

Non-algorithmic theory of randomness

Vladimir Vovk



практические выводы
теории вероятностей
могут быть обоснованы
в качестве следствий
гипотез о *предельной*
при данных ограничениях
сложности изучаемых явлений

On-line Compression Modelling Project (New Series)

Working Paper #25

First posted October 1, 2019. Last revised March 10, 2020.

Project web site:
<http://alrw.net>

Abstract

This paper proposes an alternative language for expressing results of the algorithmic theory of randomness. The language is more precise in that it does not involve unspecified additive or multiplicative constants, making mathematical results, in principle, applicable in practice. Our main testing ground for the proposed language is the problem of defining Bernoulli sequences, which was of great interest to Andrei Kolmogorov and his students.

The conference version is to appear in *Fields of Logic and Computation III: Essays Dedicated to Yuri Gurevich on the Occasion of His 80th Birthday* (ed. by Andreas Blass, Patrick Cégilski, Nachum Dershowitz, Manfred Droste, and Bernd Finkbeiner), Springer 2020.

Contents

1	Introduction	1
2	Testing simple statistical hypotheses	2
3	Testing composite statistical hypotheses	4
4	The validity of Bayesian statistics	5
5	Parametric Bayesian models	8
6	IID vs exchangeability: general case	9
7	IID vs exchangeability: Bernoulli sequences	11
8	Conclusion	14
	References	16

1 Introduction

There has been a great deal of criticism of the notion of p-value lately, and in particular, Glenn Shafer [20] defended the use of betting scores instead. This paper refers to betting scores as e-values and demonstrates their advantages by establishing results that become much more precise when they are stated in terms of e-values instead of p-values.

Both p-values and e-values have been used, albeit somewhat implicitly, in the algorithmic theory of randomness: Martin-Löf's tests of algorithmic randomness [16] are an algorithmic version of p-functions (i.e., functions producing p-values [6]) while Levin's tests of algorithmic randomness [13, 2] are an algorithmic version of e-functions (this is the term we will use in this paper for functions producing e-values). Levin's tests are a natural modification of Martin-Löf's tests leading to simpler mathematical results; similarly, many mathematical results stated in terms of p-values become simpler when stated in terms of e-values.

The algorithmic theory of randomness is a powerful source of intuition, but strictly speaking, its results are not applicable in practice since they always involve unspecified additive or multiplicative constants. The goal of this paper is to explore ways of obtaining results that are more precise; in particular, results that may be applicable in practice. The price to pay is that our results may involve more quantifiers (usually hidden in our notation) and, therefore, their statements may at first appear less intuitive.

In Section 2 we define p-functions and e-functions in the context of testing simple statistical hypotheses, explore relations between them, and explain the intuition behind them. In Section 3 we generalize these definitions, results, and explanations to testing composite statistical hypotheses.

Section 4 is devoted to testing in Bayesian statistics and gives non-algorithmic results that are particularly clean and intuitive. They will be used as technical tools later in the paper. In Section 5 these results are slightly extended and then applied to clarifying the difference between statistical randomness and exchangeability. (In this paper we use "statistical randomness" to refer to being produced by an IID probability measure; there will always be either "algorithmic" or "statistical" standing next to "randomness" in order to distinguish between the two meanings.)

Section 7 explores the question of defining Bernoulli sequences, which was of great interest to Kolmogorov [8], Martin-Löf [16], and Kolmogorov's other students. Kolmogorov defined Bernoulli sequences as exchangeable sequences, but we will see that another natural definition is narrower than exchangeability. A precise relation between the two definitions is deduced from a general result in Section 6, which can be regarded as another finitary analogue of de Finetti's theorem.

Kolmogorov paid particular attention to algorithmic randomness with respect to uniform probability measures on finite sets. On one hand, he believed that his notion of algorithmic randomness in this context "can be regarded as definitive" [10], and on the other hand, he never seriously suggested any gene-

realizations of this notion (and never endorsed generalizations proposed by his students). In Section 7 we state a simple result in this direction that characterizes the difference between Bernoulliness and exchangeability.

In Sections 4 and 7 we state our results first in terms of e-functions and then p-functions. Results in terms of e-functions are always simpler and cleaner, supporting Glenn Shafer’s recommendation in [20] to use betting scores more widely.

