

# PREDICTION WITH EVENTUAL ALMOST SURE GUARANTEES

BY CHANGLONG WU<sup>1,\*</sup> AND NARAYANA SANTHANAM<sup>1,†</sup>

<sup>1</sup>*Department of Electrical Engineering, University of Hawai'i at Manoa, \*wuchangl@hawaii.edu; †nsanthan@hawaii.edu*

We study the problem of sequentially predicting properties of a probabilistic model and its next outcome over an infinite horizon, with the goal of ensuring that the predictions incur only finitely many errors with probability 1. We introduce a general framework that models such prediction problems, provide general characterizations for the existence of successful prediction rules, and demonstrate the application of these characterizations through several concrete problem setups, including hypothesis testing, online learning, and risk domination. In particular, our characterizations allow us to recover the findings of Dembo and Peres (1994) with simple and elementary proofs, and offer a partial resolution to an open problem posed therein.

**1. Introduction.** Suppose we sequentially observe a sample  $X_1, X_2, \dots$  from an unknown probabilistic model. At each time step  $n$ , we make a prediction about either a property of the underlying model or the next sample  $X_{n+1}$ , based on the current observation  $X_1, \dots, X_n$  of the sample. There is a predefined binary  $\{0, 1\}$ -loss that measures whether the prediction is correct/acceptable or not. The natural questions then are: under what conditions will we be able to incur only finitely many losses with probability 1, so that the predictions are eventually correct almost surely? If we can do so, can we identify from what point we will not incur any more losses?

This problem setup was initiated by Cover [8]. In Cover's paper, the objective is to predict the irrationality of the mean of an underlying distribution over  $[0, 1]$  using *i.i.d.* samples of it. The prediction can be updated after every observation, but the learner is allowed only finitely many errors with probability 1. It is shown in [8] that such a prediction rule exists for distributions with means in a subset of  $[0, 1]$  that contains all rational numbers and has Lebesgue measure 1. Cover's setup was generalized substantially by Dembo and Peres [9] to identify general properties of distributions over  $\mathbb{R}^d$  using a topological criterion. Similar problem setups were continued in [13, 12, 15, 19]. However, all of these setups considered only the classification of the underlying distribution using *i.i.d.* samples.

In [30], the authors considered a problem setup of predicting upper bounds on the *next* observation of *i.i.d.* samples from distributions over  $\mathbb{N}$ , such that the next observation violates the bound only finitely often with probability 1. In [23], the authors obtain a stopping rule that additionally indicates when such a prediction has made the last mistake.

This paper introduces a general framework that encompasses the previously mentioned problem setups within a unified analysis framework. Informally, we consider the following general prediction setup. The observations are modeled as a general discrete-time random process  $X_1, X_2, \dots$  whose underlying probability measure (not necessarily *i.i.d.*)  $p$  belongs to a known collection  $\mathcal{P}$ . The learning process is a game between Nature and a learner, where the learner attempts to predict a property of  $p$  or future observations. Nature fixes a random process  $p \in \mathcal{P}$  at the beginning of the game. At each time step  $n$ , the learner makes a prediction  $Y_n$  using  $X_1, \dots, X_{n-1}$ . The prediction, the next realization  $X_n$ , and potentially the underlying process  $p$  are associated with a binary 0-1 loss  $\ell$ , where 1 indicates

---

*MSC2020 subject classifications:* Primary 62A01, 62M20; secondary 62L12.

an error/unsatisfactory prediction. The collection  $\mathcal{P}$  together with the loss  $\ell$  is said to be *eventually almost surely* (or *e.a.s.*) predictable if there is a strategy such that the learner makes only *finitely* many errors with probability 1 no matter what the underlying  $p \in \mathcal{P}$  is.

Our contributions in this paper establish a comprehensive theoretical framework for the *e.a.s.*-predictability. We summarized our main results in the follows.

1.1. *Summary of results.* Our first main result, Theorem 1, demonstrates that the *e.a.s.*-predictability of a model class is, in many cases, equivalent to a decomposition of the class into nested  $\eta$ -predictable subclasses. These are model classes where we can predict with a *bounded* number of errors and with confidence  $\geq 1 - \eta$  uniformly over all models in the class. Therefore, this result has the context of matching sample size to complexity—as the sample size increases, we address larger subsets of the model class. Put another way, this implies that such *e.a.s.*-predictable estimators can be obtained through regularization based on the sample at hand.

We illustrate how Theorem 1 can be applied to transform seemingly impossible tasks, as posed by Cover [8], into essentially trivial problems in Example 3. In Section 7, we demonstrate further applications of Theorem 1 across various contexts—hypothesis testing, risk domination, and online learning, to name a few. Notably, Theorem 1 enables us to rederive all the main results of Dembo and Peres [9] with straightforward and elementary proofs (see Section 7.1 and Appendix A). Moreover, we partially resolve an open problem posed by Dembo and Peres [9] in Section 7.1.

Theorem 1, while broadly applicable, does not always characterize *e.a.s.*-predictability in the most general settings—indeed, we present counterexamples for the same. We then refine Theorem 1 for several special cases. Theorem 2 completely characterizes the *e.a.s.*-predictability in the *supervised setting*, i.e., a setting when the losses are observable from samples, using methods reminiscent of structural risk minimization. In Theorem 4, we establish a different tight characterization of the *e.a.s.*-predictability for *pure estimation* problems with i.i.d. sampling and finite prediction domain. This result forms the basis of the almost sure hypothesis testing setup of Dembo and Peres [9] that we will study in detail in Section 7.1.

Finitely many errors do not imply that there exists a method, a stopping rule, that identifies when we are past the point of the final error. Indeed, in many cases, such a stopping rule is impossible. Section 5 addresses this refinement with the notion of *e.a.s.*-learnability. Informally, *e.a.s.*-learnability of a model class requires that for any given confidence parameter  $\eta > 0$ , we are able to find a prediction rule and stopping rule such that the probability of the prediction rule making errors after the stopping rule has ceased is upper bounded by  $\eta$ . We provide a characterization of *e.a.s.*-learnability in Theorem 8 through a notion of *identifiability*. This characterization is tight when the processes are i.i.d.

Lastly, we introduce other natural variations of the *e.a.s.*-predictability framework in Section 6, including a notion of weak *e.a.s.*-predictability and the notion of predictability with finite expected loss. We show that predictability with finite expected loss implies *e.a.s.*-predictability, and that *e.a.s.*-predictability implies weak *e.a.s.*-predictability. The converses do not hold, and these formulations are essentially different—we also provide examples that illustrate these nuances. We conclude the paper with several open problems that concern the exact characterizations of the above variations.

1.2. *Related work.* Perhaps the most related non-uniform consistency setup in the learning theory literature is non-uniform PAC learning [6, 4] and the concept of *Structural Risk Minimization* (SRM) [28, 27]. In this setup, the goal is to obtain a *distribution-free* sample complexity in the PAC model that could depend on the underlying hypothesis. It has been

shown by [16, 4] that a hypothesis class can be learned in such a sense in the *realizable* setting if and only if the class can be decomposed into a countable union of finite VC-dimensional classes. [25] generalized the result to the *agnostic* setup using SRM. There has also been research in the *distribution-dependent* setting along this line, aiming to quantify the *learning rates* of a hypothesis class with respect to the sample size [24, 1, 7]. In particular, [7] recently showed that under mild assumptions on the realizability of the underlying generating model, the learning rates can only asymptotically approach  $e^{-n}$ ,  $\frac{1}{n}$ , or be arbitrarily slow. Our framework is also related to the learning of *recursive functions* (a.k.a. inductive inference), introduced by Gold [10]. A typical setup starts with a set  $\mathcal{H}$  of functions that map  $\mathbb{N} \rightarrow \{0, 1\}$ . Nature fixes a function  $h$  from  $\mathcal{H}$  at the beginning. At each time step  $n$ , the learner attempts to predict  $h(n)$  using the history  $h(1), \dots, h(n-1)$  observed thus far. Nature then reveals the true value  $h(n)$  after the learner has made the prediction. The goal is a *computable* learner that makes only *finitely* many errors no matter what  $h \in \mathcal{H}$  is. It is shown by [3] that a finite error computable predictor exists if and only if there exists a computable function  $g : \mathbb{N} \rightarrow \mathbb{N}$  such that the time complexity of  $\mathcal{H}$  is *eventually dominated* by  $g$ . We refer the reader to [10, 3, 5, 32] for more results along this line. One may view the learning of recursive functions as a special case of our *e.a.s.*-predictability where the underlying process is *deterministic* but restricts the learner to be computable.

**2. Problem setup.** Let  $\mathcal{X}$  be a Polish space (separable completely metrizable topological space) with Borel  $\sigma$ -algebra and  $\mathcal{P}$  be a collection of probability measures over the cylinder  $\sigma$ -algebra over  $\mathcal{X}^\infty$ . We consider a discrete time random process  $\mathbf{X} = \{X_n\}_{n \in \mathbb{N}^+}$  generated by sampling from a probability law  $p \in \mathcal{P}$ . We will denote  $X_i^j = (X_i, X_{i+1}, \dots, X_j)$  in the sequel.

Prediction is modeled as a measurable function  $\Phi : \mathcal{X}^* \rightarrow \mathcal{Y}$ , where  $\mathcal{X}^*$  denotes the set of all finite strings of sequences from  $\mathcal{X}$ , and  $\mathcal{Y}$  is the set of all predictions. The *loss* is a function  $\ell : \mathcal{P} \times \mathcal{X}^* \times \mathcal{Y} \rightarrow \{0, 1\}$  such that  $\ell(p, \cdot, \cdot)$  is measurable for all  $p \in \mathcal{P}$ . We consider the property we are estimating to be defined implicitly by the subset of  $\mathcal{P} \times \mathcal{X}^* \times \mathcal{Y}$  where  $\ell = 0$ , and therefore, in a slight abuse of notation sometimes refer to  $\ell$  as a *property* as well.

We consider the following game that proceeds in time indexed by  $\mathbb{N}^+$ . The game has two parties: the Learner and Nature. Nature chooses some model  $p \in \mathcal{P}$  to begin the game. At each time step  $n$ , the Learner makes a prediction  $Y_n = \Phi(X_1^{n-1})$  based on the current observation  $X_1^{n-1}$  generated according to  $p$ . Nature then generates  $X_n$  based on  $p$  and  $X_1^{n-1}$ .

The Learner fails at step  $n$  if  $\ell(p, X_1^n, Y_n) = 1$ . The Learner targets a strategy that minimizes the cumulative loss in the infinite horizon, without knowledge of the model that Nature chooses at the beginning.

The loss in general can be any function of the probability model in addition to the sample observed, and our prediction on the sample. When the loss depends on the probability model, there may be no direct way to estimate the loss incurred at say, step  $n$ , from observations of the sample  $X_1^n$  even after the prediction  $Y_n$  is made. We call such setups the *unsupervised setting* borrowing from learning theory. A special case is the *supervised setting*, where we define the loss to be a function from  $\mathcal{X}^* \times \mathcal{Y}$  to  $\{0, 1\}$ .

**DEFINITION 1** ( $\eta$ -predictability). *A collection  $(\mathcal{P}, \ell)$  is  $\eta$ -predictable, if there exists a prediction rule  $\Phi : \mathcal{X}^* \rightarrow \mathcal{Y}$  and a sample size  $n$  such that for all  $p \in \mathcal{P}$ ,*

$$p \left( \sum_{i=n}^{\infty} \ell(p, X_1^i, \Phi(X_1^{i-1})) > 0 \right) \leq \eta,$$

*i.e. the probability that the learner makes errors after step  $n$  is at most  $\eta$  uniformly over  $\mathcal{P}$ .*

EXAMPLE 1. Let  $\mathcal{P}^\epsilon$  be the class of all i.i.d. Bernoulli processes with parameters in  $\{p \in [0, 1] : |p - \frac{1}{2}| \geq \epsilon\} \cup \{\frac{1}{2}\}$  for some absolute constant  $\epsilon > 0$ . We define the loss to be  $\ell(p, X_1^n, Y_n) = 1\{Y_n \neq 1\{p = \frac{1}{2}\}\}$ , i.e., we would like to determine whether the underlying parameter equals  $\frac{1}{2}$  or not. By Chernoff bound, we know that  $(\mathcal{P}^\epsilon, \ell)$  is  $\eta$ -predictable for all  $\eta, \epsilon > 0$ , by predicting the (same) empirical mean of  $X_1^N$  for all time step  $> N$  where  $N$  is any constant  $\geq \frac{2 \log 1/\eta}{\epsilon^2}$ .

DEFINITION 2 (e.a.s.-predictable). A collection  $(\mathcal{P}, \ell)$  is said to be eventually almost surely (e.a.s.)-predictable, if there exists a prediction rule  $\Phi$ , such that for all  $p \in \mathcal{P}$

$$p \left( \sum_{n=1}^{\infty} \ell(p, X_1^n, \Phi(X_1^{n-1})) < \infty \right) = 1.$$

We need a technical definition that will help simplify notation further.

DEFINITION 3. A nesting of  $\mathcal{P}$  is a collection of subsets of  $\mathcal{P}$ ,  $\{\mathcal{P}_i : i \geq 1\}$  such that  $\mathcal{P}_1 \subset \mathcal{P}_2 \subset \dots$  and  $\bigcup_{i \geq 1} \mathcal{P}_i = \mathcal{P}$ .

The following lemma characterize immediate connections between the above definitions.

LEMMA 1. Let  $\mathcal{P}$  be a collection of models,  $\{\mathcal{P}_i, i \geq 1\}$  be a nesting of  $\mathcal{P}$ . If for all  $\eta > 0$  and  $i \in \mathbb{N}^+$ ,  $(\mathcal{P}_i, \ell)$  is  $\eta$ -predictable. Then  $(\mathcal{P}, \ell)$  is e.a.s.-predictable.

PROOF. From the definition of  $\eta$ -predictability, we can choose an increasing sequence  $\{b_i, i \geq 1\}$ , and predictors  $\Phi_i$  for  $\mathcal{P}_i$  respectively as follows. For all  $i$  and for all  $p \in \mathcal{P}_i$ , the probability  $\Phi_i$  makes errors after step  $b_i$  is at most  $2^{-i}$ .

The predictor  $\Phi$  is then constructed from  $\{\Phi_i, i \geq 1\}$  as follows: use predictor  $\Phi_i$  when the length  $T$  of the observed sample satisfies  $b_i \leq T < b_{i+1}$ .

Let  $p \in \mathcal{P}_k \subset \mathcal{P}$ . Because the collections  $\mathcal{P}_i$  are nested, for all  $i \geq k$ ,  $p \in \mathcal{P}_i$ . During the phase  $\Phi$  coincides with  $\Phi_i$ , the probability of  $\Phi$  making an error is  $\leq 2^{-i}$ . The result follows using the Borel-Cantelli lemma.  $\square$

EXAMPLE 2. Let  $\mathcal{P}$  be the class of all iid Bernoulli processes with parameters in  $[0, 1]$  and  $\ell$  is the same loss as in Example 1. We claim that  $(\mathcal{P}, \ell)$  is e.a.s.-predictable using Lemma 1. For any  $i \geq 1$ , we define  $\mathcal{P}_i = \mathcal{P}^{1/i}$ , where  $\mathcal{P}^{1/i}$  is defined as in Example 1 with  $\epsilon = \frac{1}{i}$ . Clearly,  $\{\mathcal{P}_i, i \geq 1\}$  forms a nesting of  $\mathcal{P}$ . By Example 1, we also know that each  $\mathcal{P}_i$  is  $\eta$ -predictable for all  $\eta > 0$ , the claim follows.

*Refined notions of nestings.* Motivated by the above lemma, we introduce the following two refinements on the notion of nestings.

DEFINITION 4 (Universal nestings). A nesting  $\{\mathcal{P}_i, i \geq 1\}$  of  $\mathcal{P}$  is a universal nesting w.r.t loss  $\ell$ , if for all  $\eta > 0$  and  $i \geq 1$ ,  $(\mathcal{P}_i, \ell)$  is  $\eta$ -predictable.

DEFINITION 5 ( $\eta$ -nestings). For any  $\eta > 0$ , a nesting  $\{\mathcal{P}_i^\eta, i \geq 1\}$  of  $\mathcal{P}$  is an  $\eta$ -nesting w.r.t. loss  $\ell$ , if for all  $i \geq 1$ ,  $(\mathcal{P}_i^\eta, \ell)$  is  $\eta$ -predictable.

Clearly, if a class  $\mathcal{P}$  has a universal nesting then it has  $\eta$ -nesting for all  $\eta > 0$ . However, the converse is not true, as we will see in Section 3. Indeed, as we will see in Theorem 1 that these two concepts will form the basis for characterizing e.a.s.-predictability.

The following technical lemma will be useful in our following proofs.

LEMMA 2. *Let  $\mathcal{P}$  be a collection of probability measures and let  $\ell$  be a loss function. Suppose  $\{\mathcal{P}_i : i \geq 1\}$  is a nesting of  $\mathcal{P}$  such that for all  $i \geq 1$ ,  $(\mathcal{P}_i, \ell)$  is  $\eta$ -predictable for some  $\eta > 0$ .*

*Then there exists a relabeling of the sets in a nesting  $\{\mathcal{P}'_i : i \geq 1\}$  of  $\mathcal{P}$  such that  $(\mathcal{P}'_i, \ell)$  is  $\eta$ -predictable with sample size  $i$ . Namely, there is a predictor  $\Phi_i$  such that for all  $p \in \mathcal{P}'_i$ , the probability  $\Phi_i$  incurs non-zero  $\ell$ -loss on samples with size larger than  $i$  is  $\leq \eta$ .*

PROOF. Since  $\mathcal{P}_i$  is  $\eta$ -predictable, there exists a number  $n'_i$  and a predictor  $\Phi_i$  such that for all  $p \in \mathcal{P}_i$ , the probability  $\Phi_i$  incurs non-zero  $\ell$ -loss on samples larger than  $n'_i$  is  $\leq \eta$ . We can therefore choose an increasing sequence  $\{n_i, i \geq 1\}$  with  $n_i \geq n'_i$ . Note that each  $\mathcal{P}_i$  is  $\eta$ -predictable with sample size  $n_i$ . For  $n_k \leq i < n_{k+1}$ , set  $\mathcal{P}'_i = \mathcal{P}_k$ .  $\{\mathcal{P}'_i : i \geq 1\}$  is the desired nesting and the lemma follows.  $\square$

**3. Characterization of *e.a.s.*-predictability.** We first provide a general characterization of the *e.a.s.*-predictability without any assumption on the class  $\mathcal{P}$  and loss  $\ell$ .

