martingales are constants (and if we fix their initial value to one, then it must simply remain at one). This fact is not obvious. Ramdas et al. [18] prove that for $\mathcal{P}^{01} = \{\text{Ber}(p)^\infty : p \in [0, 1]\}$ on the space $\Omega = \{0, 1\}^\mathbb{N}$, meaning the set of all possible i.i.d. Bernoulli sequences, the only nonnegative \mathcal{P}^{01} -martingales are constants.

If (for example) $\Omega = \mathbb{R}^\mathbb{N}$, then as long as $\mathcal{P}^{01} \subset \mathcal{P}$, there can still be no nonnegative \mathcal{P} -martingale that can grow to infinity on A_{div} . Indeed, A_{div} contains many binary sequences; on these, no \mathcal{P} -martingale can grow, since no \mathcal{P}^{01} -martingale can grow.

If the constraint of being a martingale for every $\mathbb{P} \in \mathcal{P}$ is too hard to meet, one could ask if it can be relaxed to a nonnegative \mathcal{P} -supermartingale. The answer again is no. The aforementioned paper also proves that the only nonnegative \mathcal{P}^{01} -supermartingales are decreasing sequences, and thus the same conclusion holds in our setting.

However, in that paper, the authors show that it is possible to construct a nontrivial e-process that tests \mathcal{P}^{01} (equivalently, since the e-process property is unaffected by taking convex combinations of distributions, it tests whether the distribution of the data

is exchangeable), and that such an e-process is not condemned to stay small: it increases to infinity under a variety of alternatives, such as Markovian ones.

Unlike nonnegative \mathcal{P} -(super)martingales, the \mathcal{P} -e-process constructed by Theorem 3.7 is also not condemned to stay small, it increases to infinity on any erratic sequence. The above paragraphs demonstrate that e-processes are not only sufficient, but also (in some sense) necessary, in the sense that existing standard martingale concepts are not satisfactory in such richly composite settings.

5 Summary

This paper derived a composite generalization of Ville's fundamental martingale theorem that relates measure-theoretic and game-theoretic probability for single distributions. This required the development of an inverse price outer measure of any event, which is the appropriate generalization of the probability of that event in the singleton case. It also required the concept of an e-process, which generalizes an optional stopping property of martingales and supermartingales, and appropriately extends those fundamental concepts to the composite setting. Then, we proved that for any event of inverse price measure zero, there exists an e-process which increases to infinity if that event occurs. Said differently, our results give a rigorous way to find a fair price for indicators of events in a composite setting.

There exist composite sets \mathcal{P} for which a nonnegative (super)martingale could have sufficed, and e-processes are not required. Examining interesting special classes \mathcal{P} , and understanding when supermartingales do or do not suffice, has been the subject of much intense research recently [9, 10, 7, 17, 18, 12]. In this paper, we show that the composite strong law of large numbers provided a setting where e-processes were necessary and sufficient. Future work could extend game-theoretic constructions for the law of the iterated logarithm [15, 21] to composite settings.

We end with a quote by Ville himself in an interview from the 1980s [3], reflecting on his seminal work:

My goal, given the X_n and a set to which the sequence of X_n belongs with probability zero, was to define (*martingales*) s_n so that they tend to infinity on that set. I insist on the point because it took me so long to make this way of proceeding understood.

Our work extends Ville's research program to the realm of composite, nonparametric sets of distributions, and expands the deep connections between measure-theoretic and game-theoretic probability, as developed by [23, 24].

References

- [1] Patrick Billingsley, *Probability and Measure*, Wiley Series in Probability and Mathematical Statistics, John Wiley & Sons, New York-Chichester-Brisbane, 1979. MR534323
- [2] Kai Lai Chung, *The strong law of large numbers*, Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability, 1950, University of California Press, Berkeley-Los Angeles, Calif., 1951, pp. 341–352. MR0045328
- [3] Pierre Crépel, *Jean Ville's recollections, in 1984 and 1985, concerning his work on martingales*, Electronic Journal for History of Probability and Statistics **5** (2009), no. 1.
- [4] J. L. Doob, *Regularity properties of certain families of chance variables*, Trans. Amer. Math. Soc. **47** (1940), 455–486. MR2052
- [5] Lester E. Dubins and Leonard J. Savage, *A Tchebycheff-like inequality for stochastic processes*, Proc. Nat. Acad. Sci. U.S.A. **53** (1965), 274–275. MR182042