Remark 1. There is no standard terminology for what we call e-values and e-functions. In addition to Shafer’s use of “betting scores” for our e-values:

- Grünwald et al. [5] refer to e-values as “s-values” (“s” for “safe”; the expression “s-values” has been used [4] in a completely different sense, as the minus binary log of p-values).
- Gammerman and Vovk [3] refer to the reciprocals of e-values as “i-values” (“i” for “integral”). This term and its variations were used widely in discussions in the Department of Computer Science at Royal Holloway, University of London, around 2000: cf., e.g., “i-test” [18] and “i-randomness” [28]. They are natural in view of the common term “integral tests” for Levin-type tests (used, e.g., in [15] starting from the first edition in 1993).

Our “e-value” is motivated by expectation playing a role similar to that of probability in “p-value” [21, Section 3.8].

No formal knowledge of the algorithmic theory of randomness will be assumed in this paper; the reader can safely ignore all comparisons between our results and results of the algorithmic theory of randomness.

Notation

Our notation will be mostly standard or defined at the point where it is first used. If \mathcal{F} is a class of $[0, \infty]$ -valued functions on some set Ω and $g : [0, \infty] \rightarrow [0, \infty]$ is a function, we let $g(\mathcal{F})$ stand for the set of all compositions $g(f) = g \circ f$, $f \in \mathcal{F}$ (i.e., g is applied to \mathcal{F} element-wise). We will also use obvious modifications of this definition: e.g., $0.5\mathcal{F}^{-0.5}$ would be interpreted as $g(\mathcal{F})$, where $g(u) := 0.5u^{-0.5}$ for $u \in [0, \infty]$.

2 Testing simple statistical hypotheses

Let P be a probability measure on a measurable space Ω . A *p-function* [6] is a measurable function $f : \Omega \rightarrow [0, \infty]$ such that, for any $\epsilon > 0$, $P\{f \leq \epsilon\} \leq \epsilon$. (Since P is a probability measure, we can assume, without loss of generality, that f takes values in $[0, 1]$.) An *e-function* is a measurable function $f : \Omega \rightarrow [0, \infty]$ such that $\int f dP \leq 1$. (As already mentioned, e-functions have been promoted in [20] and [5], and also in [22, Section 11.5], using different terminology.)

Let \mathcal{P}_P be the class of all p-functions and \mathcal{E}_P be the class of all e-functions, where the underlying measure P is shown as subscript. We can define p-values

and e-values as values taken by p-functions and e-functions, respectively. The intuition behind p-values and e-values will be discussed later in this section.

The following is an algorithm-free version of the standard relation (see, e.g., [15, Lemma 4.3.5] or [23, Theorem 43]) between Martin-Löf's and Levin's algorithmic notions of randomness deficiency.

Proposition 2. *For any probability measure P and $\kappa \in (0, 1)$,*

$$\kappa \mathcal{P}_P^{\kappa-1} \subseteq \mathcal{E}_P \subseteq \mathcal{P}_P^{-1}. \quad (1)$$

Proof. The right inclusion in (1) follows from the Markov inequality: if f is an e-function,

$$P\{f^{-1} \leq \epsilon\} = P\{f \geq 1/\epsilon\} \leq \epsilon. \quad (2)$$

The left inclusion in (1) follows from [22, Section 11.5]. The value of the constant in front of the $\mathcal{P}_P^{\kappa-1}$ on the left-hand side of (1) follows from $\int_0^1 p^{\kappa-1} dp = 1/\kappa$. \square

Both p-functions and e-functions can be used for testing statistical hypotheses. In this section we only discuss *simple statistical hypotheses*, i.e., probability measures. Observing a large e-value or a small p-value with respect to a simple statistical hypothesis P entitles us to rejecting P as the source of the observed data, provided the e-function or p-function were chosen in advance. The e-value can be interpreted as the amount of evidence against P found by our chosen e-function. Similarly, the p-value reflects the amount of evidence against P on a different scale; small p-values reflect a large amount of evidence against P .

Remark 3. Proposition 2 tells us that using p-values and using e-values are equivalent, on a rather crude scale. Roughly, a p-value of p corresponds to an e-value of $1/p$. The right inclusion in (2) says that any way of producing e-values e can be translated into a way of producing p-values $1/e$. On the other hand, the left inclusion in (2) says that any way of producing p-values p can be translated into a way of producing e-values $\kappa p^{\kappa-1} \approx 1/p$, where the “ \approx ” assumes that we are interested in the asymptotics as $p \rightarrow 0$, $\kappa > 0$ is small, and we ignore positive constant factors (as customary in the algorithmic theory of randomness).