THEOREM 1. *Consider a collection  $\mathcal{P}$  with a loss  $\ell : \mathcal{P} \times \mathcal{X}^* \times \mathcal{Y} \rightarrow \{0, 1\}$ .  $(\mathcal{P}, \ell)$  is *e.a.s.*-predictable if there exists a nesting  $\{\mathcal{P}_i : i \geq 1\}$  of  $\mathcal{P}$  such that for all  $\eta > 0$ ,  $i \geq 1$ ,  $(\mathcal{P}_i, \ell)$  is  $\eta$ -predictable.*

*Conversely, if  $(\mathcal{P}, \ell)$  is *e.a.s.*-predictable, then for all  $\eta > 0$ , there is a nesting  $\{\mathcal{P}^\eta_i : i \geq 1\}$  of  $\mathcal{P}$  such that for all  $i \geq 1$ ,  $(\mathcal{P}^\eta_i, \ell)$  is  $\eta$ -predictable.*

PROOF. The sufficiency follows directly from Lemma 1. We now prove the necessary condition. Suppose  $\mathcal{P}$  is *e.a.s.*-predictable, we show that  $\mathcal{P}$  can be decomposed into a nesting of  $\eta$ -predictable collections. By Definition 2, there exists a predictor  $\Phi$  such that for all  $p \in \mathcal{P}$ ,  $\Phi$  makes finitely many errors with probability 1. For  $\eta > 0$ , we define

$$\mathcal{P}^\eta_n = \{p \in \mathcal{P} \mid p(\Phi \text{ makes errors after time } n) < \eta\},$$

so that for all  $n$ ,  $\mathcal{P}^\eta_n$  is  $\eta$ -predictable by definition. Further, by definition,  $\forall n \in \mathbb{N}^+$ ,  $\mathcal{P}^\eta_n \subset \mathcal{P}^\eta_{n+1}$ . To see that the union of  $\mathcal{P}^\eta_n$  over all  $n$  is  $\mathcal{P}$ , for all  $p \in \mathcal{P}$ , we consider the event

$$A_k^p = \{X_1^\infty \mid \sum_{n=k}^{\infty} \ell(p, X_1^n, \Phi(X_1^{n-1})) > 0\}.$$

We have  $p(A_k^p) \rightarrow 0$  as  $k \rightarrow \infty$  by *e.a.s.*-predictability of  $\Phi$ . Therefore, there must be some  $k$  such that  $p(A_k^p) < \eta$ , and for such a number  $k$  we have  $p \in \mathcal{P}^\eta_k$ . Therefore,  $\mathcal{P} = \bigcup_{n \in \mathbb{N}^+} \mathcal{P}^\eta_n$ . This completes the proof.  $\square$

Theorem 1 shows that to prove a collection  $(\mathcal{P}, \ell)$  is *e.a.s.*-predictable, it is sufficient to find a *universal nesting* of  $\mathcal{P}$ . To prove that a collection  $(\mathcal{P}, \ell)$  is not *e.a.s.*-predictable it is sufficient to show that for some  $\eta > 0$  there is no  $\eta$ -nesting of  $\mathcal{P}$ .

We provide an example below, which illustrates how Theorem 1 can be used to derive the result of Cover [8]. See also Section 7 for more applications in different contexts.

EXAMPLE 3. *The task is to predict whether the parameter  $p$  of an iid Bernoulli( $p$ ) process is rational or not using samples from it.*

*Therefore our predictor*

$$\Phi : \{0, 1\}^* \rightarrow \{\text{rational}, \text{irrational}\}.$$

In [8], Cover showed a scheme that predicted accurately with only finitely many errors for all rational sources, and for a set of irrationals with Lebesgue measure 1. Here we show a more transparent version of Cover's proof as well as subsequent refinements in [12] using Theorem 1 above and an argument evocative of regularization.

Define the loss  $\ell(p, X_1^n, Y_n) = 0$  iff  $Y_n$  gives the correct irrationality of  $p$ . Note that the setting is what we would call the "unsupervised" case and that there is no way to judge if our predictions thus far are right or wrong.

Let  $r_1, r_2, \dots$  be an enumeration of rational numbers in  $[0, 1]$ . Let  $B(p, \epsilon)$  be the set of numbers in  $[0, 1]$  whose  $L_1$  distance from  $p$  is  $< \epsilon$ . For all  $k$ , let

$$\mathcal{S}_k = \left( [0, 1] \setminus \bigcup_{i=1}^{\infty} B(r_i, \frac{1}{k2^i}) \right) \cup \{r_1, \dots, r_k\}$$

be the set that excludes a ball centered on each rational number, but throws back in the first  $k$  rational numbers. Note that the Lebesgue measure of  $\mathcal{S}_k$  is  $\geq 1 - \frac{1}{k}$ . Now  $\mathcal{S}_k$  contains exactly  $k$  rational numbers, the rest being irrational. Moreover,  $\mathcal{S}_k$  contains no irrational number within distance  $\leq 2^{-k}/k$  from any of the included rationals. Hence, the set  $\mathcal{B}_k$  of Bernoulli processes with parameters in  $\mathcal{S}_k$  is  $\eta$ -predictable for all  $\eta > 0$ .

From Theorem 1, we can conclude that the collection  $\mathcal{B} \stackrel{\text{def}}{=} \bigcup_{k \in \mathbb{N}} \mathcal{B}_k$  is e.a.s.-predictable. Note that every rational number belongs to  $\mathcal{S} = \bigcup_{k \in \mathbb{N}} \mathcal{S}_k$ , and the set of irrational numbers in  $\mathcal{S}$  has Lebesgue measure 1, proving [8, Theorem 1].

Conversely, let  $\mathcal{S} \subset [0, 1]$  and  $\mathcal{B}$  be the Bernoulli variables with parameters in  $\mathcal{S}$ . We show that if  $\mathcal{B}$  is e.a.s.-predictable for rationality of the underlying parameter, then  $\mathcal{S} = \bigcup_{k \in \mathbb{N}} \mathcal{S}_k$  such that  $\mathcal{S}_k \subset \mathcal{S}_{k+1}$  and

$$\inf\{|r - x| : r, x \in \mathcal{S}_k, r \text{ is rational, } x \text{ is irrational}\} > 0.$$

Since  $\mathcal{B}$  is e.a.s.-predictable, Theorem 1 yields that for any  $\eta > 0$ , the collection  $\mathcal{B}$  can be decomposed as  $\mathcal{B} = \bigcup_k \mathcal{B}_k$  where each  $\mathcal{B}_k$  is  $\eta$ -predictable and  $\forall k, \mathcal{B}_k \subset \mathcal{B}_{k+1}$ . Let  $\mathcal{S}_k$  be the set of parameters of the sources in  $\mathcal{B}_k$ . Intuitively,  $\eta$ -predictability of  $\mathcal{B}_k$  implies that we must have

$$\inf\{|u - v| : u, v \in \mathcal{S}_k, u \text{ rational, } v \text{ irrational}\} > 0,$$

or else we would not be able to universally attest to rationality with confidence  $\geq 1 - \eta$  using a bounded number of samples. See Lemma 12 for a formal proof.

Suppose we want  $\mathcal{S}$  to contain all rational numbers in  $[0, 1]$ . Then it follows (see Lemma 13) that the subset of irrational numbers of  $\mathcal{S}_k$  must be nowhere dense. Therefore, the set of irrationals in  $\mathcal{S}$  is a meager or Baire first category set [22, Chapter 2.1], proving [12, Theorem 2].

Note that the necessary and sufficient conditions in Theorem 1 do not necessarily match in general. Indeed, the following example shows that the necessary (respectively sufficient) condition in Theorem 1 is not sufficient (respectively necessary).

EXAMPLE 4. Let  $\mathcal{P}$  be a class of binary (taking values 0 or 1) random processes that converge to either 0 or 1 in probability. Formally,  $\mathcal{P}$  is the collection of all probability measures  $p_b, b \in \{0, 1\}$ , defined on the Borel  $\sigma$ -algebra of  $\{0, 1\}^\infty$ , that satisfies

$$\lim_{n \rightarrow \infty} p_b(X_n = b) = 1.$$

The task of the prediction  $Y_n$  is to predict the parameter  $b$  associated with the process, and takes values in  $\{0, 1\}$ . The loss  $\ell$  associated with the prediction is defined as

$$\ell(p_b, X_1^n, Y_n) = 1\{Y_n \neq b\}.$$

We now show that the condition deemed necessary in Theorem 1 holds for  $(\mathcal{P}, \ell)$ . To see this, let  $\mathcal{P}_i^\eta$  be the class of processes  $p_b \in \mathcal{P}$  such that for all  $n \geq i$

$$p_b(X_n = b) \geq 1 - \eta.$$

The  $\eta$ -predictability of  $\mathcal{P}_i^\eta$  follows because  $p_b(X_i = b) \geq 1 - \eta$ , and a predictor that predicts  $X_i$  for all time steps  $\geq i$  will incur loss 0 past time step  $i$  whenever  $X_i = b$ .

We show that the collection  $(\mathcal{P}, \ell)$  is, however, not e.a.s.-predictable. To see this, suppose such a prediction rule  $\Phi$  exists. We first observe that there exists a number  $N(m)$  such that for all finite binary sequences  $x_1, \dots, x_m$  of length  $m$  and all  $b \in \{0, 1\}$

$$(3.1) \quad \Phi(x_1, \dots, x_m, b, \dots, b) = b,$$

whenever the number of  $b$ 's is larger than  $N(m)$ . This holds because each of the  $2 \cdot 2^m$  semi-infinite strings  $x_1, \dots, x_m, b, \dots$  corresponds to a process in  $\mathcal{P}$  that assigns probability 1 to that string. If (3.1) did not hold,  $\Phi$  would make an infinite number of errors on one of these processes, contradicting the e.a.s.-predictability of  $\Phi$  on  $(\mathcal{P}, \ell)$ .

We now construct the following process  $p_0$  in  $\mathcal{P}$  that will break  $\Phi$ . Let  $M_0 = 0$ ,  $M_1 = N(0) + 1$ , and recursively define  $M_n = N(M_{n-1}) + 1$ . The process  $p_0$  is partitioned into independent sample blocks, where the  $n$ th block ranges from  $X_{M_{n-1}+1}$  to  $X_{M_n}$  such that  $X_{M_{n-1}+1} = X_{M_{n-1}+2} = \dots = X_{M_n}$  and

$$p_0(X_{M_{n-1}+1} = 0) = 1 - \frac{1}{n}.$$

Let  $A_n$  be the event that  $X_{M_{n-1}+1} = 1$ . We have  $A_n$  happens infinitely often almost surely by the converse of the Borel-Cantelli lemma, since  $\sum_n p_0(A_n) = \sum_n \frac{1}{n} = \infty$  and  $A_n$ 's are independent. By construction,  $\Phi$  makes errors in sample block  $n$  if  $A_n$  happens, hence  $\Phi$  makes infinitely many errors almost surely. But clearly  $p_0 \in \mathcal{P}$ , contradicting the e.a.s.-predictability of  $\Phi$ .

With similar construction, we can also show that a collection  $(\mathcal{P}, \ell)$  is e.a.s.-predictable does not imply the existence of a universal nesting of  $\mathcal{P}$ . See Example 11 for a detailed construction.

While Theorem 1 may look rather innocuous, it offers a partial resolution to an open problem in [9]. Let  $\mathcal{H}_1$  and  $\mathcal{H}_2$  be disjoint classes of distributions over  $\mathbb{R}^d$ . Let  $\mathcal{H}$  be the class of all iid random processes with marginal distributions from  $\mathcal{H}_1 \cup \mathcal{H}_2$ . Dembo and Peres [9] considered the problem of identifying whether the marginal of a iid random process in  $\mathcal{H}$  comes from  $\mathcal{H}_1$  or  $\mathcal{H}_2$  by observing samples from it. The prediction domain now is  $\mathcal{Y} = \{1, 2\}$  and loss is  $\ell^{HT}(p, X_1^n, Y_n) = 1\{p \notin \mathcal{H}_{Y_n}\}$ . Dembo and Peres [9] showed that if the distributions in  $\mathcal{H}_1 \cup \mathcal{H}_2$  have densities and there exists some  $r > 1$  such that the  $r$ -th norm of the densities are finite, then  $(\mathcal{H}, \ell^{HT})$  is e.a.s.-predictable iff the distributions in  $\mathcal{H}_1$  and  $\mathcal{H}_2$  are  $F_\sigma$ -separable (see Section 7.1 for definition) under any metric consistent with weak convergence topology.

Dembo and Peres [9] asked whether the condition  $r > 1$  can be removed. We give a positive answer to this problem as follows.

**COROLLARY 1.** *Suppose there exists a monotonically increasing function  $G : \mathbb{R}^+ \rightarrow \mathbb{R}^+$  with  $\lim_{x \rightarrow \infty} G(x) = \infty$  such that for any distribution  $p \in \mathcal{H}_1 \cup \mathcal{H}_2$  with density  $f_p(x)$ , we have  $\mathbb{E}_{X \sim p}[G(f_p(X))] < \infty$ . Then  $(\mathcal{H}, \ell^{\text{HT}})$  is *e.a.s.-predictable* iff the distributions in  $\mathcal{H}_1$  and  $\mathcal{H}_2$  are  $F_\sigma$ -separable under weak convergence topology.*

The proof and a discussion of the Corollary above is left to Section 7.1.

**4. Specialized settings.** While the necessary and sufficient conditions in Theorem 1 do not match up in general as shown above, the characterization of *e.a.s.-predictability* can be tightened in several natural settings. We split our further analysis into two settings—one where the loss can be observed (the supervised setting) and where the loss cannot be observed (unsupervised setting). We provide tight characterizations for the supervised setting, and for *iid* generated data when the loss is only a function of the source and the prediction in the unsupervised setting.

In this section, the distinction between  $\eta$ -nesting and universal nesting will become clear as well—we will also encounter examples that allow  $\eta$ -nestings for all  $\eta > 0$ , but which lack a universal nesting even in the supervised setting.

**4.1. Supervised setting.** We now characterize *e.a.s.-predictability* in *supervised* setting, i.e. the loss  $\ell : \mathcal{X}^* \times \mathcal{Y} \rightarrow \{0, 1\}$  is independent of the underlying source  $p$  (or equivalently the loss can be gauged from the samples). We show that the existence of  $\eta$ -nesting for all  $\eta > 0$  is necessary and sufficient to achieve *e.a.s.-predictability* in the supervised case. As we have demonstrated in Example 4 the existence of  $\eta$ -nesting for all  $\eta > 0$  does not imply *e.a.s.-predictability* in the most general case. Theorem 2 shows that the implication holds when the loss is supervised.

**THEOREM 2.** *Consider a collection  $\mathcal{P}$  with a loss  $\ell : \mathcal{X}^* \times \mathcal{Y} \rightarrow \{0, 1\}$  (i.e. the supervised setting).  $(\mathcal{P}, \ell)$  is *e.a.s.-predictable* iff for all  $\eta > 0$ , there exists a  $\eta$ -nesting of  $\mathcal{P}$ .*

**PROOF.** The necessity follows directly from Theorem 1. We now prove that the existence of  $\eta$ -nesting for all  $\eta > 0$  is also sufficient. Suppose that for all  $j \in \mathbb{N}$ , there exists a nesting  $\{\mathcal{P}_n^j : n \geq 1\}$  of  $\mathcal{P}$  such that  $\mathcal{P}_n^j$  is  $2^{-j}$ -predictable for all  $n \geq 1$ . Furthermore, from Lemma 2, we can choose a decomposition such that  $\mathcal{P}_n^j$  is  $2^{-j}$  predictable with sample size  $n$ . Therefore, there exist predictors  $\Phi_{n,j}$  such that for all  $p \in \mathcal{P}_n^j$

$$p(\Phi_{n,j} \text{ makes errors after time } n) \leq 2^{-j}.$$

We construct a predictor  $\Phi$  for  $\mathcal{P}$  as follows. At each time step  $T$ , let  $I(n, j)$  be the indicator that  $\Phi_{n,j}$  makes no error on  $X_1^{T-1}$  after time  $n$ . Let

$$(4.1) \quad (k, i) = \arg \min_{(n,j) \in \mathbb{N} \times \mathbb{N}} \{j + n \mid I(n, j) = 1\}.$$

The prediction is defined as  $\Phi(X_1^{T-1}) = \Phi_{k,i}(X_1^{T-1})$ .

We claim that the predictor  $\Phi$  will make only finitely many errors with probability 1 for all models in  $\mathcal{P}$ . Fix some  $p \in \mathcal{P}$ . Let  $n_j = \min\{n \mid p \in \mathcal{P}_n^j\}$ . Define the event

$$A_j = \{\Phi_{n_j, j} \text{ makes errors after time step } n_j\}.$$

We have  $\sum_{j=1}^{\infty} p(A_j) \leq \sum_{j=1}^{\infty} 2^{-j} < \infty$ .

Therefore, the Borel-Cantelli lemma implies that there is a set with probability 1, such that on every semi-infinite sequence  $X_1^\infty$  in that set, there is a  $J$  such that  $\Phi_{n_j, J}$  makes no



errors after step  $n_J$ . By construction of  $\Phi$ , for  $X_1^\infty$  in the set of probability 1 above, we will therefore never choose an estimator  $\Phi_{n,j}$  with  $n + j > n_J + J$  in step (4.1). If some  $\Phi_{n,j}$  with  $n + j \leq n_J + J$  makes infinitely many errors, it will no longer appear in the feasible set in (4.1) after some time step  $T \geq n$ . Since there are only finitely many predictors  $\Phi_{n,j}$  with  $n + j \leq n_J + J$ , the procedure will eventually choose some predictor that makes finitely errors.  $\square$

REMARK 1. *Note that, the approach in the proof of Theorem 2 can be viewed as regularization. The measure of complexity here is  $n + j$ . Here  $j$  equals the negative log of the confidence probability of the  $2^{-j}$ -nesting  $\{\mathcal{P}_n^j, n \geq 1\}$  and  $n$  denotes that we restrict the class to  $\mathcal{P}_n^j$ , i.e., it is a measure of the complexity of the class considered. The smaller  $n + j$  is, the less the complexity the  $\mathcal{P}_n^j$  has. The indicator  $I(n, j)$  is interpreted as consistency of the sample with class  $\mathcal{P}_n^j$ . At each step, equation (4.1) selects a subclass with minimum complexity that is consistent with current sample.*

While *e.a.s.*-predictability is completely characterized by the existence of  $\eta$ -nesting for all  $\eta > 0$  in the supervised setting, *e.a.s.*-predictability does not imply the existence of a *universal nesting* as the following example demonstrates. A corollary of this is that the existence of  $\eta$ -nesting for all  $\eta$  is not equivalent to the existence of a universal nesting even in the supervised setting—the two are different concepts.

EXAMPLE 5. *Let  $\mathcal{P}$  be the class of the following processes. For any infinite sequence  $M_1 < M_2 < \dots < M_n < \dots \in \mathbb{N}^\infty$ , we define process  $p^M$  over  $\{0, 1\}^\infty$  as follows:*

$$X_{M_n} \sim \text{Bernoulli}\left(\frac{1}{2}\right) \text{ independently for all } n \in \mathbb{N},$$

and  $X_t = 0$  for all other time step  $t$ . The loss  $\ell$  is defined as follows:

$$\ell(p^M, X_1^t, Y_t) = 0 \text{ iff there exists at least one } 1 \text{ in } X_1^t.$$

*Note that the loss only depends on the sample but not the underlying source and predictions. Clearly, the loss is observable from the sample, thus the setup is supervised. Furthermore,  $(\mathcal{P}, \ell)$  is *e.a.s.*-predictable since for any source  $p^M$ ,  $p^M(\exists n \geq 1, X_{M_n} = 1) = 1$ .*

*Let  $\mathcal{P}_1 \subset \mathcal{P}_2 \subset \dots \subset \mathcal{P}$  be an arbitrary nesting of  $\mathcal{P}$ . We show that  $\{\mathcal{P}_k, k \geq 1\}$  cannot be a universal nesting of  $\mathcal{P}$  by contradiction.*

*Suppose the above were a universal nesting, so each  $\mathcal{P}_k$  in the nesting is  $\eta$ -predictable for all  $\eta > 0$ . For all  $k$ , let  $B_k$  be the sample size for  $\mathcal{P}_k$  to achieve  $\frac{1}{2^{k+1}}$ -predictability. Now let  $p^B \in \mathcal{P}$  be the source that is associated with sequence  $\{B_k + 1\}_{k \in \mathbb{N}^+}$ .*

*We claim that  $p^B \notin \mathcal{P}_k$  for all  $k \in \mathbb{N}^+$ . To see this, note that  $p^B(\forall i \leq k, X_{B_i+1} = 0) = \frac{1}{2^k}$  which in turn implies that with probability at least  $\frac{1}{2^k}$ , we make an error at step  $B_k + 1$ . If for any  $k$ ,  $p^B$  were in  $\mathcal{P}_k$ ,  $\mathcal{P}_k$  would not be  $\frac{1}{2^{k+1}}$ -predictable. Therefore  $p^B \notin \bigcup_k \mathcal{P}_k$ . However, we also know that  $p^B \in \mathcal{P}$ , meaning that  $\{\mathcal{P}_k, k \geq 1\}$  cannot be a nesting of  $\mathcal{P}$ .*

4.2. *Unsupervised setting.* While  $\eta$ -nestings (for all  $\eta$ ) characterizes the supervised setting, it does not carry over for the unsupervised case in the most general cases. In this subsection, we introduce several special cases for the unsupervised setting where tight characterizations are available.