- [6] Hans Föllmer and Alexander Schied, *Stochastic finance*, De Gruyter Graduate, De Gruyter, Berlin, 2016, An introduction in discrete time, Fourth revised and extended edition of [MR1925197]. MR3859905
- [7] Peter Grünwald, Rianne de Heide, and Wouter M. Koolen, *Safe testing*, Journal of the Royal Statistical Society, Series B (to appear) (2023).
- [8] Wassily Hoeffding, *Probability inequalities for sums of bounded random variables*, J. Amer. Statist. Assoc. **58** (1963), 13–30. MR144363
- [9] Steven R. Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon, *Time-uniform Chernoff bounds via nonnegative supermartingales*, Probab. Surv. **17** (2020), 257–317. MR4100718
- [10] ———, *Time-uniform, nonparametric, nonasymptotic confidence sequences*, Ann. Statist. **49** (2021), no. 2, 1055–1080. MR4255119
- [11] Rasul A. Khan, *L_p -version of the Dubins-Savage inequality and some exponential inequalities*, J. Theoret. Probab. **22** (2009), no. 2, 348–364. MR2501324
- [12] Wouter M. Koolen and Peter Grünwald, *Log-optimal anytime-valid E-values*, Internat. J. Approx. Reason. **141** (2022), 69–82. MR4364896
- [13] Masayuki Kumon and Akimichi Takemura, *On a simple strategy weakly forcing the strong law of large numbers in the bounded forecasting game*, Ann. Inst. Statist. Math. **60** (2008), no. 4, 801–812. MR2453572
- [14] Laurent Mazliak and Glenn Shafer (eds.), *The Splendors and Miseries of Martingales—Their History from the Casino to Mathematics*, Trends in the History of Science, Birkhäuser/Springer, Cham, [2022] ©2022. MR4516531
- [15] Kenshi Miyabe and Akimichi Takemura, *The law of the iterated logarithm in game-theoretic probability with quadratic and stronger hedges*, Stochastic Process. Appl. **123** (2013), no. 8, 3132–3152. MR3062440
- [16] Aaditya Ramdas, Peter Grünwald, Vladimir Vovk, and Glenn Shafer, *Game-theoretic statistics and safe anytime-valid inference*, arXiv:2210.01948 (2022).
- [17] Aaditya Ramdas, Johannes Ruf, Martin Larsson, and Wouter M. Koolen, *Admissible anytime-valid sequential inference must rely on nonnegative martingales*, arXiv:2009.03167 (2020).
- [18] Aaditya Ramdas, Johannes Ruf, Martin Larsson, and Wouter M. Koolen, *Testing exchangeability: fork-convexity, supermartingales and e-processes*, Internat. J. Approx. Reason. **141** (2022), 83–109. MR4364897
- [19] Herbert Robbins, *Statistical methods related to the law of the iterated logarithm*, The Annals of Mathematical Statistics **41** (1970), no. 5, 1397–1409 (EN). MRMR277063
- [20] Herbert Robbins and David Siegmund, *Boundary crossing probabilities for the Wiener process and sample sums*, The Annals of Mathematical Statistics **41** (1970), no. 5, 1410–1429 (EN). MRMR277059
- [21] Takeyuki Sasai, Kenshi Miyabe, and Akimichi Takemura, *Erdos-Feller-Kolmogorov-Petrowsky law of the iterated logarithm for self-normalized martingales: a game-theoretic approach*, Ann. Probab. **47** (2019), no. 2, 1136–1161. MR3916944
- [22] Glenn Shafer, *Testing by betting: a strategy for statistical and scientific communication*, J. Roy. Statist. Soc. Ser. A **184** (2021), no. 2, 407–478. MR4255905
- [23] Glenn Shafer and Vladimir Vovk, *Probability and Finance*, Wiley Series in Probability and Statistics. Financial Engineering Section, Wiley-Interscience, New York, 2001, It's only a game! MR1852450
- [24] ———, *Game-Theoretic Foundations for Probability and Finance*, vol. 455, John Wiley & Sons, 2019.
- [25] Galen R. Shorack, *Probability for Statisticians*, Springer Texts in Statistics, Springer-Verlag, New York, 2000. MR1762415
- [26] J Ville, *Étude Critique de la Notion de Collectif.*, Gauthier-Villars, Paris, 1939.
- [27] A. Wald, *Sequential tests of statistical hypotheses*, Ann. Math. Statistics **16** (1945), 117–186. MR13275