Remark 4. Proposition 2 can be greatly strengthened, under the assumptions of Remark 3. For example, we can replace (1) by

$$H_\kappa(\mathcal{P}_P) \subseteq \mathcal{E}_P \subseteq \mathcal{P}_P^{-1},$$

where

$$H_\kappa(v) := \begin{cases} \infty & \text{if } v = 0 \\ \kappa(1 + \kappa)^\kappa v^{-1} (-\ln v)^{-1-\kappa} & \text{if } v \in (0, e^{-1-\kappa}] \\ 0 & \text{if } v \in (e^{-1-\kappa}, 1] \end{cases} \quad (3)$$

and $\kappa \in (0, \infty)$ (see [22, Section 11.1]). The value of the coefficient $\kappa(1 + \kappa)^\kappa$ in (3) follows from

$$\int_0^{e^{-1-\kappa}} v^{-1} (-\ln v)^{-1-\kappa} dv = \frac{1}{\kappa(1 + \kappa)^\kappa}.$$

We can rewrite (1) in Proposition 2 as

$$\kappa^{-1}\mathcal{P}_P^{1-\kappa} \subseteq \mathcal{E}_P^{-1} \subseteq \mathcal{P}_P, \quad (4)$$

as

$$\mathcal{E}_P^{-1} \subseteq \mathcal{P}_P \subseteq \kappa^{\frac{1}{1-\kappa}} \mathcal{E}_P^{\frac{1}{\kappa-1}}, \quad (5)$$

and as

$$\mathcal{E}_P \subseteq \mathcal{P}_P^{-1} \subseteq \kappa^{\frac{1}{\kappa-1}} \mathcal{E}_P^{\frac{1}{1-\kappa}}. \quad (6)$$

3 Testing composite statistical hypotheses

Let Ω be a measurable space, which we will refer to as our *sample space*, and Θ be another measurable space (our *parameter space*). We say that $P = (P_\theta \mid \theta \in \Theta)$ is a *statistical model* on Ω if P is a Markov kernel with source Θ and target Ω : each P_θ is a probability measure on Ω , and for each measurable $A \subseteq \Omega$, the function $P_\theta(A)$ of $\theta \in \Theta$ is measurable.

The notions of an e-function and a p-function each split in two. We are usually really interested only in the outcome ω , while the parameter θ is an auxiliary modelling tool. This motivates the following pair of simpler definitions. A measurable function $f : \Omega \rightarrow [0, \infty]$ is an *e-function* with respect to the statistical model P (which is our *composite statistical hypothesis* in this context) if

$$\forall \theta \in \Theta : \int_{\Omega} f(\omega) P_\theta(d\omega) \leq 1.$$

In other words, if $P^*(f) \leq 1$, where P^* is the upper envelope

$$P^*(f) := \sup_{\theta \in \Theta} \int_{\Omega} f(\omega) P_\theta(d\omega) \quad (7)$$

(in Bourbaki's [1, IX.1.1] terminology, P^* is an encumbrance provided the integral in (7) is understood as the upper integral). Similarly, a measurable function $f : \Omega \rightarrow [0, 1]$ is a *p-function* with respect to the statistical model P if, for any $\epsilon > 0$,

$$\forall \theta \in \Theta : P_\theta\{\omega \in \Omega \mid f(\omega) \leq \epsilon\} \leq \epsilon.$$

In other words, if, for any $\epsilon > 0$, $P^*(1_{\{f \leq \epsilon\}}) \leq \epsilon$.

Let \mathcal{E}_P be the class of all e-functions with respect to the statistical model P , and \mathcal{P}_P be the class of all p-functions with respect to P . We can easily generalize Proposition 2 (the proof stays the same).

Proposition 5. *For any statistical model P and $\kappa \in (0, 1)$,*

$$\kappa \mathcal{P}_P^{\kappa-1} \subseteq \mathcal{E}_P \subseteq \mathcal{P}_P^{-1}.$$

For $f \in \mathcal{E}_P$, we regard the e-value $f(\omega)$ as the amount of evidence against the statistical model P found by f (which must be chosen in advance) when the outcome is ω . The interpretation of p-values is similar.