We first prove the following theorem, which shows that if  $\mathcal{P}$  is *countable* then existence of universal nesting  $\Leftrightarrow$  *e.a.s.*-predictability  $\Leftrightarrow$  existence of  $\eta$ -nesting for all  $\eta > 0$ .

**THEOREM 3.** *Let  $\mathcal{P}$  be a countable class and  $\ell$  be an arbitrary loss. Then there exists a universal nesting of  $\mathcal{P} \Leftrightarrow (\mathcal{P}, \ell)$  is e.a.s.-predictable  $\Leftrightarrow$  there exist  $\eta$ -nesting of  $\mathcal{P}$  for all  $\eta > 0$ .*

**PROOF.** Let  $p_1, p_2, \dots$  be an arbitrary enumeration of  $\mathcal{P}$ .

We first show that if  $(\mathcal{P}, \ell)$  is e.a.s.-predictable then there is a universal nesting for  $\mathcal{P}$ . Define  $\mathcal{P}_k = \{p_1, \dots, p_k\}$ . We show that  $(\mathcal{P}_k, \ell)$  is  $\eta$ -predictable for all  $\eta > 0$  and  $k \geq 1$ . To see this, let  $\Phi$  be the e.a.s.-prediction rule for  $(\mathcal{P}, \ell)$ . For any given  $\eta > 0$  and  $i \geq 1$ , there exists a number  $N_{i,\eta}$  such that  $\Phi$  makes errors after step  $N_{i,\eta}$  w.p.  $\leq \eta$  when  $p_i$  is in force. Now, for any  $\mathcal{P}_k$ , we know that  $\Phi$  makes no errors after step  $\max_{1 \leq i \leq k} \{N_{i,\eta}\}$  w.p.  $\geq 1 - \eta$  no matter what source  $\in \mathcal{P}_k$  is in force. Therefore  $(\mathcal{P}_k, \ell)$  is  $\eta$ -predictable for all  $\eta > 0$  and  $k \geq 1$ .

We now prove that if there exists  $\eta$ -nesting of  $\mathcal{P}$  for all  $\eta > 0$  then  $(\mathcal{P}, \ell)$  is e.a.s.-predictable. To see this, let  $\{\mathcal{P}_k^j, k \geq 1\}$  be the  $2^{-j}$ -nesting of  $\mathcal{P}$ . By Lemma 2, we can assume that there exists prediction rule  $\Phi_k^j$  such that  $\Phi_k^j$  makes no errors after step  $k$  w.p.  $\geq 1 - 2^{-j}$  for all  $p \in \mathcal{P}_k^j$ . We can assume, w.l.o.g., that  $\mathcal{P}_k^j$  contains only the sources in  $\{p_1, \dots, p_k\}$ . For all  $j \geq 1$ , let  $K_j$  be the minimum index of class in  $\{\mathcal{P}_k^j\}$  such that  $\mathcal{P}_{K_j}^j$  contains all sources in  $\{p_1, \dots, p_j\}$ . We now construct the e.a.s.-prediction rule as follows. We partition the prediction into phases, initially at phase 0. We are in phase  $j \geq 1$  if the time step  $n$  satisfies  $K_j \leq n < K_{j+1}$ . At phase  $j \geq 1$  we use  $\Phi_{K_j}^j$  to make the prediction and predict arbitrarily at phase 0.

Suppose the underlying source is  $p_t$ , we know that  $\mathcal{P}_{K_j}^j$  contains  $p_t$  for all  $j \geq t$ . Meaning that, the probability we make errors at phase  $j \geq t$  is at most  $2^{-j}$ . The theorem follows by the Borel-Cantelli lemma.  $\square$

We now introduce another unsupervised setting where we obtain a full characterization of e.a.s.-predictability—the *pure estimation* setting with *finite* prediction domain and *iid* sampling, where the loss is a function of only the source and prediction. This includes, e.g., the almost sure hypothesis testing framework of Dembo and Peres [9].

**THEOREM 4.** *Let  $\mathcal{P}$  be a model collection of iid measures over  $\mathcal{X}^\infty$ ,  $\ell$  is a loss function that only depends on the prediction and underlying source but not samples, i.e.  $\ell : \mathcal{P} \times \mathcal{Y} \rightarrow \{0, 1\}$ . If  $|\mathcal{Y}|$  is finite, then  $(\mathcal{P}, \ell)$  is e.a.s.-predictable iff there exists a universal nesting  $\{\mathcal{P}_i, i \geq 1\}$  of  $\mathcal{P}$  such that for all  $\eta > 0$ ,  $(\mathcal{P}_i, \ell)$  is  $\eta$ -predictable.*

**PROOF.** Applying Theorem 1 to  $(\mathcal{P}, \ell)$ , we know that if the universal nesting exists, then  $(\mathcal{P}, \ell)$  is e.a.s.-predictable. Theorem 1 also guarantees that if  $(\mathcal{P}, \ell)$  is e.a.s.-predictable, then for all  $\eta > 0$ , there is a nesting  $\{\mathcal{P}_i^\eta\}$  where  $(\mathcal{P}_i^\eta, \ell)$  is  $\eta$ -predictable.

We prove the theorem by showing that if there exists some  $\eta > 0$  for which there is a nesting  $\{\mathcal{P}_i\}$  where  $(\mathcal{P}_i, \ell)$  is  $\eta$ -predictable, then this nesting is also universal for all  $\eta > 0$ , i.e.  $(\mathcal{P}_i, \ell)$  is also  $\eta$ -predictable for all  $\eta > 0$ .

To do so, we show that if  $(\mathcal{P}_i, \ell)$  is  $(\frac{1}{|\mathcal{Y}|} - \epsilon)$ -predictable for any  $\epsilon > 0$  then it is  $\eta$ -predictable for all  $\eta > 0$ . Suppose  $N$  is the sample size such that a predictor  $\Phi$  makes no errors past  $N$  with probability  $\geq 1 - \frac{1}{|\mathcal{Y}|} + \epsilon$ . For any  $\eta < \frac{1}{|\mathcal{Y}|} - \epsilon$ , let  $M = \frac{2 \log(\eta/2)}{\epsilon^2}$  and consider a sample of size  $MN$ . We split this sample into  $M$  blocks of size  $N$  each, and apply the predictor  $\Phi$  to each block, obtaining a prediction  $Y_i \in \mathcal{Y}$  for the  $i$ 'th block. Our prediction for the sample of size  $MN$  is an element of  $\mathcal{Y}$  that repeats most often in  $Y_1, \dots, Y_M$ .

By Hoeffding bound with probability  $\geq 1 - \eta$ , any element of  $\mathcal{Y}$  that incurs loss 1 against the underlying distribution will appear at most  $\frac{1}{|\mathcal{Y}|} - \epsilon/2$  frequency among  $Y_1, \dots, Y_M$ . But at least one element of  $\mathcal{Y}$  appears more than  $1/|\mathcal{Y}|$  frequency among  $Y_1, \dots, Y_M$ , and any such an element must incur 0 loss. The theorem follows.  $\square$

REMARK 2. *Note that the finiteness of the prediction domain  $\mathcal{Y}$  is essential for the proof of the above Theorem to work. We leave it as an open problem to determine if the theorem above can be extended to countable domain  $\mathcal{Y}$  and general loss  $\ell : \mathcal{P} \times \mathcal{Y} \rightarrow \{0, 1\}$ .*

**5. Capturing the final error.** While *e.a.s.*-predictability is an attractive setup when considering rich model classes, we would like to see if a predictor that makes finitely many errors has finished making the errors. Namely, can we obtain a stopping rule that identifies the last error? Recall that a stopping rule is a function  $\tau : \mathcal{X}^* \rightarrow \{0, 1\}$ , such that  $\tau(y) \leq \tau(x)$  if  $y$  is a prefix of  $x$ . We interpreted  $\tau = 0$  as the waiting period, and  $\tau = 1$  as the rule has stopped waiting.

DEFINITION 6 (*e.a.s.*-learnable). *A collection  $(\mathcal{P}, \ell)$  is said to be eventually almost surely (*e.a.s.*)-learnable, if for any  $\eta > 0$ , there exists a universal prediction rule  $\Phi_\eta$  together with a stopping rule  $\tau_\eta$ , such that for all  $p \in \mathcal{P}$*

$$p \left( \sum_{n=1}^{\infty} \ell(p, X_1^n, \Phi_\eta(X_1^{n-1})) \tau_\eta(X_1^{n-1}) > 0 \right) < \eta,$$

and

$$p \left( \lim_{n \rightarrow \infty} \tau_\eta(X_1^n) = 1 \right) = 1.$$

Clearly, *e.a.s.*-learnability implies *e.a.s.*-predictability.

THEOREM 5. *Any *e.a.s.*-learnable  $(\mathcal{P}, \ell)$  is *e.a.s.*-predictable.*

PROOF. Suppose  $(\mathcal{P}, \ell)$  is *e.a.s.*-learnable. Then for each  $i$ , we let  $\Phi_i$  and  $\tau_i$  be the predictor and stopping rule pair respectively that *e.a.s.*-learns  $(\mathcal{P}, \ell)$  with  $\eta = 1/2^i$ . By definition, we have that the probability  $\Phi_i$  makes an error after  $\tau_i$  stops (i.e.  $\tau_i = 1$ ) is  $\leq \frac{1}{2^i}$ . Let  $\Phi_0$  be an arbitrary predictor.

Now, there are countably many stopping rules (one for each natural number  $i \geq 0$ ) and each such rule stops at a finite time with probability 1, we conclude that with probability 1 all of them would have stopped simultaneously at some finite time by a union bound.

We initialize  $t = 1$  ( $t$  will stand for the stage). As we see more of the sample, at any stage  $t$ , we predict using the prediction rule  $\Phi_{t-1}$ , till  $\tau_t$  halts (i.e.  $\tau_t = 1$ ). At that point, we move to stage  $t + 1$ . For  $t \geq 2$ , the probability of making an error in stage  $t$  is  $\leq 2^{-t}$ . Invoking the Borel-Cantelli lemma, we conclude that we make errors in finitely many stages almost surely, and the Theorem follows.  $\square$

However, *e.a.s.*-predictability does not imply *e.a.s.*-learnability.

EXAMPLE 6. *Let  $\mathcal{P}$  be the class of all iid Bernoulli random processes with parameters in  $[0, 1]$ . For any  $p \in \mathcal{P}$ , we would like to determine if the parameter of the process equals 1/2 or not (prediction is 1 if parameter is 1/2, 0 else), namely, the loss is  $\ell(p, X_1^n, Y_n) = 1\{Y_n = 1\{p = \frac{1}{2}\}\}$ , where in a slight abuse of notation, we use  $p$  to denote both the iid Bernoulli*

source and its parameter. In Example 2, we have shown that  $(\mathcal{P}, \ell)$  is *e.a.s.-predictable*. We now show that  $(\mathcal{P}, \ell)$  is not *e.a.s.-learnable*.

Suppose otherwise, let  $\Phi, \tau$  be the prediction rule and stopping rule that *e.a.s.-learns*  $(\mathcal{P}, \ell)$  with  $\eta = \frac{1}{4}$ . Consider the Bernoulli 1/2 source. Since  $\tau$  stops finitely almost surely on all sources, there exists a number  $N$  such that  $\tau$  has stopped before step  $N$  w.p.  $\geq \frac{3}{4}$  when the Bernoulli 1/2 source is in force.

Let  $A$  be the event that  $\tau$  has stopped before  $N$  and  $\Phi$  takes value 1 at step  $N + 1$ . Observe that  $p(A) \geq \frac{1}{2}$  by union bound. Let  $p'$  be any source other than the Bernoulli 1/2 source such that  $\|p' - p\|_{TV} \leq \frac{1}{8N}$ . We have  $p'(A) \geq \frac{1}{2} - \frac{1}{8} > \frac{1}{4}$  by Lemma 10. But  $\Phi$  incurs an error at step  $N + 1$  on any sequence in  $A$  whenever  $p'$  is in force, which contradicts that  $\Phi$  incurs errors after the stopping rule has halted on a set of probability  $< 1/4$ .

5.1. *Unions of learnable classes.* Even finite unions of *e.a.s.-learnable* classes need not be *e.a.s.-learnable*, as the stopping rule for one class need not even stop with probability 1 on sources of the other. It is more interesting to consider nested unions of classes. Clearly finite nested unions of *e.a.s.-learnable* classes are trivially *e.a.s.-learnable*, but countable unions of nested *e.a.s.-learnable* classes may not necessarily be *e.a.s.-learnable*. To see this, consider the subclasses  $\mathcal{P}_k$  we constructed in Example 2. We know that each of the  $\mathcal{P}_k$  is *e.a.s.-learnable* but  $\bigcup_{k \geq 1} \mathcal{P}_k$  is not *e.a.s.-learnable* per Example 6 above.

The following theorem however shows that a countable union of nested *e.a.s.-learnable* classes is always *e.a.s.-predictable*.

**THEOREM 6.** *Let  $\mathcal{P}$  be a class of distributions, and  $\{\mathcal{P}_i, i \geq 1\}$  be a nesting of  $\mathcal{P}$ . If for all  $i \geq 1$ ,  $(\mathcal{P}_i, \ell)$  is *e.a.s.-learnable*, then  $(\mathcal{P}, \ell)$  is *e.a.s.-predictable*.*

**PROOF.** The proof is similar to the proof of Lemma 1. For any  $k, i$ , let  $\Phi_i^k$  and  $\tau_i^k$  be the prediction rule and stopping rule respectively that achieves *e.a.s.-learnability* for  $(\mathcal{P}_i, \ell)$  with parameter  $\eta = \frac{1}{k^2 2^i}$ . We now construct the *e.a.s.-prediction* rule for  $(\mathcal{P}, \ell)$  as follows. We partition the prediction into phases, at the beginning, we are at phase 0 and the prediction is arbitrary. We will go from phase  $k - 1$  to phase  $k$  if at least one of  $\tau_i^k$  with  $i \geq k$  has stopped. Denote  $I_k$  to be the index of the stopping rule in  $\{\tau_i^k\}_{i \geq k}$  that stops earliest during phase  $k - 1$ . We will use the prediction rule  $\Phi_{I_k}^k$  to make the prediction during phase  $k$ . Suppose the underlying distribution is  $p \in \mathcal{P}_j$  for some  $j$ . We know that for each phase  $k \geq 1$  w.p. 1 there must be some stopping rule in  $\{\tau_i^k\}_{i \geq 1}$  with  $i \geq j$  that stops finitely almost surely. Using a union bound, we therefore conclude that all the phases will be finite simultaneously almost surely. Now, by definition of *e.a.s.-learnability*, the probability of making errors at phase  $k \geq j$  is upper bounded by  $\frac{1}{k^2}$  using a union bound. By the Borel-Cantelli lemma, we will make errors in finitely many phases almost surely, proving that  $\mathcal{P}$  is *e.a.s.-predictable*.  $\square$

**REMARK 3.** *Note that the nesting property is required for Theorem 6 to hold, that is, it is possible that countable unions of non-nested *e.a.s.-learnable* classes are not *e.a.s.-predictable*. To see this, let  $r_1, r_2, \dots$  be an arbitrary enumeration of rationals in  $[0, 1]$ . We denote  $\mathcal{B}_{i,j} = \{x \in [0, 1] : x \text{ is irrational and } |x - r_i| \geq \frac{1}{j}\}$ , and  $\mathcal{P}_{i,j}$  be the class of iid Bernoulli processes with parameters in  $\mathcal{B}_{i,j}$ . Denote  $\mathcal{P} = \bigcup_{i,j \geq 1} \mathcal{P}_{i,j}$ . Let  $\ell$  be the rationality testing loss of Cover [8] as we introduced in Example 3. It is easy to see that  $(\mathcal{P}_{i,j}, \ell)$  is *e.a.s.-learnable* (in fact it is  $\eta$ -predictable for all  $\eta > 0$ ). Moreover, the stopping rule for each  $\mathcal{P}_{i,j}$  stops finitely almost surely on all sources in  $\mathcal{P}$ . However,  $(\mathcal{P}, \ell)$  is not *e.a.s.-predictable*. This follows from Example 3 and the fact that the set of irrational numbers in  $[0, 1]$  is not of Baire first category [22, Chapter 2.1].*

5.2. *Characterization of e.a.s.-learnability.* The characterization of *e.a.s.-learnability* is captured by the notion of *identifiability* as follows.

DEFINITION 7. *Let  $\mathcal{U}$  be a collection of probability measures over  $\mathcal{X}^\infty$ ,  $\mathcal{V} \subset \mathcal{U}$ . The class  $\mathcal{V}$  is said to be *identifiable* in  $\mathcal{U}$  if for any  $\eta > 0$  there exists a stopping rule  $\tau_\eta$ , such that*

1.  $p \left( \lim_{n \rightarrow \infty} \tau_\eta(X_1^n) = 1 \right) = 1$  for  $p \in \mathcal{V}$ ;
2.  $p \left( \lim_{n \rightarrow \infty} \tau_\eta(X_1^n) = 1 \right) \leq \eta$  for  $p \in \mathcal{U} \setminus \mathcal{V}$ .

In other words, the rule almost surely stops on sources in  $\mathcal{V}$ , but does not stop on sources in  $\mathcal{U} \setminus \mathcal{V}$  with the prescribed confidence. If we want to distinguish whether a source is in  $\mathcal{V}$  (the null hypothesis) or in  $\mathcal{U} \setminus \mathcal{V}$  (alternate), we can do so with asymptotically zero type-I error and arbitrarily small type-II error, hence the term *identifiability*.

EXAMPLE 7. *Let  $\mathcal{U}$  be the collection of all iid processes with marginal distributions over  $[0, 1]$ , and let  $\mathcal{V} \subset \mathcal{U}$  be the set of distributions whose marginal mean is not equal to  $t$  for some fixed  $t \in [0, 1]$ .*

*We show  $\mathcal{V}$  is identifiable in  $\mathcal{U}$ . To see this, let  $\epsilon_n = \frac{1}{n}$ . Consider the following stopping rule. At stage  $n$ , we obtain a sample of size  $\frac{2 \log(2^{n+1}/\eta)}{\epsilon_n^2}$  and check whether the empirical mean is within  $\epsilon_n$  distance of  $t$ . If not, we stop, else we continue to stage  $n + 1$ .*

*We show that this stopping rule identifies  $\mathcal{V}$  in  $\mathcal{U}$  using Definition 7. Suppose the underlying process has marginal mean equal to  $t$ . By Hoeffding bound, with probability at most  $\eta/2^n$ , the empirical mean will be outside distance  $\epsilon_n$  to  $t$ . Therefore, the stopping rule stops with probability at most  $\eta$  by a union bound. If the marginal mean does not equal  $t$ , since  $\epsilon_n \rightarrow 0$ , the probability that the empirical mean will be within distance  $\epsilon_n$  to  $t$  is at most  $\frac{\eta}{2^n}$ . By Borel-Cantelli lemma, this happens only finitely many times since  $\sum \frac{\eta}{2^n} < \infty$ , and the stopping rule stops almost surely.*

The following theorem shows that for *iid* processes over  $\mathbb{N}$ , *identifiability* is equivalent to *relative openness* under total variation distance. See Appendix C.1 for details on the proof.

THEOREM 7. *Let  $\mathcal{U}$  be a collection of i.i.d. processes over  $\mathbb{N}^\infty$ ,  $\mathcal{V} \subset \mathcal{U}$ . Then  $\mathcal{V}$  is identifiable in  $\mathcal{U}$  iff the marginals of  $\mathcal{V}$  are relatively open in the marginals of  $\mathcal{U}$  under total variation distance.*

REMARK 4. *Note that Theorem 7 cannot be generalized to the iid processes over  $\mathbb{R}^\infty$ . To see this, let  $\mathcal{A}$  be the collection of all uniform distributions supported on a finite set of  $[0, 1]$ , and  $\mathcal{V}$  contains the single uniform distribution  $p$  over  $[0, 1]$ . Define  $\mathcal{U} = \mathcal{A} \cup \mathcal{V}$ . Clearly,  $\mathcal{V}$  is relatively open in  $\mathcal{U}$  under total variation distance. We show that  $\mathcal{V}$  is not identifiable in  $\mathcal{U}$ . Suppose otherwise, let  $\tau_\eta$  be a stopping rule as in Definition 7. Let  $N$  be a number such that  $\tau_\eta$  stops on  $p$  before step  $N$  w.p.  $\geq 1 - \eta/4$ . Consider a mixture  $\tilde{q}$  of  $\mathcal{A}$  as follows, we first sample a set  $S \subset [0, 1]$  of size  $CN^2$  from  $p$ , then we sample from  $S$  uniformly, where  $C$  is a constant. By birthday paradox, for sufficient large  $C$ , w.p.  $\geq 1 - \eta/4$ , there will be no repeats on a sample of size  $N$  from  $\tilde{q}$ . Conditioning on the event that there is no repeats, the statistic of a sample of size  $N$  from  $\tilde{q}$  will be the same as the sample of size  $N$  from  $p$ . By union bound, there must be some distribution in  $q \in \mathcal{A}$  such that  $\tau_\eta$  stops on  $q$  before step  $N$  w.p.  $\geq 1 - \frac{\eta}{2} > \eta$  for all  $\eta < 1/2$ . Which contradicts to Definition 7.*

We now provide the following characterization of *e.a.s.*-learnability. We leave the details of the proof to Appendix C.1.

**THEOREM 8.** *Let  $\mathcal{P}$  be a collection of probability measures. Then  $(\mathcal{P}, \ell)$  is *e.a.s.*-learnable if for all  $\eta > 0$  there exists an  $\eta$ -nesting  $\{\mathcal{P}_i, i \geq 1\}$ , where in addition for all  $i$ ,  $\mathcal{P}_i^\eta$  is identifiable in  $\mathcal{P}$ . Moreover, the condition is necessary if the measures in  $\mathcal{P}$  are iid over  $\mathcal{X}^\infty$ .*

The tight characterization for *iid* sources in Theorem 8 can not be extended to arbitrary sources. We provide an example below that shows that the characterization in Theorem 8 for *iid* sources will not extend to Markov processes with even two states.

**EXAMPLE 8.** *We consider the Markov processes with state space  $\{0, 1\}$ . Let  $\mathcal{P}$  be the class that contains the single state 1 process  $p_0$ , and processes  $p_\epsilon$  with transition probability  $p_\epsilon(1|0) = p_\epsilon(0|1) = \epsilon$  for all  $\epsilon \in (0, 1)$ . We assume the initial state of  $p_\epsilon$  to be uniformly sampled from  $\{0, 1\}$ .*

*We define the loss  $\ell(p, X_1^n, Y_n) = 0$  if there exists  $k \leq n$  such that  $X_k = 1$ . Else the loss is 1. Note that the loss only depends on the samples  $X_1^n$  but is independent of the prediction  $Y_n$ . Thus the prediction does not affect the loss.*

*We now observe that the class is *e.a.s.*-learnable. One simply stops if the initial state is 1, else we wait until we see 1.*

*We now show that the decomposition of Theorem 8 does not exist for  $(\mathcal{P}, \ell)$ . Suppose otherwise, we have a decomposition  $\{\mathcal{P}_n\}_{n \geq 1}$  such that each  $(\mathcal{P}_n, \ell)$  is  $1/4$ -predictable. We know that for all  $n \geq 1$  there exists a number  $\epsilon_n$  such that for all  $p_\epsilon \in \mathcal{P}_n$  we have  $\epsilon \geq \epsilon_n$ . Otherwise, we will not see state 1 in the sample before a bounded time step if the initial state is 0 (which happens w.p.  $1/2$ ), thus violating the  $1/4$ -predictability of  $\mathcal{P}_n$ . We now assume  $p_0 \in \mathcal{P}_k$  for some  $k$ . We show that  $\mathcal{P}_k$  is not identifiable in  $\mathcal{P}$ . Taking the parameter  $\eta = 1/4$  in Definition 7, we know that any  $\tau$  must stop on the all 1 sequence at some point  $N_0$ , since  $p_0 \in \mathcal{P}_k$ . Now, taking any process  $p_\epsilon$  that is not in  $\mathcal{P}_k$  with  $\epsilon < \epsilon_k$  and small enough so that  $(1 - \epsilon)^{N_0} \geq 3/4$ , we have  $\tau$  stops on  $p_\epsilon$  with probability at least  $3/8 > 1/4$ , contradicting identifiability.*

Note that the reason why construction of Example 8 is possible is because the number of states of the process is not known *a-priori* (either 1 or 2). Indeed, with arguments similar to the necessity part of Theorem 8 in the *iid* case, we can show that if  $\mathcal{P}$  is a collection of irreducible finite state Markov processes with the same number of states then the necessary condition in Theorem 8 still holds. We should also emphasize that the stopping rule derived from the sufficient condition of Theorem 8 is not meant to be optimal. For specific problems, there will often be more natural stopping rules.

**6. Variations.** In addition to *e.a.s.*-predictability and *e.a.s.*-learnability that we established in the previous sections, we define the natural extensions of weakly *e.a.s.*-predictability and prediction with finite expected loss in this section, and study their relationships with *e.a.s.*-predictability.

### 6.1. Weakly *e.a.s.*-predictable.

DEFINITION 8. A collection  $(\mathcal{P}, \ell)$  is said to be weakly *e.a.s.-predictable* if for any  $\eta > 0$  there exists a prediction rule  $\Phi^\eta$ , such that for all  $p \in \mathcal{P}$

$$p \left( \sum_{n=1}^{\infty} \ell(p, X_1^n, \Phi^\eta(X_1^{n-1})) < \infty \right) \geq 1 - \eta.$$

It is easy to see that the characterization in Theorem 1 holds for weakly *e.a.s.-predictability* as well.

THEOREM 9. If a collection  $(\mathcal{P}, \ell)$  is weakly *e.a.s.-predictable* then for all  $\eta > 0$  there exists a  $\eta$ -nesting of  $\mathcal{P}$ .

Conversely, if there exists a universal nesting of  $\mathcal{P}$ , then  $(\mathcal{P}, \ell)$  is weakly *e.a.s.-predictable*.

PROOF. The sufficiency follows immediately from Theorem 1, since *e.a.s.-predictability* implies weakly *e.a.s.-predictability*. To prove the necessary condition, one could replicate the same argument as in the proof of Theorem 1.  $\square$

The following example shows that weakly *e.a.s.-predictable* does not imply *e.a.s.-predictable*.

EXAMPLE 9. For any  $m \in \mathbb{N}$  we define a set  $S_m \subset \mathbb{N}$  such that  $|S_m| = m$ , and we require  $S_m \cap S_{m'} = \emptyset$  for all  $m \neq m' \in \mathbb{N}$ . Clearly, such a family of sets exists. Now, for any sequences  $\mathbf{a} = (a_1, a_2, \dots) \in \mathbb{N}^\infty$  and  $\mathbf{b} = (b_1, b_2, \dots) \in \mathbb{N}^\infty$  with  $b_m \in S_m$  for all  $m \in \mathbb{N}$ , we define a source  $p^{(\mathbf{a}, \mathbf{b})}$  as follows.

1. The process  $p^{(\mathbf{a}, \mathbf{b})}$  has no outcome, i.e. the learner can't observe any sample;
2. The loss  $\ell(p^{(\mathbf{a}, \mathbf{b})}, X_1^n, Y_n) = 0$  if and only if there exist some  $m \in \mathbb{N}$  such that  $n \geq a_m$ ,  $Y_n \in S_m$  and  $Y_n \neq b_m$ .

Let  $\mathcal{P}$  be the class of all such sources. We now show that there exist randomized prediction rules that achieve weakly *e.a.s.-predictability* for  $(\mathcal{P}, \ell)$ . For any  $\eta > 0$ , the predictor  $\Phi^\eta$  is defined as follows. Let  $M$  be a number such that  $M \geq \frac{1}{\eta}$ . We select a number  $c \in S_M$  uniform randomly at the beginning of the game, and define  $\Phi^\eta(X_1^{n-1}) = c$  for all  $n \geq 1$ . Clearly, there exists some time step  $N$  such that  $N \geq a_M$ , and the probability that we make errors after step  $N$  is at most  $\frac{1}{M} \leq \eta$ .

We now show that any randomized predictor  $\Phi$  can't achieve *e.a.s.-predictability* for  $(\mathcal{P}, \ell)$ . Let  $\Omega$  be the probability space of the internal randomness of  $\Phi$ , i.e., for any  $\omega \in \Omega$ ,  $\Phi^\omega$  is a deterministic strategy. Let  $\mu$  be the corresponding probability measure over  $\Omega$ .

Let  $c_1, c_2, \dots$  be a sequence of numbers such that the probability of  $\Phi$  making predictions in  $\bigcup_{i \leq m} S_i$  after step  $c_m$  is upper bounded by  $\frac{\eta}{2^m}$  for some  $\eta < 1$ . We show that such sequences exist. Otherwise, there exists some  $m$  such that  $\Phi$  makes predictions on some number  $s \in S_m$  infinite often with positive probability. However, this will violate the *e.a.s.-predictability* of  $\Phi$ , since it will make infinite errors for the source that takes  $b_m = s$ .

Let  $A$  be the event that,  $\Phi$  makes no prediction on  $\bigcup_{i \leq m} S_i$  after step  $c_m$  for all  $m \in \mathbb{N}$ . By union bound, we have  $\mu(A) \geq 1 - \eta$ . We now take  $a_m = c_m$  for all  $m \in \mathbb{N}$ . And let  $b_m$  to be choosing uniform randomly from  $S_m$  and independently for different  $m$ . Denote  $\Gamma$  to be the probability space of the selection processes, and  $\nu$  be the corresponding probability measure. For any  $\gamma \in \Gamma$  we denote  $p^\gamma$  to be the source that is selected from the process.

For any  $(\omega, \gamma) \in \Omega \times \Gamma$ , we define  $1\{\omega, \gamma\}$  to be the indicator that  $\Phi^\omega$  makes infinite errors on  $p^\gamma$ . We claim that for any  $\omega \in A$  we have

$$\mathbb{E}_{\gamma \sim \nu} 1\{\omega, \gamma\} = 1.$$

W.o.l.g., we can assume that  $\Phi^\omega$  makes predictions only on  $S_{m+1}$  during phase  $m$  of steps from  $a_m + 1$  to  $a_{m+1}$ . Now, we have the probability of making errors at phase  $m$  is at least  $\frac{1}{m}$ . Since the  $b_m$  is selected independently, we know that the errors at different phases are independent. The claim follows by the converse Borel-Cantelli lemma.

Now, we have shown that

$$\mathbb{E}_{\omega \sim \mu} [\mathbb{E}_{\gamma \sim \nu} 1\{\omega, \gamma\}] \geq 1 - \eta.$$

By switching order of expectation, we have

$$\mathbb{E}_{\gamma \sim \nu} \mathbb{E}_{\omega \sim \mu} 1\{\omega, \gamma\} \geq 1 - \eta.$$

This implies that there exists some  $\gamma$  such that the probability of  $\Phi$  making infinite errors on  $p^\gamma$  is at least  $1 - \eta > 0$ . This violates the e.a.s.-predictability of  $\Phi$ .

REMARK 5. The above example shows that if the prediction rule is randomized, then weakly e.a.s.-predictability does not imply e.a.s.-predictability. However, it is still open whether such a separation exists if the prediction rule is required to be deterministic.

## 6.2. Prediction with finite expected loss.

DEFINITION 9. A collection  $(\mathcal{P}, \ell)$  is said to be predictable with finite expected loss if there exists a prediction rule  $\Phi$ , such that for all  $p \in \mathcal{P}$

$$\mathbb{E}_p \left[ \sum_{n=1}^{\infty} \ell(p, X_1^n, \Phi(X_1^{n-1})) \right] < \infty.$$

By the Borel-Cantelli lemma, it is clear that if a collection is predictable with finite expected loss, then it is also e.a.s.-predictable. The following example shows that the converse does not hold in general.

EXAMPLE 10. Let  $\mathcal{P}$  be the class with a single source  $p$  defined as follows. Let  $U$  be the uniform distribution over  $[0, 1]$ , we define

$$\forall n \in \mathbb{N}, X_n = U,$$

i.e. all the  $X_n$  have the same value that is sampled from  $U$ . The loss is defined as follows

$$\ell(p, X_1^n, Y_n) = 1 \text{ if and only if } X_n \leq \frac{1}{n}.$$

Clearly,  $(\mathcal{P}, \ell)$  is e.a.s.-predictable, since with probability 1 we have  $U \neq 0$ , meaning that we will make no error after step  $\lceil \frac{1}{U} \rceil$ . However, the collection is not predictable with finite expected loss. To see this, we define  $A_n$  to be the event that an error occurred at step  $n$ . We have

$$\mathbb{E}_p \left[ \sum_{n=1}^{\infty} \ell(p, X_1^n, \Phi(X_1^{n-1})) \right] = \sum_{n=1}^{\infty} \mathbb{E}_p[A_n] = \sum_{n=1}^{\infty} \frac{1}{n} = \infty.$$



Clearly, if a collection  $(\mathcal{P}, \ell)$  is predictable with finite expected loss, then there exists a nesting  $\{\mathcal{P}_i, i \geq 1\}$  of  $\mathcal{P}$  such that  $\forall i \geq 1$   $(\mathcal{P}_i, \ell)$  is predictable with *bounded* expected loss. To do so, we can simply define

$$\mathcal{P}_i = \{p \in \mathcal{P} : p \text{ has expected loss } \leq i \text{ under prediction rule } \Phi\},$$

where  $\Phi$  is a prediction rule that achieves finite expected loss on  $(\mathcal{P}, \ell)$ . However, we are unaware if the converse is true or not. We have the following open problem.

**PROBLEM 1 (Open Problem).** *Suppose that there is a nesting  $\{\mathcal{P}_i, i \geq 1\}$  for a class  $\mathcal{P}$ , such that  $\forall i \geq 1$   $(\mathcal{P}_i, \ell)$  is predictable with bounded expected loss for the same loss function  $\ell$ . Is the collection  $(\mathcal{P}, \ell)$  predictable with finite expected loss?*

**REMARK 6.** *We say a collection  $(\mathcal{P}, \ell)$  is predictable with bounded expected tail loss if there exists a prediction rule  $\Phi$ , function  $\rho : \mathbb{N}^+ \rightarrow \mathbb{R}^+$  and constant  $N$ , such that for all  $p \in \mathcal{P}$  and  $n \geq N$ , we have*

$$\mathbb{E}_p \left[ \sum_{k=n}^{\infty} \ell(p, X_1^k, \Phi(X_1^{k-1})) \right] \leq \rho(n),$$

and  $\rho(n) \rightarrow 0$  as  $n \rightarrow \infty$ . It is not hard to show that **Problem 1** is true if we assume  $\forall i \geq 1$  the collection  $(\mathcal{P}_i, \ell)$  is predictable with bounded expected tail loss. This is satisfied, e.g., by the setting in **Theorem 4**.

Note that **Problem 1** holds for supervised setting. Let  $\mathcal{P}_1, \mathcal{P}_2, \dots$  be countably many classes such that each  $(\mathcal{P}_i, \ell)$  is predictable with finite expected loss using prediction rule  $\Phi_i$ . We can construct a prediction rule  $\Phi$  using the Randomized Weight Majority Algorithm [17, Section 4] with experts pool  $\{\Phi_1, \Phi_2, \dots\}$ . Suppose the underlying distribution  $p \in \mathcal{P}_k$  for some  $k$  with expected errors of  $E_i < \infty$ , [17, Section 4] shows that the expected errors of  $\Phi$  is upper bounded by  $E_i + \log k < \infty$ .

**7. Applications.** In this section we will characterize several concrete setups with the application of our general characterizations in previous sections.

**7.1. Hypothesis Testing.** For any two disjoint collections  $\mathcal{H}_1, \mathcal{H}_2$  of distributions over  $\mathbb{R}^d$ , the *hypothesis testing* problem is to decide whether some underlying source  $p \in \mathcal{H}_1 \cup \mathcal{H}_2$  is from  $\mathcal{H}_1$  or  $\mathcal{H}_2$ , by observing *iid* samples from  $p$ . Therefore, using the notations we developed in **Section 2**, the class under consideration will include the *iid* processes in  $\mathcal{H} = \mathcal{H}_1 \cup \mathcal{H}_2$ , and for any  $p \in \mathcal{H}$  the loss is defined as

$$\ell^{\text{HT}}(p, X_1^n, Y_n) = 1\{p \notin \mathcal{H}_{Y_n}\}.$$

With slight abuse of notation, for any distribution  $p$  over  $\mathbb{R}^d$ , we will use  $p$  to denote the *iid* process of  $p$  as well, and we denote  $p^n$  to be the  $n$ -fold *i.i.d.* distribution of  $p$ .

This problem was considered extensively in [8, 9]. Informally, Dembo and Peres [9] showed that the existence of *e.a.s.*-prediction rule for a hypothesis testing problem is *essentially* captured by the notion of what they call  $F_\sigma$ -separability.

Let  $A, B$  be two sets in some metric space with metric  $d$ . We say  $A, B$  are  $F_\sigma$ -separable if  $A \subset A'$  and  $B \subset B'$ ,  $A' \cap B' = \emptyset$  and both  $A'$  and  $B'$  can be written as countable unions of closed sets (with topology induced by  $d$ ). One of the main results of [9] is stated below using our notations in **Section 2**.