In some case we would like to take the parameter θ into account more seriously. A measurable function $f : \Omega \times \Theta \rightarrow [0, \infty]$ is a *conditional e-function* with respect to the statistical model P if

$$\forall \theta \in \Theta : \int_{\Omega} f(\omega; \theta) P_{\theta}(d\omega) \leq 1.$$

Let $\bar{\mathcal{E}}_P$ be the class of all such functions. And a measurable function $f : \Omega \times \Theta \rightarrow [0, 1]$ is a *conditional p-function* with respect to P if

$$\forall \epsilon > 0 \forall \theta \in \Theta : P_{\theta} \{ \omega \in \Omega \mid f(\omega; \theta) \leq \epsilon \} \leq \epsilon.$$

Let $\bar{\mathcal{P}}_P$ be the class of all such functions.

We can embed \mathcal{E}_P (resp. \mathcal{P}_P) into $\bar{\mathcal{E}}_P$ (resp. $\bar{\mathcal{P}}_P$) by identifying a function f on domain Ω with the function f' on domain $\Omega \times \Theta$ that does not depend on $\theta \in \Theta$, $f'(\omega; \theta) := f(\omega)$.

For $f \in \bar{\mathcal{E}}_P$, we can regard $f(\omega; \theta)$ as the amount of evidence against the specific probability measure P_{θ} in the statistical model P found by f when the outcome is ω .

We can generalize Proposition 5 further as follows.

Proposition 6. *For any statistical model P and $\kappa \in (0, 1)$,*

$$\kappa \bar{\mathcal{P}}_P^{\kappa-1} \subseteq \bar{\mathcal{E}}_P \subseteq \bar{\mathcal{P}}_P^{-1}. \quad (8)$$

Remarks 3 and 4 are also applicable in the context of Propositions 5 and 6.

4 The validity of Bayesian statistics

In this section we establish the validity of Bayesian statistics in our framework, mainly as a sanity check. We will translate the results in [31], which are stated in terms of the algorithmic theory of randomness, to our algorithm-free setting. It is interesting that the proofs simplify radically, and become almost obvious. (And remarkably, one statement also simplifies.)

Let $P = (P_{\theta} \mid \theta \in \Theta)$ be a statistical model, as in the previous section, and Q be a probability measure on the parameter space Θ . Together, P and Q form a *Bayesian model*, and Q is known as the *prior measure* in this context.

The joint probability measure T on the measurable space $\Omega \times \Theta$ is defined by

$$T(A \times B) := \int_B P_{\theta}(A) Q(d\theta),$$

for all measurable $A \subseteq \Omega$ and $B \subseteq \Theta$. Let Y be the marginal distribution of T on Ω : for any measurable $A \subseteq \Omega$, $Y(A) := T(A \times \Theta)$.

The *product* $\bar{\mathcal{E}}_P \mathcal{E}_Q$ of $\bar{\mathcal{E}}_P$ and \mathcal{E}_Q is defined as the class of all measurable functions $f : \Omega \times \Theta \rightarrow [0, \infty]$ such that, for some $g \in \bar{\mathcal{E}}_P$ and $h \in \mathcal{E}_Q$,

$$f(\omega, \theta) = g(\omega; \theta) h(\theta) \quad T\text{-a.s.} \quad (9)$$

Such f can be regarded as ways of finding evidence against (ω, θ) being produced by the Bayesian model (P, Q) : to have evidence against (ω, θ) being produced by (P, Q) we need to have evidence against θ being produced by the prior measure Q or evidence against ω being produced by P_θ ; we combine the last two amounts of evidence by multiplying them. The following proposition tells us that this product is precisely the amount of evidence against T found by a suitable e-function.

Proposition 7. *If $(P_\theta \mid \theta \in \Theta)$ is a statistical model with a prior probability measure Q on Θ , and T is the joint probability measure on $\Omega \times \Theta$, then*

$$\mathcal{E}_T = \bar{\mathcal{E}}_P \mathcal{E}_Q. \quad (10)$$

Proposition 7 will be deduced from Theorem 14 in Section 5. It is the analogue of Theorem 1 in [31], which says, in the terminology of that paper, that the level of impossibility of a pair (θ, ω) with respect to the joint probability measure T is the product of the level of impossibility of θ with respect to the prior measure Q and the level of impossibility of ω with respect to the probability measure P_θ . In an important respect, however, Proposition 7 is simpler than Theorem 1 in [31]: in the latter, the level of impossibility of ω with respect to P_θ has to be conditional on the level of impossibility of θ with respect to Q , whereas in the former there is no such conditioning. Besides, Proposition 7 is more precise: it does not involve any constant factors (specified or otherwise).