**THEOREM 10** ([9, Theorem 2]). *Let  $\mathcal{H}_1$  and  $\mathcal{H}_2$  be disjoint collections of distributions over  $\mathbb{R}^d$  that are absolutely continuous w.r.t. Lebesgue measure. Denote  $\mathcal{H} = \mathcal{H}_1 \cup \mathcal{H}_2$ . We have*

1. *If  $\mathcal{H}_1$  and  $\mathcal{H}_2$  are  $F_\sigma$ -separable under the weak convergence topology, then  $(\mathcal{H}, \ell^{\text{HT}})$  is e.a.s.-predictable.*
2. *If for any  $p \in \mathcal{H}$ , there exists  $r > 1$  (possibly depends on  $p$ ) such that*

$$\int_{\mathbb{R}^d} f_p^r(\mathbf{x}) d\mu < \infty,$$

*where  $f_p$  is the density of  $p$  and  $\mu$  is Lebesgue measure. Then  $(\mathcal{H}, \ell^{\text{HT}})$  is e.a.s.-predictable iff  $\mathcal{H}_1$  and  $\mathcal{H}_2$  are  $F_\sigma$ -separable under the weak convergence topology.*

Moreover, Dembo-Peres [9] asked whether the condition  $r > 1$  in Theorem 10(2) can be removed. We now give a positive answer to this open problem by showing that the condition of the finite  $r$ th norm of the density can be relaxed to a much weaker condition as defined in the following.

**DEFINITION 10.** *Let  $\mathcal{H}$  be a collection of distributions over  $\mathbb{R}^d$  that are absolutely continuous w.r.t. Lebesgue measure. We say  $\mathcal{H}$  is uniformly bounded if for any  $\epsilon > 0$  there exists a number  $M_\epsilon$  such that*

$$\forall p \in \mathcal{H}, p(f_p(\mathbf{x}) \geq M_\epsilon) \leq \epsilon,$$

*where  $f_p$  is the density of  $p$ .*

Before showing why such a condition is sufficient to establish Theorem 10(2), we first provide a simple alternative proof to Theorem 10(1) by using our general characterization in Section 3. To begin with, we introduce the following lemma, which relates  $F_\sigma$ -separability with the nesting of sets. The proof is left to Appendix C.2.

**LEMMA 3.** *Let  $A, B$  be two sets of some metric space with metric  $d$ . The  $A, B$  are  $F_\sigma$ -separable iff there exist nestings  $\{A_i, i \geq 1\}$  and  $\{B_i, i \geq 1\}$  for  $A, B$  respectively, such that*

$$\forall i \geq 1, \inf\{d(x, y) : x \in A_i, y \in B_i\} > 0.$$

**REMARK 7.** *Note that, with similar argument as in the proof of Lemma 3 (see Appendix C.2), we can also prove that two sets  $A, B$  are  $F_\sigma$ -separable iff there exist nestings  $\{A_i, i \geq 1\}$  and  $\{B_i, i \geq 1\}$  for  $A, B$  respectively, such that for all  $i \geq 1$  there is no limit point of  $A_i$  in  $B_i$  (and vice versa).*

The following lemma will be useful in our following analysis.

**LEMMA 4.** *Let  $\{\{A_{i,j}\}_{i \in \mathbb{N}}, \{B_{i,j}\}_{i \in \mathbb{N}}\}_{j \in \mathbb{N}}$  be nesting such that  $A_{i,j} \subset A_{i',j'}$  and  $B_{i,j} \subset B_{i',j'}$  whenever  $i' \geq i$  and  $j' \geq j$ . If  $d(A_{i,j}, B_{i,j}) > 0$  for all  $i, j \in \mathbb{N}$ , where  $d$  is some metric. Then  $A = \bigcup_{i,j \in \mathbb{N}} A_{i,j}$  and  $B = \bigcup_{i,j \in \mathbb{N}} B_{i,j}$  are  $F_\sigma$ -separable.*

**PROOF.** We only need to show that for all  $i, j$ ,  $B$  and  $A$  contain no limit point of  $A_{i,j}$  and  $B_{i,j}$  respectively. Suppose otherwise that there exists some  $y \in B$  such that  $d(y, A_{i,j}) = 0$ . We have  $y \in B_{i',j'}$  for some  $i', j'$ . By nesting property, we may assume  $i = i'$  and  $j = j'$ . However, this will imply  $d(B_{i,j}, A_{i,j}) = 0$ , a contradiction.  $\square$

For any two distributions  $p_1, p_2$  over  $\mathbb{R}^d$ , the Kolmogorov-Smirnov distance (abbreviate as KS-distance) is defined as

$$|p_1 - p_2|_{KS} = |F_{p_1}(\mathbf{x}) - F_{p_2}(\mathbf{x})|_{KS} \stackrel{\text{def}}{=} \sup_{\mathbf{x} \in \mathbb{R}^d} |F_{p_1}(\mathbf{x}) - F_{p_2}(\mathbf{x})|,$$

where  $F_{p_i}$  is the CDF of  $p_i$  for  $i \in \{1, 2\}$ . The following lemma shows that the empirical distribution uniformly estimates any distribution under KS-distance, and is known as Dvoretzky-Kiefer-Wolfowitz Inequality, see e.g., [18, 11, 20].

**LEMMA 5 (Dvoretzky-Kiefer-Wolfowitz Inequality).** *Let  $X_1, X_2, \dots, X_n$  be iid samples of distribution  $p$  over  $\mathbb{R}^d$ ,  $F_n(\mathbf{x})$  is the CDF of the empirical distribution. Then there exists a constant  $C_d$  depends only on  $d$  such that*

$$p(|F_n(\mathbf{x}) - F_p(\mathbf{x})|_{KS} \geq \epsilon) \leq C_d \exp(-n\epsilon^2).$$

Note that the convergence under KS-distance is consistent with weak convergence under the assumption of absolutely continuous w.r.t. Lebesgue measure by Polya's theorem (see, e.g., [2, Theorem 9.1.4]).

**PROOF OF THEOREM 10(1).** By Lemma 3, we have nestings  $\{A_i, i \geq 1\}, \{B_i, i \geq 1\}$  of  $\mathcal{H}_1, \mathcal{H}_2$  respectively, such that

$$\forall i \in \mathbb{N}, \inf\{|p_1 - p_2|_{KS} : p_1 \in A_i, p_2 \in B_i\} \geq \frac{1}{i}.$$

By Theorem 1, we only need to show that  $(A_i \cup B_i, \ell^{\text{HT}})$  is  $\eta$ -predictable for all  $i \geq 1$  and  $\eta > 0$ . Let  $F_n$  be the CDF of the empirical distribution with sample of size  $n$ . By Lemma 5, we can simultaneously make  $|F_n(\mathbf{x}) - F_p(\mathbf{x})|_{KS} \leq 1/4i$  with confidence  $\geq 1 - \eta$  for all  $p \in A_i \cup B_i$  by choosing the sample size  $n$  large enough. By triangle inequality of KS-distance, one can classify the distributions in  $A_i \cup B_i$  successfully with probability at least  $1 - \eta$ , by predicting the class that is closer to  $F_n$  under KS-distance.

The theorem follows by observing that convergence under KS-distance is consistent with weak convergence with the absolutely continuous assumption.  $\square$

We now show that if the class  $\mathcal{H}_1 \cup \mathcal{H}_2$  is uniformly bounded as defined in Definition 10, then the  $F_\sigma$ -separability will be necessary to achieve *e.a.s.*-predictability under the assumption of Theorem 10.

**THEOREM 11.** *Let  $\mathcal{H}_1, \mathcal{H}_2$  be collections of distributions that are absolutely continuous w.r.t. Lebesgue measure on  $\mathbb{R}^d$ , and  $\mathcal{H}_1 \cup \mathcal{H}_2$  is uniformly bounded. Then  $(\mathcal{H}_1 \cup \mathcal{H}_2, \ell^{\text{HT}})$  is *e.a.s.*-predictable if and only if  $\mathcal{H}_1, \mathcal{H}_2$  are  $F_\sigma$ -separable under KS-distance.*

**PROOF.** The sufficiency follows directly from Theorem 10(1). We now only prove the necessity.

By Theorem 1, *e.a.s.*-predictability implies that there exist nesting  $\{A_i, i \geq 1\}, \{B_i, i \geq 1\}$  of  $\mathcal{H}_1, \mathcal{H}_2$  respectively, such that for all  $i \geq 1$ ,  $(A_i \cup B_i, \ell^{\text{HT}})$  is  $\frac{1}{8}$ -predictable. By Lemma 3, it is sufficient to show that for all  $i \geq 1$ , there is no limit point of  $A_i$  in  $B_i$  or vice versa. Suppose otherwise, there exist  $p_1, p_2, \dots \in A_i$  and  $p \in B_i$ , such that  $|p_n - p|_{KS} \rightarrow 0$  as  $n \rightarrow \infty$ . Let  $\Phi$  be an arbitrary predictor that achieves  $\frac{1}{8}$ -predictability of  $A_i \cup B_i$  with sample size  $i$ . Denote  $\Phi^i$  as the prediction function at step  $i$ . By absolute continuity,  $p_n$  weakly converges to  $p$ . Therefore, there exists a compact set  $S \subset \mathbb{R}^{i \cdot d}$ , such that  $p(S) \geq \frac{7}{8}$  and  $p_n(S) \geq \frac{7}{8}$

for all  $n \in \mathbb{N}$  by Prokhorov's theorem (see [26, Theorem 5.2]), where we have used  $p, p_n$  to also denote  $i$ -fold measures of  $p, p_n$  over  $\mathbb{R}^{i \cdot d}$ . By uniform boundedness of  $\mathcal{H}_1 \cup \mathcal{H}_2$ , there exists a number  $M$ , such that

$$(7.1) \quad \forall p \in \mathcal{H}_1 \cup \mathcal{H}_2, p(\{\mathbf{x} \in \mathbb{R}^{i \cdot d}, f_p(\mathbf{x}) \geq M\}) \leq \frac{1}{16}.$$

By Lusin's theorem [21, Theorem 2.24], there exists a continuous function  $g$  over  $\mathbb{R}^{i \cdot d}$  and a set  $E \subset S$ , such that  $\sup_{\mathbf{x} \in E} |\Phi^i(\mathbf{x}) - g(\mathbf{x})| \leq \frac{1}{4}$  and  $\mu(S \setminus E) \leq \frac{1}{16M}$  where  $\mu(\cdot)$  is the Lebesgue measure over  $\mathbb{R}^{i \cdot d}$ . Let  $\Omega = \{\mathbf{x} \in \mathbb{R}^{i \cdot d} : g(\mathbf{x}) > \frac{3}{2}\}$ , we have  $\Omega$  is open and  $\{\mathbf{x} \in E : \Phi^i(\mathbf{x}) = 2\} \subset \Omega$ . By (7.1) and  $\mu(S \setminus E) \leq \frac{1}{16M}$ , we have  $p(S \setminus E) \leq \frac{1}{8}$  and  $p_n(S \setminus E) \leq \frac{1}{8}$  for all  $n \in \mathbb{N}$ . By  $\frac{1}{8}$ -predictability, we have  $p(\Omega) \geq p(\Omega \cap E) \geq \frac{7}{8} - \frac{1}{4} = \frac{5}{8}$ , since  $p(\bar{E}) \leq \frac{1}{4}$ . By weak convergence, we have  $\liminf p_n(\Omega) \geq p(\Omega)$ , since  $\Omega$  is open (see [26, Theorem 3.2]). There exists some  $p_n$  such that  $p_n(\Omega) \geq \frac{1}{2}$ , which implies  $p_n(\Omega \cap E) \geq \frac{1}{2} - \frac{1}{4} = \frac{1}{4} > \frac{1}{8}$ , contradicting the  $\frac{1}{8}$ -predictability.  $\square$

REMARK 8. Note that, the assumption of uniform boundedness on  $\mathcal{H}_1 \cup \mathcal{H}_2$  is necessary for the proof of Theorem 11 to work. To see this, let  $p'_n$  be the uniform distribution over  $\{\frac{1}{n}, \frac{2}{n}, \dots, \frac{n-1}{n}\}$ , and  $p$  be the uniform distribution over  $[0, 1]$ . Clearly, we have that  $p'_n$  converges to  $p$  in distribution. But  $p'_n$  is not continuous. This can be easily resolved by choosing some continuous distribution  $p_n$  such that  $|p_n - p'_n|_{KS} \leq \frac{1}{n}$  and there exists a set  $S_n$  concentrated on  $\{\frac{1}{n}, \frac{2}{n}, \dots, \frac{n-1}{n}\}$  so that we have  $p_n([0, 1] \setminus S_n) \leq \frac{\eta}{2^n}$  and  $p(S_n) \leq \frac{\eta}{2^n}$ . Let  $A = \{p_1, p_2, \dots\}$  and  $B = \{p\}$ , we have  $p_n$  converges to  $p$  in distribution. However, since

$$\sum_{n=1}^{\infty} p(S_n) \leq \eta$$

we have  $(A \cup B, \ell^{\text{HT}})$  is  $\eta$ -predictable with only one sample by predicting the underlying distribution in  $A$  if the sample is in  $\bigcup S_n$  and in  $B$  otherwise.

We prove the following corollary, which establishes Corollary 1.

COROLLARY 2. Let  $G : \mathbb{R}^+ \rightarrow \mathbb{R}^+$  be a monotone increasing function such that  $\lim_{x \rightarrow \infty} G(x) \rightarrow \infty$ ,  $\mathcal{H}_1, \mathcal{H}_2$  be distributions over  $\mathbb{R}^d$  that are absolutely continuous w.r.t. Lebesgue measure. If for all  $p \in \mathcal{H}_1 \cup \mathcal{H}_2$ , we have  $\mathbb{E}_{\mathbf{x} \sim p}[G(f_p(\mathbf{x}))] < \infty$ , where  $f_p$  is the density of  $p$ . Then  $(\mathcal{H}_1 \cup \mathcal{H}_2, \ell^{\text{HT}})$  is e.a.s.-predictable iff  $\mathcal{H}_1, \mathcal{H}_2$  are  $F_\sigma$  separable with KS-distance.

PROOF. By Theorem 10 we only need to prove the necessary condition. By Lemma 3, Lemma 4 and breaking  $\mathcal{H}_1 \cup \mathcal{H}_2$  into countably subcollections, one may assume, w.l.o.g.,  $\forall p \in \mathcal{H}_1 \cup \mathcal{H}_2, \mathbb{E}_{\mathbf{x} \sim p}[G(f_p(\mathbf{x}))] \leq M$  for some constant  $M$ . By Theorem 11, we only need to show that  $\mathcal{H}_1 \cup \mathcal{H}_2$  is uniformly bounded. For any  $p \in \mathcal{H}_1 \cup \mathcal{H}_2$ , we define random variable  $Y_p = G(f_p(\mathbf{x}))$ . We have, by Markov inequality,  $p(Y_p \geq T) \leq \frac{M}{T}$ . Note that the upper bound is independent of  $p$ . By letting  $T = \frac{M}{\epsilon}$ , one can make the probability upper bounded by  $\epsilon$ . Since  $G$  is monotone increasing and goes to infinity, thus invertible on  $\mathbb{R}^+$ . We now have  $p(f_p(\mathbf{x}) \geq G^{-1}(M/\epsilon)) \leq \epsilon$  for all  $p \in \mathcal{H}_1 \cup \mathcal{H}_2$  and  $\epsilon > 0$ .  $\square$

REMARK 9. Note that, Theorem 10(2) follows directly from Corollary 2 by simply taking  $G(x) = x^{r-1}$  for some  $r > 1$ . However, the conditions in Corollary 2 are much weaker, since  $G(x)$  can be an arbitrary monotone function so long as  $G(x) \rightarrow \infty$  when  $x \rightarrow \infty$ . One may, e.g., take  $G(x) = \log(x)$ , which is asymptotically upper bounded by  $x^{r-1}$  for all  $r > 1$ .

We now have the following conjecture, which asserts that a condition such as the uniform boundedness is necessary for Theorem 10(2) to hold. See also Appendix A a *complete* characterization for distributions over  $\mathbb{N}$ .

CONJECTURE 1. *There exist collections  $\mathcal{H}_1, \mathcal{H}_2$  of distributions over  $\mathbb{R}^d$  that are absolutely continuous w.r.t. Lebesgue measure, such that  $(\mathcal{H}_1 \cup \mathcal{H}_2, \ell^{\text{HT}})$  is e.a.s.-predictable but  $\mathcal{H}_1, \mathcal{H}_2$  are not  $F_\sigma$ -separable under KS-distance.*

7.2. *The Insurance Problem.* We now consider a problem that is introduced in [23] and [30]. Let  $\mathcal{P}$  be a class of iid processes of distributions over  $\mathbb{N}$ . The loss function is defined as

$$\ell^{\text{IP}}(p, X_1^n, Y_n) = 1\{X_n > Y_n\},$$

i.e., we would like predict an upper bound for the next sample  $X_n$  by observing  $X_1, \dots, X_{n-1}$ . With slight abuse of notations, for any iid process  $p \in \mathcal{P}$  we will use  $p$  and  $X_p$  to denote the underlying marginal distribution and marginal random variable respectively as well.

We say a class  $\mathcal{P}$  of iid processes over  $\mathbb{N}$  is tight if for any  $\epsilon > 0$  there exist  $N_\epsilon$  such that for all  $p \in \mathcal{P}$  we have  $p(X_p \geq N_\epsilon) \leq \epsilon$ . The following theorems were shown in [30].

THEOREM 12 ([30, Theorem 1]). *The collection  $(\mathcal{P}, \ell^{\text{IP}})$  is e.a.s.-predictable iff there are countably many tight collections  $\{\mathcal{P}_i, i \geq 1\}$ . such that*

$$\mathcal{P} = \bigcup_{i=1}^{\infty} \mathcal{P}_i.$$

THEOREM 13 ([30, Corollary 3]). *The collection  $(\mathcal{P}, \ell^{\text{IP}})$  is e.a.s.-learnable iff*

$$\mathcal{P} = \bigcup_{i=1}^{\infty} \mathcal{P}_i,$$

where each  $\mathcal{P}_i$  is tight and relatively open (under total variation distance) in  $\mathcal{P}$ .

We now show that the above theorems can be naturally derived from our general characterizations in the previous sections. The crucial part of the proof is the following technical lemma.

LEMMA 6. *If a collection  $(\mathcal{P}, \ell^{\text{IP}})$  is  $\eta$ -predictable with  $\eta < \frac{1}{4}$ , then there exists nesting  $\{\mathcal{P}_i, i \geq 1\}$  of  $\mathcal{P}$  such that each  $\mathcal{P}_i$  is tight and relatively open in  $\mathcal{P}$  under total variation distance.*

PROOF. We show that if collection  $(\mathcal{P}, \ell^{\text{IP}})$  is  $\eta$ -predictable with  $\eta < 1/4$ , then for any  $p \in \mathcal{P}$  there is an neighborhood  $\mathcal{N}_p$  of  $p$  in  $\mathcal{P}$  such that  $\mathcal{N}_p$  is tight. Let  $m$  be the samples size that achieves  $\eta$ -predictability of  $\mathcal{P}$  with predictor  $\Phi$ . Suppose for some  $p \in \mathcal{P}$ , any neighborhood of  $p$  is not tight. Let  $N = \min\{n : p^m(\max\{X_1^m\} > n) \leq 1 - 4\eta\}$ . Let  $\epsilon \leq \frac{\eta}{m}$ , we have by Lemma 10 that  $\sup_{q \in \mathcal{B}(p, \epsilon)} q^m(\max\{X_1^m\} \geq N) \leq 1 - 3\eta$ . Since  $\mathcal{B}(p, \epsilon)$  is not tight, there exist  $\delta > 0$ , so that for any  $M \in \mathbb{N}$  there exist  $q \in \mathcal{B}(p, \epsilon)$  such that  $q(X \geq M) \geq \delta$ . Let

$$a_1 = \Phi(\underbrace{N, \dots, N}_m),$$

and  $a_n = \max_{X_1^{n+m} \in [a_{n-1}]^{m+n}} \Phi(X_1^{n+m})$ . Let  $M_\delta$  be a number such that  $(1 - \delta)^{M_\delta} < \eta$ , and  $q_\delta \in \mathcal{B}(p, \epsilon)$  with  $q_\delta(X \geq a_{M_\delta}) \geq \delta$ . Now, by construction, with probability at least  $2\eta$ ,  $\Phi$  makes error at step  $m + M_\delta$  when  $q_\delta$  is in force. Contradicting to the  $\eta$ -predictability of  $\mathcal{P}$ .

Since the topology of distributions over  $\mathbb{N}$  induced by total variation distance is separable, any open covering of  $\mathcal{P}$  has a countable subcover. The lemma follows.  $\square$

We now ready to prove Theorem 12 and Theorem 13.

**PROOF OF THEOREM 12.** We first show that if a collection  $\mathcal{P}_i$  is tight then  $(\mathcal{P}_i, \ell^{\text{IP}})$  is  $\eta$ -predictable for all  $\eta > 0$ . To see this, by definition of tightness, for any  $n \geq 1$  there exist a number  $N_n$  such that for all  $p \in \mathcal{P}$  we have  $p(X_p \geq N_n) \leq \frac{\eta}{2^n}$ . The prediction rule goes as follows, at step  $n$  we predict  $N_n$ . The  $\eta$ -predictability of  $(\mathcal{P}_i, \ell^{\text{IP}})$  follows by a union bound. The sufficiency now follows by Theorem 1.

To prove the necessity, by Theorem 1, we know that if  $(\mathcal{P}, \ell^{\text{IP}})$  is *e.a.s.*-predictable then there exist nesting  $\{\mathcal{P}_i, i \geq 1\}$  of  $\mathcal{P}$  such that  $(\mathcal{P}_i, \ell^{\text{IP}})$  is  $\eta$ -predictable for some  $\eta < \frac{1}{4}$ . By Lemma 6, each of  $\mathcal{P}_i$  is a countable union of tight classes, the theorem follows.  $\square$

**PROOF OF THEOREM 13.** By Theorem 7, we know that a class  $\mathcal{P}_i$  is identifiable in  $\mathcal{P}$  if and only if  $\mathcal{P}_i$  is relatively open in  $\mathcal{P}$  under total variation distance. The theorem follows by Theorem 8, Lemma 6 and the fact that  $(\mathcal{P}_i, \ell^{\text{IP}})$  is  $\eta$ -predictable for all  $\eta > 0$  if  $\mathcal{P}$  is tight.  $\square$

**7.3. Online Learning.** Let  $\mathcal{H}$  be a set of binary measurable functions from  $\mathbb{R}^d \rightarrow \{0, 1\}$ ,  $\mu$  is an arbitrary distributions on  $\mathbb{R}^d$ . We consider the following prediction game with two parties, Nature and the learner, both know  $\mathcal{H}$  and  $\mu$ . At the beginning Nature chooses some  $h \in \mathcal{H}$ . At times step  $n$ , Nature independently samples  $X_n \sim \mu$  and provides it to the learner. The learner outputs a guess  $Y_n$  of  $h(X_n)$  based on his previous observations  $(X_1, h(X_1)), \dots, (X_{n-1}, h(X_{n-1}))$  and the new sample  $X_n$ . Nature reveals  $h(X_n)$  after the guess  $Y_n$  has been made by the learner. The learner incurs a loss at step  $n$  if  $Y_n \neq h(X_n)$ .

To put the setup into our general framework, for any  $\mu$  and  $h \in \mathcal{H}$ , one may define the underlying random process to be  $Z_1, Z_2, \dots$  such that  $Z_i = (X_{i+1}, h(X_i))$  (with  $Z_1 = (0, h(X_1))$  by convention) where  $X_1, X_2, \dots$  are *iid* samples of  $\mu$ . The loss is defined as

$$\ell^{\text{OL}}(h, Z_1^n, \Phi(Z_1^{n-1})) = 1\{\Phi(Z_1^{n-1}) \neq h(X_n)\}.$$

For any function classes  $\mathcal{H}$ , we will denote  $\mathcal{H}$  to be the corresponding random processes as well when the underlying distribution  $\mu$  is fixed.

The following result was shown in [31]. See Appendix C.3 for a proof with the application of Theorem 1.

**THEOREM 14.** *Let  $\mu$  be an arbitrary distribution over  $\mathbb{R}^d$ , and  $\mathcal{H}$  be a class of measurable functions from  $\mathbb{R}^d \rightarrow \{0, 1\}$ . Then  $(\mathcal{H}, \ell^{\text{OL}})$  is *e.a.s.*-predictable if and only if  $\mathcal{H}$  is effectively countable w.r.t.  $\mu$ , meaning that there exist a countable class  $\mathcal{H}'$  such that for any  $h \in \mathcal{H}$  there exist  $h' \in \mathcal{H}'$  we have*

$$\Pr_{X \sim \mu}[h(X) \neq h'(X)] = 0.$$

**7.4. Testing matrix properties.** We now apply our framework of testing functional properties to matrix properties. To begin with, we first prove the following theorem, which is a direct corollary of Theorem 20(1).

**THEOREM 15.** *Let  $\mathcal{H}$  the set of all iid processes with marginal distributions over  $[0, 1]^d$  for some  $d \geq 1$ . For all  $A \subset [0, 1]^d$ , we define loss  $\ell_A(p, X_1^n, Y_n) = 1\{1\{\mathbb{E}_{X \sim p}[X] \in A\} \neq Y_n\}$ , where the prediction  $Y_n$  tries to decide whether  $\mathbb{E}_{X \sim p}[X] \in A$  or not. We have  $(\mathcal{H}, \ell_A)$  is *e.a.s.*-predictable if  $A$  is closed in  $[0, 1]^d$ .*

PROOF. This follows directly from Theorem 20(1), since  $A, \bar{A}$  are  $F_\sigma$ -separable if  $A$  is closed, where  $\bar{A}$  is the complement of  $A$  in  $[0, 1]^d$ .  $\square$

For any function  $f : [0, 1]^d \rightarrow \{0, 1\}$ , we will be able to identify a set  $A_f = \{\mathbf{x} \in [0, 1]^d : f(\mathbf{x}) = 1\}$ . Let  $A_1, \dots, A_n \subset [0, 1]^d$  be finitely many sets such that  $(\mathcal{H}, \ell_{A_i})$  is *e.a.s.*-predictable for all  $1 \leq i \leq n$ . Let  $g : [0, 1]^d \rightarrow \{0, 1\}$  be an arbitrary function, denote  $f(\mathbf{x}) = g(1_{A_1}(\mathbf{x}), \dots, 1_{A_n}(\mathbf{x}))$ . It is easy to show that  $(\mathcal{H}, \ell_{A_f})$  is also *e.a.s.*-predictable.

We now consider the following problem setup. Let  $\mathbf{X}$  be a  $d \times d$  random matrix such that each entry  $\mathbf{X}(i, j)$  is a Bernoulli random variable. We denote  $\mathbb{E}[\mathbf{X}]$  to be a (deterministic) matrix that takes expectation entry-wise on  $\mathbf{X}$ . Let  $\mathbf{X}_1, \mathbf{X}_2, \dots$  be *iid* realization of  $\mathbf{X}$ , which are binary matrices. We will try to identify the properties of  $\mathbb{E}[\mathbf{X}]$  by observing the samples  $\mathbf{X}_1, \mathbf{X}_2, \dots$ . Clearly, we can associate properties of  $\mathbb{E}[\mathbf{X}]$  with subsets of  $[0, 1]^{d \times d}$ . We will denote  $\mathcal{H}$  to be the class of all *iid* Bernoulli random matrices process. We say a property of  $\mathbb{E}[\mathbf{X}]$  is *e.a.s.*-predictable if  $(\mathcal{H}, \ell_A)$  is *e.a.s.*-predictable where  $A \subset [0, 1]^{d \times d}$  is the subset corresponding to the property.

THEOREM 16. *The singularity of  $\mathbb{E}[\mathbf{X}]$  is e.a.s.-predictable.*

PROOF. Note that the determinant is a continuous function w.r.t. the entries of the matrix. Thus subsets in  $[0, 1]^{d \times d}$  corresponding to the determinant being 0 are closed. The theorem follows by Theorem 15.  $\square$

THEOREM 17. *To determine if  $\mathbb{E}[\mathbf{X}]$  has rank  $k$  is e.a.s.-predictable for all  $k$ .*

PROOF. Note that to check whether a matrix has rank  $k$ , one only needs to check the maximum non-singular square submatrix is of dimension  $k$ . Thus the property can be expressed as a function of finite singularity tests. By Theorem 16 we know that the property is still *e.a.s.*-predictable.  $\square$

The above theorem is surprising because all of  $\mathbf{X}_1, \dots, \mathbf{X}_n$  will have rank  $\geq n - 2 \log n$  with high probability even when the entries are Bernoulli( $\frac{1}{2}$ ) (i.e.  $\mathbb{E}[X]$  has rank 1). To see this, for each  $\mathbf{X}_i$ , note that the probability of being full rank for the first  $n - 2 \log n$  rows is

$$\prod_{i=2 \log n}^n (1 - 2^{-i}) \geq 1 - \sum_{i=\log 2n}^n 2^{-i} = 1 - O\left(\frac{1}{n^2}\right),$$

and the probability that none of  $\mathbf{X}_i$ ,  $1 \leq i \leq n$  have rank  $\leq n - 2 \log n$  is  $\geq 1 - O(1/n)$ . Yet, Theorem 17 shows that it is still possible to infer the rank of  $\mathbb{E}[\mathbf{X}]$  from the realizations  $\mathbf{X}_1, \mathbf{X}_2, \dots$

To unravel the proof in Theorem 17, we also give a more algorithmic way of estimating the rank that is correct eventually almost surely. Let  $n$  be the sample size we observed, and  $\hat{p}_{i,j}$  be the empirical mean of the  $(i, j)$ th entry. The estimation of rank is given by the following optimization problem.

$$\begin{aligned} & \min_X \text{rank}(X) \\ & \text{s.t. } \forall i, j, |X_{i,j} - \hat{p}_{i,j}| \leq \frac{\log n}{\sqrt{n}}. \end{aligned}$$

We show that the estimation given by the rule above converges to  $\text{rank}(\mathbb{E}[\mathbf{X}])$  w.p. 1. Note that, by the selection of the threshold of the constraints at  $\log n / \sqrt{n}$  and using the Chernoff

bound, the probability that the matrix  $\mathbb{E}[\mathbf{X}]$  is not in the feasible set of the optimization above is at most  $O\left(\frac{1}{n^2}\right)$ . By the Borel-Cantelli lemma,  $\mathbb{E}[\mathbf{X}]$  will be in the feasible set eventually almost surely. Therefore, the rule will never over estimate the rank in the limit.

Denote  $k = \text{rank}(\mathbb{E}[\mathbf{X}])$ . There exists a full rank  $k \times k$  submatrix  $\mathbf{Y}$  of  $\mathbb{E}[\mathbf{X}]$ . By the continuity of the determinant and since  $\det(\mathbf{Y}) \neq 0$ , there is a neighborhood  $\mathcal{N}$  of  $\mathbf{Y}$  such that any  $\mathbf{T} \in \mathcal{N}$  has rank  $k$ . Since  $\log n / \sqrt{n} \rightarrow 0$ , there is some constant  $N$  depending on the neighborhood  $\mathcal{N}$  such that for all  $n > N$ , the feasible set of the optimization restricted to coordinates of  $\mathbf{Y}$  will be a subset of  $\mathcal{N}$ . Combining this with the observation that  $\mathbb{E}[\mathbf{X}]$  can be out of the feasible set only finitely many times, we conclude that the rule will eventually output  $k$  almost surely.

**THEOREM 18.** *To determine if  $\mathbb{E}[\mathbf{X}]$  has eigenvalues of multiplicity more than 1 is e.a.s.-predictable.*

**PROOF.** For any matrix  $A$ , consider the characteristic polynomial  $p_A$  of  $A$ . We know that the coefficients of  $p_A$  are polynomials of the entries of  $A$ . We now only need to check if  $\text{GCD}(p_A, p'_A) = 1$ , where  $p'_A$  is the derivative of  $p_A$ . Note that this can be done by checking the resultant of  $p_A, p'_A$  is zero. Since resultant is continuous functions of the coefficients, the theorem follows using Theorem 15.  $\square$

While one should expect most properties of matrices to be e.a.s.-predictable, we have the following open problem.

**PROBLEM 2 (Open Problem).** *Is determining whether a matrix is diagonalizable e.a.s.-predictable?*

**8. Conclusion.** In this paper, we introduced a general framework for the infinite horizon prediction problems that require to be correct eventually almost surely. We provide general characterization for the existence of the e.a.s.-prediction rules. Our characterization has a natural interpretation using regularization that relates e.a.s.-predictability with the decomposition of a class into uniformly predictable subclass. While our characterization does not completely capture the e.a.s.-predictability in the most general case, it allows us to recover prior results in the literature with simple and elementary proofs, and resolves conjectures posed therein. We also provide tight characterisations for several special cases. There are also several open problems that we have mentioned through out this paper which concerns the exact characterization for certain cases.

## APPENDIX A: MORE ON HYPOTHESIS TESTING

In this section, we present more results on the almost surely hypothesis testing framework (see Section 7.1) using our general machinery.

**A.1. Hypothesis testing with discrete domain.** As we have demonstrated in Section 7.1, the characterization of hypothesis testing for continuous distributions is involved, e.g., we are unaware if  $F_\sigma$ -separability under  $KS$ -distance is necessarily to achieve e.a.s.-predictability even assumed absolute continuity. However, we have the following full characterization for discrete distributions.

**THEOREM 19.** *Let  $\mathcal{H}_1, \mathcal{H}_2$  be collections of distributions over  $\mathbb{N}$ , then  $(\mathcal{H}_1 \cup \mathcal{H}_2, \ell^{\text{HT}})$  is e.a.s.-predictable if and only if  $\mathcal{H}_1, \mathcal{H}_2$  are  $F_\sigma$ -separable under total variation distance.*



We first prove a general lemma.

LEMMA 7. *Let  $\mathcal{H}_1, \mathcal{H}_2$  be classes of distributions over  $\mathbb{R}^d$ . If  $(\mathcal{H}_1 \cup \mathcal{H}_2, \ell^{\text{HT}})$  is  $\eta$ -predictable for some  $\eta \leq \frac{1}{4}$ , then*

$$\inf\{\|p_1 - p_2\|_{TV} : p_1 \in \mathcal{H}_1, p_2 \in \mathcal{H}_2\} > 0.$$

PROOF. Let  $i$  be the sample complexity for  $(\mathcal{H}_1 \cup \mathcal{H}_2, \ell^{\text{HT}})$  to achieve  $\eta$ -predictability, i.e., there exists a predictor such that the probability of making errors after step  $i$  is upper bounded by  $\eta$  for all  $p \in \mathcal{H}_1 \cup \mathcal{H}_2$ . Suppose the lemma does not hold. We can find  $p_1 \in \mathcal{H}_1$  and  $p_2 \in \mathcal{H}_2$ , such that

$$\|p_1 - p_2\|_{TV} \leq \frac{1}{4i}.$$

We have, by Lemma 10,  $\|p_n^i - p^i\|_{TV} \leq \frac{1}{4}$ . However, this implies that any predictor will make an error at step  $i + 1$  w.p.  $\geq \frac{3}{8} > \frac{1}{4}$  by Lemma 9. This will violate the  $\eta$ -predictability of  $(\mathcal{H}_1 \cup \mathcal{H}_2, \ell^{\text{HT}})$ .  $\square$

PROOF OF THEOREM 19. We first prove the necessity. By Theorem 1, there exists a nesting  $\{\mathcal{G}_i, i \geq 1\}$  for  $\mathcal{H}_1 \cup \mathcal{H}_2$  such that  $(\mathcal{G}_i, \ell^{\text{HT}})$  is  $\frac{1}{4}$ -predictable for all  $i \geq 1$ . The necessity is followed now by Lemma 7.

To prove the sufficiency, by  $F_\sigma$ -separability, there exist nestings  $\{A_i, i \geq 1\}$  and  $\{B_i, i \geq 1\}$  for  $\mathcal{H}_1$  and  $\mathcal{H}_2$  respectively, such that

$$\forall i \geq 1, \inf\{\|p_1 - p_2\|_{TV} : p_1 \in A_i, p_2 \in B_i\} \geq \frac{1}{i}.$$

We now show that  $(A_i \cup B_i, \ell^{\text{HT}})$  is *e.a.s.*-learnable (see Definition 6). The theorem will now follow by Theorem 6.

For any  $\eta > 0$ , the stopping rule  $\tau_\eta$  goes as follows. We partition  $\tau_\eta$  into two phases, at phase one,  $\tau_\eta$  tries to determine if the missing probability mass of the underlying distribution is less than  $\frac{1}{4i}$  with confidence at least  $1 - \frac{\eta}{2}$ . This can be done by requesting additional samples to test if new symbols appear in the new coming samples. At phase two, we will estimate the probability mass only on the symbols we have seen so far at the end of phase one, so that the  $L_1$  error of the estimation is upper bounded by  $\frac{1}{4i}$  with confidence  $\geq 1 - \frac{\eta}{2}$ . This can be done since the symbols we have seen are finite. Now, the prediction at the end of phase two is as follows, if the total variation distance of the estimated distribution is closer to  $A_i$  we predict  $p \in \mathcal{H}_1$ , otherwise, we predict  $p \in \mathcal{H}_2$ . We will keep the prediction for all steps after phase two of  $\tau_\eta$ . It is easy to verify that such a rule satisfies the requirements of *e.a.s.*-learnability and  $\tau_\eta$  stops finitely almost surely on all distributions over  $\mathbb{N}$ .  $\square$

REMARK 10. *Note that, the proof of Theorem 19 above does not work for distributions over  $\mathbb{R}$ . Since for distribution classes  $A, B$  over  $\mathbb{R}$ , if*

$$\inf\{\|p_1 - p_2\|_{TV} : p_1 \in A, p_2 \in B\} > 0,$$

*we cannot conclude that  $(A \cup B, \ell^{\text{HT}})$  is *e.a.s.*-learnable. See construction in Remark 4.*

*Note that, this does not imply that Theorem 19 does not hold for distributions over  $\mathbb{R}$ . We leave it as an open problem to determine whether Theorem 19 holds for arbitrary distributions over  $\mathbb{R}$ .*

**A.2. Testing properties of first moment.** For any two sets  $A, B \subset \mathbb{R}^d$ , we would like to determine whether the underlying distribution has first moment in  $A$  or  $B$  with the premise that the first moment is in  $A \cup B$ . The following theorem is due to Dembo and Peres [9].

**THEOREM 20** ([9, Theorem 1]). *Let  $A, B \subset \mathbb{R}^d$  be disjoint sets. Denote  $\mathcal{H}_1^r$  and  $\mathcal{H}_2^r$  to be classes of all distributions over  $\mathbb{R}^d$  with first moments in  $A$  and  $B$  respectively that have finite  $r$ th moment. We have:*

1. *If  $r > 1$ , then  $(\mathcal{H}_1^r \cup \mathcal{H}_2^r, \ell^{\text{HT}})$  is e.a.s.-predictable iff  $A$  and  $B$  are  $F_\sigma$ -separable under  $L_2$  distance.*
2. *If  $r = 1$ , then  $(\mathcal{H}_1^r \cup \mathcal{H}_2^r, \ell^{\text{HT}})$  is e.a.s.-predictable iff  $A$  and  $B$  are contained in disjoint open sets of  $\mathbb{R}^d$ .*

We now provide an alternative proof of the theorem above by applying our general characterizations. We need the following lemma.

**LEMMA 8.** *Let  $X_1, \dots, X_n$  be iid random variables with  $\mathbb{E}[|X_1|^r] = M < \infty$  where  $r > 1$  and  $\mathbb{E}[X_1] = u$ ,  $\bar{X} = \frac{X_1 + \dots + X_n}{n}$ , then for any  $\epsilon > 0$  we have*

$$\Pr[|\bar{X} - u| \geq \epsilon] \leq \frac{C_{\epsilon, M}}{n^{r-1}},$$

where  $C_{\epsilon, M}$  is a constant that depends only on  $\epsilon$  and  $M$ .

**PROOF.** W.l.o.g., we assume  $u = 0$ . By Bahr-Essen inequality [29], we have

$$\mathbb{E}[|\bar{X}|^r] \leq 2 \frac{\sum_{i=1}^n \mathbb{E}[|X_i|^r]}{n^r} = \frac{2Mn}{n^r} = \frac{2M}{n^{r-1}}.$$

The lemma follows by a simple application of Markov inequality.  $\square$

**PROOF OF THEOREM 20(1).** We first prove the sufficiency. Since  $A, B$  are  $F_\sigma$ -separable, we have by Lemma 3 that there exist nestings  $\{A_i, i \geq 1\}$ ,  $\{B_i, i \geq 1\}$  of  $A, B$  respectively, such that

$$\forall i \geq 1, \inf\{\|\mathbf{x} - \mathbf{y}\|_2 : \mathbf{x} \in A_i, \mathbf{y} \in B_i\} \geq \frac{1}{i}.$$

We now define  $\mathcal{G}_i = \{p \in \mathcal{H}_1^r \cup \mathcal{H}_2^r : \mathbb{E}[|X_p|^r] \leq i \text{ and } \mathbb{E}[X_p] \in A_i \cup B_i\}$ , where  $X_p$  is the random variable governed by distribution  $p$ . Clearly, we have  $\mathcal{G}_i \subset \mathcal{G}_{i+1}$  for all  $i \geq 1$ . We now show that  $\mathcal{H}_1^r \cup \mathcal{H}_2^r = \bigcup_{i \geq 1} \mathcal{G}_i$ , thus proving that  $\{\mathcal{G}_i, i \geq 1\}$  forms a nesting of  $\mathcal{H}_1^r \cup \mathcal{H}_2^r$ . To see this, let  $p \in \mathcal{H}_1^r \cup \mathcal{H}_2^r$  be an arbitrary distribution with  $\mathbb{E}[X_p] \in A$  (respectively  $\in B$ ), we have  $\mathbb{E}[X_p] \in A_i$  (respectively  $\in B_i$ ) for some  $i$ . Since  $X_p$  also has finite  $r$ th moment, we have  $\mathbb{E}[|X_p|^r] \leq j$  for some  $j$ . Therefore, we have  $p \in \mathcal{G}_{\max\{i, j\}}$  by nesting of  $A_i$  (respectively  $B_i$ ).

By Theorem 1, we only need to show that  $(\mathcal{G}_i, \ell^{\text{HT}})$  is  $\eta$ -predictable for all  $i \geq 1$  and  $\eta > 0$ . By Lemma 8, we know that by letting sample size  $n$  large enough one can make

$$p(|\bar{X}_p - \mathbb{E}[X_p]| \geq \frac{1}{2i}) \leq \eta$$

for all distributions  $p \in \mathcal{G}_i$ . To achieve the  $\eta$ -predictability, we will predict at step  $n$  that the first moment is in  $A_i$  (respectively  $B_i$ ) if  $\bar{X}_p$  is closer to  $A_i$  (respectively  $B_i$ ), and *retain* the same prediction for all time step  $\geq n$ .

We now prove the necessity. Let  $\mathcal{G}$  be the class of all Gaussian distributions with mean vectors in  $A \cup B$  and covariance matrix  $I$  (i.e., the identity matrix). We have  $(\mathcal{G}, \ell^{\text{HT}})$  is *e.a.s.*-predictable, since  $\mathcal{G}$  is a subset of  $\mathcal{H}_1^r \cup \mathcal{H}_2^r$ . Now, by Theorem 1, there exists a nesting  $\{\mathcal{G}_i, i \geq 1\}$  of  $\mathcal{G}$  such that for all  $i \geq 1$ ,  $(\mathcal{G}_i, \ell^{\text{HT}})$  is  $\frac{1}{4}$ -predictable. Define  $A_i = \{\mathbf{x} \in A : \exists p \in \mathcal{G}_i \text{ s.t. } \mathbb{E}[X_p] = \mathbf{x}\}$  (and  $B_i$  similarly). Clearly,  $\{A_i, i \geq 1\}$  forms a nesting of  $A$ , and  $\{B_i, i \geq 1\}$  forms a nesting of  $B$ . We show that there is no limit point of  $A_i$  (respectively  $B_i$ ) in  $B$  (respectively  $A$ ). This will complete the proof by Lemma 3.

Suppose we have  $\mathbf{x}_1, \mathbf{x}_2, \dots \in A_i$  and  $\mathbf{x} \in B_i$  such that  $\|\mathbf{x}_n - \mathbf{x}\|_2 \rightarrow 0$  as  $n \rightarrow \infty$ . Denote  $p_n$  to be the Gaussian distribution with mean  $\mathbf{x}_n$  and  $p$  be the Gaussian distribution with mean  $\mathbf{x}$ . We have  $\|p_n - p\|_{TV} \rightarrow 0$  as  $n \rightarrow \infty$ . By Lemma 7, this will violate the  $\frac{1}{4}$ -predictability of  $\mathcal{G}_i$ .  $\square$

## APPENDIX B: TECHNICAL LEMMAS

In this section we collect some technical lemmas that are used in this paper. The following lemma is due to Le Cam [14].

**LEMMA 9 (Le Cam's Two Point Method).** *Let  $p_1, p_2$  be two distributions over the same probability space  $\mathcal{X}$ . Then for any estimator  $\Phi : \mathcal{X} \rightarrow \{1, 2\}$  we have*

$$\max_{i \in \{1, 2\}} \Pr_{X \sim p_i}(\Phi(X) \neq i) \geq \frac{1 - \|p_1 - p_2\|_{TV}}{2}.$$

The following lemma bounds the total variation of  $n$ -fold *iid* distributions from the marginal distributions, which is "folklore" in the literature.

**LEMMA 10.** *Let  $p_1, p_2$  be distributions over the same Polish space with Borel  $\sigma$ -algebra, and  $p_1^n, p_2^n$  be the  $n$ -fold *iid* distributions of  $p_1, p_2$  respectively. Then*

$$\|p_1^n - p_2^n\|_{TV} \leq 1 - (1 - \|p_1 - p_2\|_{TV})^n.$$

*In particular, we have  $\|p_1^n - p_2^n\|_{TV} \leq n\|p_1 - p_2\|_{TV}$ .*

**PROOF.** Consider a perfect coupling  $(X_1, X_2)$  of  $(p_1, p_2)$ . We have

$$\Pr[X_1 \neq X_2] = \|p_1 - p_2\|_{TV}.$$

Let  $(X_1^n, X_2^n)$  be the  $n$ -fold *iid* copy of  $(X_1, X_2)$ . We have

$$\begin{aligned} \Pr[X_1^n = X_2^n] &= \Pr[X_1 = X_2]^n, \text{ by independence} \\ &= (1 - \|p_1 - p_2\|_{TV})^n. \end{aligned}$$

Now, we have

$$\begin{aligned} \|p_1^n - p_2^n\|_{TV} &\leq \Pr[X_1^n \neq X_2^n], \text{ by coupling lemma} \\ &= 1 - \Pr[X_1^n = X_2^n] \\ &= 1 - (1 - \|p_1 - p_2\|_{TV})^n. \end{aligned}$$

The last part follows since  $(1 - x)^n \geq 1 - nx$  for all  $x \in [0, 1]$  and  $n \geq 1$ .  $\square$

The following example demonstrates that *e.a.s.*-predictability does not imply the existence of universal nesting.

EXAMPLE 11. Let  $\mathcal{P}$  be the class of all random processes over  $\{0, 1\}^\infty$  such that for any  $p \in \mathcal{P}$  there exists a parameter  $b \in \{0, 1\}$  and strictly monotonically increasing integer sequence  $M_1, M_2, \dots$  with  $M_1 = 1$  satisfying for all  $j \geq 1$ ,  $X_{M_j} = X_{M_{j+2}} = \dots = X_{M_{j+1}-1}$ , and  $X_{M_j}^{M_{j+1}-1}$  is independent of all other random variables in the process, and

$$p(X_{M_j} = b) = 1 - \frac{1}{(j+1)^2}.$$

Note that a process in  $\mathcal{P}$  is purely determined by the parameter  $b$  and sequence  $\{M_j\}_{j \geq 1}$ . We will denote a process to be  $p_b$  if the parameter is  $b$ . Clearly, by Borel-Cantelli lemma  $\mathcal{P}$  is e.a.s.-predictable under the loss  $\ell(p_b, X_1^n, Y_n) = 1\{Y_n \neq b\}$  by predicting  $X_{n-1}$  at each step  $n$ . We now show that there is no nesting  $\{\mathcal{P}_i, i \geq 1\}$  of  $\mathcal{P}$  such that  $(\mathcal{P}_i, \ell)$  is  $\eta$ -predictable for all  $\eta > 0$  and  $i \geq 1$ . Our approach is a proof by contradiction. Let  $\{\mathcal{P}_i, i \geq 1\}$  be a universal nesting of  $\mathcal{P}$  w.r.t. loss  $\ell$ , i.e.,  $(\mathcal{P}_i, \ell)$  is  $\eta$ -predictable for all  $\eta > 0$  and  $i \geq 1$ . We construct a distribution in  $\mathcal{P}$  that is not in  $\bigcup_{i \geq 1} \mathcal{P}_i$ , a contradiction on our supposition that  $\mathcal{P} = \bigcup_{i \geq 1} \mathcal{P}_i$ .

For  $j \geq 1$ , let  $R_j$  be a number such that the class  $\mathcal{P}_j$  is  $\eta_j$ -predictable with a sample of size  $R_j$ , where  $\eta_j$  will be specified later. W.l.o.g., we may assume  $R_j$  is strictly increasing on  $j$ . Let  $p_0, p_1$  be two distributions in  $\mathcal{P}$  that are associated with the sequence  $M_1 = 1$  and  $(M_j = R_{j-1} + 1)_{j \geq 2}$  with parameter  $b = 0$  and  $b = 1$  respectively, i.e.  $p_0, p_1$  share the same partition of independent blocks but with different parameter for  $b$ .

Let  $\|p_0^{R_j} - p_1^{R_j}\|_{TV} = 1 - \epsilon_j$ , where  $p_b^{R_j}$  with  $b \in \{0, 1\}$  is the distribution of  $p_b$  on the initial length- $R_j$  binary strings. Observe that the probability of any length- $R_j$  binary string under either  $p_0$  or  $p_1$  is purely a function of the number  $j$  of blocks which are all-0 (or equivalently all-1), and therefore, so does  $\epsilon_j$ , the total variation distance. Specifically,  $\epsilon_j$  depends only on  $j$ , not the sequence  $\{R_j\}$ , or  $\{\eta_j\}$ , and in particular we can choose  $\eta_j < \frac{\epsilon_j}{2}$ .

By Lemma 9, we know that any prediction rule will make an error with probability at least  $\frac{\epsilon_j}{2}$  on either  $p_0$  or  $p_1$  at time step  $R_j + 1$  if they both belong to  $\mathcal{P}_j$ . Since  $\epsilon_j/2 > \eta_j$ , we conclude that for all  $j \geq 1$  at least one of  $p_0$  or  $p_1$  cannot be in  $\mathcal{P}_j$ .

But  $\mathcal{P}_j \subset \mathcal{P}_k$  for all  $k \geq j$ . Let  $t$  be the smallest number that contains one of  $p_0$  or  $p_1$ .  $\mathcal{P}_t$  cannot contain both, per the argument above, it follows that  $\mathcal{P}_k$  contains that distribution for all  $k \geq t$ . However, the argument above implies that the distribution missing from  $\mathcal{P}_t$  cannot be in  $\bigcup_{k \geq 1} \mathcal{P}_k$ , contradicting the assumption that  $\{\mathcal{P}_i, i \geq 1\}$  is a nesting of  $\mathcal{P}$ .

The following lemmas are used in Example 3. Let  $r_1, r_2, \dots$  be an arbitrary enumeration of rational numbers in  $[0, 1]$ .

LEMMA 11. Let  $\mathcal{B}_k$  be the set of all Bernoulli processes with parameters in

$$\mathcal{S}_k = \{r_1, \dots, r_k\} \cup \left( [0, 1] \setminus \bigcup_{i=1}^{\infty} B(r_i, \frac{1}{k2^i}) \right)$$

where  $B(r_i, \frac{1}{k2^i})$  is the open balls centered at  $r_i$  with radius  $\frac{1}{k2^i}$ ,  $r_1, r_2, \dots$  is an arbitrary enumeration of rational numbers in  $[0, 1]$ . Then  $\mathcal{B}_k$  is  $\eta$ -predictable with irrationality loss for any  $k \geq 1$  and  $\eta > 0$ .

PROOF. We show that for any  $k \in \mathbb{N}$  and  $\eta > 0$ , there exists  $b_\eta$  such that  $\mathcal{B}_k$  is  $\eta$ -predictable with sample size  $b_\eta$ . Let  $X_1, \dots, X_n$  be an iid sample from some  $p \in \mathcal{B}_k$  with  $\mathbb{E}[X_i] = \mu$  and  $\bar{X} = \frac{X_1 + \dots + X_n}{n}$ . We have  $\text{Var}[X_i] \leq 1$ . Chebyshev's inequality then shows that

$$p(|\bar{X} - \mu| \geq \epsilon) \leq \frac{1}{n\epsilon^2}.$$

Fix  $\epsilon = \frac{1}{k2^{k+1}}$ . Let  $b_\eta$  be a number large enough so that  $\frac{1}{b_\eta \epsilon^2} < \eta$ . Therefore for  $n > b_\eta$ ,  $p(|\bar{X} - \mu| \geq \epsilon)$  is less than  $\eta$ . Thus, we can conclude that  $\mathcal{B}_k$  is  $\eta$ -predictable by simply predicting the irrationality of element in  $\mathcal{S}_k$  that is closest to  $\bar{X}$  at step  $b_\eta$ , retaining the prediction perpetually thereafter.  $\square$

LEMMA 12. *Let  $\mathcal{B}$  be a class of Bernoulli processes with parameters in  $\mathcal{S}$ , if  $\mathcal{B}$  is  $\eta$ -predictable w.r.t. the irrationality loss for some  $0 < \eta < \frac{1}{2}$ , then*

$$(B.1) \quad \inf\{|x - r| : x, r \in \mathcal{S} \text{ and } r \in \mathbb{Q}, x \in [0, 1] \setminus \mathbb{Q}\} > 0.$$

PROOF. By definition of  $\eta$ -predictability, there exists a number  $N_\eta$  and prediction rule  $\Phi_\eta$  such that  $\Phi_\eta$  makes no errors after step  $N_\eta$  with probability at least  $1 - \eta$  for all  $p \in \mathcal{B}$ . Suppose, otherwise, that the infimum in equation (B.1) is 0. We now select two sources  $p_0, p_1$  from  $\mathcal{B}$  with parameters  $b_0, b_1$  respectively, such that  $b_0$  is rational and  $b_1$  is irrational and  $|b_0 - b_1| < \frac{1-2\eta}{2N_\eta}$ .

We now have  $\|p_0^{N_\eta} - p_1^{N_\eta}\|_{TV} < 1 - 2\eta$ , where  $p_i^{N_\eta}$  is the distribution of  $p_i$  restricted to the first  $N_\eta$  samples—using the fact that  $\|p^N - q^N\|_{TV} \leq N\|p - q\|_{TV}$  for any distributions  $p, q$  with  $N$ -fold *i.i.d.* distributions  $p^N, q^N$ . Now, by Lemma 9, any prediction rule (in particular  $\Phi_\eta$ ) will make an error at time step  $N_\eta + 1$  with probability  $> \eta$  for either  $p_0$  or  $p_1$ . This contradicts  $\eta$  predictability of  $\Phi_\eta$  on  $\mathcal{B}$ .  $\square$

LEMMA 13. *Let  $\mathcal{S}_1 \subset \mathcal{S}_2 \subset \dots \subset \mathcal{S}_k \subset \dots \subset [0, 1]$  be countably many sets, such that*

$$(B.2) \quad \forall k, \inf\{|x - r| : x, r \in \mathcal{S}_k \text{ and } r \in \mathbb{Q}, x \in [0, 1] \setminus \mathbb{Q}\} > 0.$$

*If  $\bigcup_{k \in \mathbb{N}} \mathcal{S}_k$  contains all rational numbers in  $[0, 1]$ , then the irrational numbers in  $\mathcal{S}_k$  are nowhere dense in  $[0, 1]$  for all  $k$ .*

PROOF. Suppose otherwise, the set of irrational numbers  $\mathcal{I}_k$  in  $\mathcal{S}_k$  is not nowhere dense. By definition, there exists an interval  $[a, b] \subset \text{col}(\mathcal{I}_k)$ , where  $\text{col}$  denotes for closure. Since the rational numbers in  $[0, 1]$  are dense, there exists some rational number  $r \in [a, b]$ , and therefore  $r \in \text{col}(\mathcal{I}_k)$ . Since  $r \in \bigcup_{k \in \mathbb{N}^+} \mathcal{S}_k$ , there exist some  $k' \geq k$  such that  $r \in \mathcal{S}_{k'}$ . However, we also have  $\mathcal{S}_k \subset \mathcal{S}_{k'}$ . Which implies that  $r$  is the limit point of irrational numbers in  $\mathcal{S}_{k'}$ , contradicting the assumption (B.2).  $\square$

## APPENDIX C: OMITTED PROOFS

In this section we collection some of the omitted proofs in the main paper.

### C.1. Proofs in Section 5.2.

PROOF OF THEOREM 7. Suppose there is a limit point  $p$  of  $\mathcal{U} \setminus \mathcal{V}$  in  $\mathcal{V}$ , we show that  $\mathcal{V}$  is not identifiable in  $\mathcal{U}$ . Let  $p_1, p_2, \dots, p_n, \dots \in \mathcal{U} \setminus \mathcal{V}$  such that  $\|p_n - p\|_{TV} \rightarrow 0$  as  $n \rightarrow \infty$ , where we have abused the notation  $p, p_n$  to denote the marginal distributions of processes  $p, p_n$  as well. For any  $\eta \leq \frac{1}{4}$ , we denote  $\tau_\eta$  as the stopping rule that identifies  $\mathcal{V}$  in  $\mathcal{U}$ . Let  $A_m$  be the event that  $\tau_\eta$  stops before step  $m$ . We have  $p(A_m) \rightarrow 1$  as  $m \rightarrow \infty$ . Taking  $M$  be a number such that  $p(A_M) \geq \frac{3}{4}$ . Now, since  $\|p_n - p\|_{TV} \rightarrow 0$ , we have some  $p_N$  such that  $\|p_N - p\|_{TV} \leq \frac{1}{4M}$ . Therefore, we have  $|p(A_M) - p_N(A_M)| \leq \frac{1}{4}$ , where we have used the inequality  $\|p^M - p_N^M\|_{TV} \leq M\|p - p_N\|_{TV}$  if  $p^M, p_N^M$  are the  $M$ -fold *i.i.d.* distributions of  $p, p_N$  respectively. However, this would imply that  $\tau_\eta$  stops on  $p_N$  finitely with probability  $\geq \frac{3}{4} - \frac{1}{4} \geq \frac{1}{4}$ , a contradiction.

Suppose now there are no limit points of  $\mathcal{U} \setminus \mathcal{V}$  in  $\mathcal{V}$ . We know that for any point  $p \in \mathcal{V}$ , there exists a open ball  $\mathcal{B}_p$  of  $p$  such that  $t_p \triangleq \inf\{\|p - q\|_1 : p \in \mathcal{B}_p \text{ and } q \in \mathcal{U} \setminus \mathcal{V}\} > 0$ . Since the topological space of distributions over  $\mathbb{N}$  is separable, there exist countably many distributions  $p_1, p_2, \dots \in \mathcal{V}$  such that  $\mathcal{V} = \bigcup_{i \in \mathbb{N}} \mathcal{B}_{p_i}$ . Let  $\mathcal{V}_n = \bigcup_{i=1}^n \mathcal{B}_{p_i}$ , we show that  $\mathcal{V}_n$  can be distinguished with  $\mathcal{U} \setminus \mathcal{V}$  with arbitrary high confidence by observing bounded samples. We denote  $r_i$  to be the radius of ball  $\mathcal{B}_{p_i}$  and  $t_i = t_{p_i}$ . Denote  $n_i = \min\{n : p_i(X \geq n) \leq t_i/4\}$ . Let  $q$  be the underlying distribution, we can estimate the empirical frequency  $\hat{q}$  of  $q$  on  $\{1, 2, \dots, n_i\}$  with error at most  $t_i/4$  and arbitrary high confidence using bounded samples. The key observation is that, for any  $q \in \mathcal{U} \setminus \mathcal{V}$  we have  $\|q - p_i\|_1 \geq r_i + t_i$ . Meaning that we can distinguish distributions from  $\mathcal{B}_{p_i}$  and  $\mathcal{U} \setminus \mathcal{V}$  by checking whether  $\|\hat{q} - p_i\|_1 \leq r_i + t_i/2$  with arbitrary high confidence using bounded samples. A similar argument as in Example 7 completes the proof.  $\square$

**PROOF OF THEOREM 8.** We first show that the stated conditions on a collection  $\mathcal{P}$  and loss  $\ell$  are sufficient to guarantee that  $\mathcal{P}$  is *e.a.s.*-learnable. Namely, for all  $\eta > 0$ , we will find an estimator  $\Phi$  and a stopping rule  $\tau$  such that for all  $p \in \mathcal{P}$ , the probability  $\Phi$  incurs non-zero loss after  $\tau$  stops is  $\leq \eta$ .

Now the conditions stated imply that for all  $\eta > 0$ , there is an identifiable nesting  $\mathcal{P} = \bigcup_{n \in \mathbb{N}} \mathcal{P}_n$  and a sequence of numbers  $\{m_n : n \geq 1\}$  such that each  $\mathcal{P}_n$  is  $\frac{\eta}{2}$ -predictable with sample size  $m_n$ . Since  $\mathcal{P}_n$  is identifiable, there is a stopping rule  $\sigma_n$  that stops after a finite time on  $\mathcal{P} \setminus \mathcal{P}_n$  with probability at most  $\eta/2^{n+1}$  and stops finitely almost surely on  $\mathcal{P}_n$ .

We will assume without loss of generality that  $\sigma_n$  only stops on sequences of length  $\geq m_n$ .

The stopping rule  $\tau$  for  $(\mathcal{P}, \ell)$  stops if for some  $n$ ,  $\sigma_n$  has stopped. Let  $N$  be the index of the stopping rule that stops earliest. The prediction for  $(\mathcal{P}, \ell)$  is now the  $\eta/2$ -predictor for  $\mathcal{P}_N$ , which we call  $\Phi_N$ .

For all  $n$ , define

$$A_n = \{X_1^\infty : \sigma_n \text{ stops on } X_1^n\}.$$

We claim that:

1. The stopping rule  $\tau$  stops with probability 1. This is because  $p \in \mathcal{P}_k$  for some  $k$ , we have  $\sigma_k$  stops with probability 1.
2. The probability that  $\tau$  stops but  $\Phi$  incurs non-zero loss is  $\leq \eta$ . The probability that  $\tau$  stops but  $p \notin \mathcal{P}_N$  is  $\leq \sum_{i=1}^\infty p(A_i) \leq \eta/2$ . Finally, since  $\mathcal{P}_N$  is  $\eta/2$ -predictable with sample size  $m_N$ , the probability that  $\Phi_N$  predicts incorrectly after sample size  $m_N$  (which we are guaranteed is the case since we assumed  $\sigma_N$  only stops on sequences with length  $\geq m_N$ ) is  $\leq \eta/2$ . The claim follows by union bound.

Now, suppose  $(\mathcal{P}, \ell)$  is *e.a.s.*-learnable, and  $\mathcal{P}$  are *iid* measures on  $\mathcal{X}^\infty$ . For any  $\eta > 0$ , consider the stopping rule  $\tau_{\eta/2}$  and the estimate  $\Phi_{\eta/2}$  that *e.a.s.*-learns  $(\mathcal{P}, \ell)$ . Let

$$T_n = \{X_1^\infty : \tau_{\eta/2}(X_1^n) = 1\}$$

be the set of sequences on which  $\tau$  has stopped at or before step  $n$ . Now for all  $n$ , let

$$\mathcal{P}_n = \{p \in \mathcal{P} : p(T_n) > 1 - \eta/2\}.$$

Clearly, we have  $\mathcal{P}_n \subset \mathcal{P}_{n+1}$  and that  $\bigcup_{n \geq 1} \mathcal{P}_n = \mathcal{P}$ . For all  $n$ , and for all  $p \in \mathcal{P}_n$ ,  $\Phi_{\eta/2}$  incurs non-zero loss on samples of length  $n$  with probability  $\leq \eta/2$  by construction, namely,  $\mathcal{P}_n$  is  $\eta$ -predictable by the union bound.

We will now show that  $\mathcal{P}_n$  are identifiable in  $\mathcal{P}$  to wrap up the theorem.

For any assigned confidence  $\delta$ , we construct a stopping rule  $\tau_\delta$  that stops after a finite time with probability 1 when the underlying process is from  $\mathcal{P}_n$  and with probability at most  $\delta$  on processes from  $\mathcal{P} \setminus \mathcal{P}_n$ . To do so, we choose an arbitrary sequence  $\{e_m\}$  with  $e_m \rightarrow 0$  as  $m \rightarrow \infty$ . The stopping rule is partitioned into phases. At phase  $m$  we estimate  $p(T_n)$  with confidence  $\geq 1 - \frac{\delta}{2^m}$  and error  $\leq e_m/4$ , by considering independent sample blocks of length  $n$ . If the estimate is larger than  $1 - \eta/2 + e_m$  we stop, otherwise we continue to phase  $m + 1$ . Now, if we have  $p(T_n) > 1 - \eta/2$ , then there exist some number  $M$  such that  $p(T_n) > 1 - \eta/2 + 2e_m$  for all  $m \geq M$ . Therefore, for all  $m \geq M$ , with probability at most  $\delta/2^m$  we will not stop at phase  $m$ . By Borel-Cantelli lemma, with probability 1 we will stop in a finite time. A similar argument yields that if  $p(T_n) \leq 1 - \eta/2$ , then we stop with probability at most  $\delta$ .  $\square$

### C.2. Proofs in Section 7.1.

PROOF OF LEMMA 3. By definition of  $F_\sigma$ -separability, we have collections of *closed* sets  $\{A'_i\}$  and  $\{B'_i\}$  such that  $A \subset \bigcup_i A'_i$  and  $B \subset \bigcup_i B'_i$ . Since finite unions of closed sets are closed, we may assume  $\{A'_i\}, \{B'_i\}$  to be nested. Define

$$A_i = \{x : x \in A'_k \text{ and } d(x, B'_k) \geq 1/i, k \in \mathbb{N}\}$$

and

$$B_i = \{y : y \in B'_k \text{ and } d(y, A'_k) \geq 1/i, k \in \mathbb{N}\}.$$

Since  $A'_i, B'_i$  are closed, we have  $A \subset \bigcup_i A_i$  and  $B \subset \bigcup_i B_i$ . For any  $x \in A_i$  and  $y \in B_i$ , we show that  $d(x, y) \geq 1/i$ , thus proving the necessary condition. By definition, we have  $x \in A'_{k_1}, d(x, B'_{k_1}) \geq 1/i$  and  $y \in B'_{k_2}, d(y, A'_{k_2}) \geq 1/i$  for some  $k_1, k_2 \in \mathbb{N}$ . W.l.o.g., we may assume  $k_1 \leq k_2$ . Since the collections are nested, we have  $x \in A'_{k_2}$ . Now, since  $d(y, A'_{k_2}) \geq 1/i$ , we have  $d(x, y) \geq 1/i$ .

To prove the sufficiency, it is sufficient to show that

$$\bigcup_{i \geq 1} \bar{A}_i \cap \bigcup_{i \geq 1} \bar{B}_i = \emptyset,$$

where  $\bar{S}$  denotes the closure of set  $S$  under metric  $d$ . Otherwise, there will be some  $x \in \bar{A}_i$  and  $x \in \bar{B}_j$  for some  $i, j$ . w.l.o.g., we can assume  $j \geq i$ . By nesting property, we now have  $x \in \bar{A}_j$  as well. However, this will imply that  $d(A_j, B_j) = 0$ , a contradiction.  $\square$

### C.3. Proofs in Section 7.3.

PROOF OF THEOREM 14. For any  $h_1, h_2 \in \mathcal{H}$ , we define a (pseudo)metric  $d(h_1, h_2) = \Pr_{X \sim \mu}[h_1(X) \neq h_2(X)]$ . Let  $\mathcal{H}' = \{h'_1, h'_2, \dots\}$  be an enumeration of  $\mathcal{H}'$ . We define  $\mathcal{H}_i = \{h \in \mathcal{H} : d(h, h'_i) = 0\}$ . Clearly  $\mathcal{G}_i = \bigcup_{j=1}^i \mathcal{H}_j$  forms a nesting of  $\mathcal{H}$  and  $(\mathcal{G}_i, \ell^{\text{OL}})$  is  $\eta$ -predictable for all  $\eta > 0$ , the sufficiency follows by Theorem 1.

By Theorem 1 and Lemma 2, we known that if  $(\mathcal{H}, \ell^{\text{OL}})$  is *e.a.s.*-predictable then for all  $\eta > 0$  we have nesting  $\{\mathcal{H}_i, i \geq 1\}$  of  $\mathcal{H}$  such that  $(\mathcal{H}_i, \ell^{\text{OL}})$  is  $\eta$ -predictable with sample size  $\leq i$ . Fix some  $0 < \eta < \frac{1}{4}$  and let  $\{\mathcal{H}_i, i \geq 1\}$  be the corresponding decomposition. We show that  $\mathcal{H}_i$  can be partitioned into countably collections such that any two functions within one collection differ by a measure zero set. This will complete the proof.

We claim that there exists  $\delta > 0$  such that

$$\inf\{d(h_1, h_2) : h_1, h_2 \in \mathcal{H}_i \text{ and } d(h_1, h_2) > 0\} \geq \delta.$$

Otherwise, there exist  $h_1, h_2 \in \mathcal{H}_i$  such that  $0 < d(h_1, h_2) < \delta_i$ , where  $\delta_i$  is choosing so that  $(1 - \delta_i)^i > 2\eta$ . Let  $A \subset (\mathbb{R}^d)^\infty$  be the event that  $h_1, h_2$  can't be distinguished within  $i$  samples. We have  $\Pr[A] > 2\eta$ . Fix some predictor  $\Phi$ , let  $p_j = \Pr[\Phi \text{ makes error after step } i \text{ on } h_j]$  for  $j \in \{1, 2\}$ . We show that  $\max\{p_1, p_2\} > \eta$ , which contradicts to the  $\eta$ -predictability of  $\mathcal{H}_i$ , thus establishes the claim. To do so, we use a probabilistic argument, let  $\mathbf{h}$  to be the random variable uniformly choosing from  $\{h_1, h_2\}$ . We only need to show

$$(C.1) \quad \mathbb{E}_{\mathbf{h}} \mathbb{E}_{\mathbf{x} \sim \mu^\infty} [1\{\Phi \text{ makes error after step } i \text{ on } \mathbf{h}\} \mid A] \geq \frac{1}{2}$$

Let  $B \subset (\mathbb{R}^d)^\infty$  be the event that there exist a sample  $Y$  after step  $i$  that first reveals  $h_1(Y) \neq h_2(Y)$ . We have  $\Pr[B] = 1$ , since  $d(h_1, h_2) > 0$ . Note that

$$(C.2) \quad \mathbb{E}_{\mathbf{x} \sim \mu^\infty} [\mathbb{E}_{\mathbf{h}} 1\{\Phi \text{ makes error after step } i \text{ on } \mathbf{h}\} \mid A, B] \geq \frac{1}{2},$$

since event  $C = \{\Phi \text{ makes error at sample } Y \text{ on } \mathbf{h} \mid A, B\}$  implies the above event and  $\mathbb{E}_{\mathbf{h}} [1\{C\}] = \frac{1}{2}$ . We now have

$$(C.3) \quad \mathbb{E}_{\mathbf{x} \sim \mu^\infty} [\mathbb{E}_{\mathbf{h}} 1\{\Phi \text{ makes error after step } i \text{ on } \mathbf{h}\} \mid A] \geq \frac{1}{2},$$

since  $\Pr[B] = 1$ . Note that (C.3) implies (C.1) by exchanging order of expectation.

Now,  $\forall h_1, h_2 \in \mathcal{H}_i$  with  $d(h_1, h_2) > 0$ , we have  $d(h_1, h_2) \geq \delta$  for some absolute constant  $\delta > 0$ . There exists a countable dense subset  $\mathcal{S}$  of all binary measurable functions over  $\mathbb{R}^d$  under our (pseudo)metric  $d(\cdot, \cdot)$ , since Borel  $\sigma$ -algebra on  $\mathbb{R}^d$  is countably generated thus separable. We can associate each  $h \in \mathcal{H}_i$  an unique element in  $\mathcal{S}$  by triangle inequality. Thus, the elements in  $\mathcal{H}_i$  can be partitioned into countably many collections such that all elements within one collection have metric zero. This establishes the necessary condition.  $\square$

## REFERENCES

- [1] ANTOS, A. and LUGOSI, G. (1998). Strong minimax lower bounds for learning. *Machine learning* **30** 31–56.
- [2] ATHREYA, K. B. and LAHIRI, S. N. (2006). *Measure theory and probability theory*. Springer Science & Business Media.
- [3] BARZDICIŪS, J. M. and FREIVALD, R. (1972). On the prediction of general recursive functions. *Doklady Akademii Nauk* **206** 521–524.
- [4] BENEDEK, G. M. and ITAI, A. (1994). Nonuniform learnability. *Journal of Computer and System Sciences* **48** 311–323.
- [5] BLUM, L. and BLUM, M. (1975). Toward a mathematical theory of inductive inference. *Information and control* **28** 125–155.
- [6] BLUMER, A., EHRENFEUCHT, A., HAUSSLER, D. and WARMUTH, M. K. (1989). Learnability and the Vapnik–Chervonenkis dimension. *Journal of the ACM (JACM)* **36** 929–965.
- [7] BOUSQUET, O., HANNEKE, S., MORAN, S., VAN HANDEL, R. and YEHUDAYOFF, A. (2020). A Theory of Universal Learning. *arXiv:2011.04483*.
- [8] COVER, T. M. et al. (1973). On determining the irrationality of the mean of a random variable. *The annals of Statistics* **1** 862–871.
- [9] DEMBO, A. and PERES, Y. (1994). A topological criterion for hypothesis testing. *The Annals of Statistics* 106–117.
- [10] GOLD, E. M. (1967). Language identification in the limit. *Information and control* **10** 447–474.
- [11] KIEFER, J. and WOLFOWITZ, J. (1958). On the deviations of the empiric distribution function of vector chance variables. *Transactions of the American Mathematical Society* **87** 173–186.
- [12] KOPLOWITZ, J., STEIF, J. E. and NERMAN, O. (1995). On Cover's consistent estimator. *Scandinavian journal of statistics* 395–397.



- [13] KULKARNI, S. R. and TSE, D. N. C. (1994). A paradigm for class identification problems. *IEEE Transactions on Information Theory* **40** 696–705.
- [14] LE CAM, L. (1973). Convergence of estimates under dimensionality restrictions. *The Annals of Statistics* **1** 38–53.
- [15] LESHEM, A. (2006). Cover’s test of rationality revisited: Computability aspects of hypothesis testing. In *2006 IEEE 24th Convention of Electrical & Electronics Engineers in Israel* 213–216. IEEE.
- [16] LINIAL, N., MANSOUR, Y. and RIVEST, R. L. (1991). Results on learnability and the Vapnik-Chervonenkis dimension. *Information and Computation* **90** 33–49.
- [17] LITTLESTONE, N. and WARMUTH, M. K. (1994). The weighted majority algorithm. *Information and computation* **108** 212–261.
- [18] MASSART, P. (1990). The tight constant in the Dvoretzky-Kiefer-Wolfowitz inequality. *The Annals of Probability* 1269–1283.
- [19] NAAMAN, M. (2016). Almost sure hypothesis testing and a resolution of the Jeffreys-Lindley paradox. *Electronic Journal of Statistics* **10** 1526–1550.
- [20] NAAMAN, M. (2021). On the tight constant in the multivariate Dvoretzky–Kiefer–Wolfowitz inequality. *Statistics & Probability Letters* **173** 109088.
- [21] RUDIN, W. (1974). *Real and Complex Analysis*. McGraw Hill Inc.,.
- [22] RUDIN, W. (2006). *Functional Analysis. International series in pure and applied mathematics*. McGraw-Hill.
- [23] SANTHANAM, N. and ANANTHARAM, V. (2015). Agnostic Insurability of Model Classes. *Journal of Machine Learning Research* **16** 2329–2355.
- [24] SCHUURMANS, D. (1997). Characterizing rational versus exponential learning curves. *Journal of computer and system sciences* **55** 140–160.
- [25] SHAWE-TAYLOR, J., BARTLETT, P. L., WILLIAMSON, R. C. and ANTHONY, M. (1998). Structural risk minimization over data-dependent hierarchies. *IEEE transactions on Information Theory* **44** 1926–1940.
- [26] VAN GAANS, O. (2003). Probability measures on metric spaces. *Lecture notes*.
- [27] VAPNIK, V. N. (2013). *The nature of statistical learning theory*. Springer science & business media.
- [28] VAPNIK, V. N. and CHERVONENKIS, A. Y. (1974). Theory of pattern recognition.
- [29] VON BAHR, B., ESSEEN, C.-G. et al. (1965). Inequalities for the  $r$  th Absolute Moment of a Sum of Random Variables,  $1 \leq r \leq 2$ . *The Annals of Mathematical Statistics* **36** 299–303.
- [30] WU, C. and SANTHANAM, N. (2019). Being correct eventually almost surely. In *Information Theory (ISIT), 2019 IEEE International Symposium on* 1989–1993. IEEE.
- [31] WU, C. and SANTHANAM, N. (2021). Non-uniform Consistency of Online Learning with Random Sampling. In *Algorithmic Learning Theory (ALT)* 1265–1285. PMLR.
- [32] ZEUGMANN, T. and ZILLES, S. (2008). Learning recursive functions: A survey. *Theoretical Computer Science* **397** 4–56.