Remark 8. The non-algorithmic formula (10) being simpler than its counterpart in the algorithmic theory of randomness is analogous to the non-algorithmic formula $H(x, y) = H(x) + H(y \mid x)$ being simpler than its counterpart $K(x, y) = K(x) + K(y \mid x, K(x))$ in the algorithmic theory of complexity, H being entropy and K being prefix complexity. The fact that $K(x, y)$ does not coincide with $K(x) + K(y \mid x)$ to within an additive constant, K being Kolmogorov complexity, was surprising to Kolmogorov and wasn't noticed for several years [7, 8].

The *inf-projection* onto Ω of an e-function $f \in \mathcal{E}_T$ with respect to T is the function $(\text{proj}_\Omega^{\text{inf}} f) : \Omega \rightarrow [0, \infty]$ defined by

$$(\text{proj}_\Omega^{\text{inf}} f)(\omega) := \inf_{\theta \in \Theta} f(\omega, \theta).$$

Intuitively, $\text{proj}_\Omega^{\text{inf}} f$ regards ω as typical under the model if it can be extended to a typical (ω, θ) for at least one θ . Let $\text{proj}_\Omega^{\text{inf}} \mathcal{E}_T$ be the set of all such inf-projections.

The results in the rest of this section become simpler if the definitions of classes \mathcal{E} and \mathcal{P} are modified slightly: we drop the condition of measurability on their elements and replace all integrals by upper integrals and all measures by outer measures. We will use the modified definitions only in the rest of this section (we could have used them in the whole of this paper, but they become particularly useful here since projections of measurable functions do not have to be measurable [24]).

Proposition 9. *If T is a probability measure on $\Omega \times \Theta$ and Y is its marginal distribution on Ω ,*

$$\mathcal{E}_Y = \text{proj}_\Omega^{\text{inf}} \mathcal{E}_T. \quad (11)$$

Proof. To check the inclusion “ \subseteq ” in (11), let $g \in \mathcal{E}_Y$, i.e., $\int g(\omega)Y(d\omega) \leq 1$. Setting $f(\omega, \theta) := g(\omega)$, we have $\int f(\omega, \theta)T(d\omega, d\theta) \leq 1$ (i.e., $f \in \mathcal{E}_T$) and g is the inf-projection of f onto Ω .

To check the inclusion “ \supseteq ” in (11), let $f \in \mathcal{E}_T$ and $g := \text{proj}_\Omega^{\text{inf}} f$. We then have

$$\int g(\omega)Y(d\omega) = \int g(\omega)T(d\omega, d\theta) \leq \int f(\omega, \theta)T(d\omega, d\theta) \leq 1. \quad \square$$

Proposition 9 says that we can acquire evidence against an outcome ω being produced by the Bayesian model (P, Q) if and only if we can acquire evidence against (ω, θ) being produced by the model for all $\theta \in \Theta$.

We can combine Propositions 7 and 9 obtaining

$$\mathcal{E}_Y = \text{proj}_\Omega^{\text{inf}} (\bar{\mathcal{E}}_P \mathcal{E}_Q).$$

The rough interpretation is that we can acquire evidence against ω being produced by Y if and only if we can, for each $\theta \in \Theta$, acquire evidence against θ being produced by Q or acquire evidence against ω being produced by P_θ .

The following statements in terms of p-values are cruder, but their interpretation is similar.

Corollary 10. *If $\kappa \in (0, 1)$ and (P, Q) is a Bayesian model,*

$$\kappa^{-1} \mathcal{P}_T^{1-\kappa} \subseteq \bar{\mathcal{P}}_P \mathcal{P}_Q \subseteq \kappa^{\frac{2}{1-\kappa}} \mathcal{P}_T^{\frac{1}{1-\kappa}}.$$

Proof. We will use the restatements (4) and (5) of Proposition 2 and similar restatements of Propositions 6 and 5. Therefore, by (10) in Proposition 7,

$$\kappa^{-1} \mathcal{P}_T^{1-\kappa} \subseteq \mathcal{E}_T^{-1} = (\bar{\mathcal{E}}_P \mathcal{E}_Q)^{-1} = \bar{\mathcal{E}}_P^{-1} \mathcal{E}_Q^{-1} \subseteq \bar{\mathcal{P}}_P \mathcal{P}_Q$$

and

$$\bar{\mathcal{P}}_P \mathcal{P}_Q \subseteq \kappa^{\frac{2}{1-\kappa}} (\bar{\mathcal{E}}_P \mathcal{E}_Q)^{\frac{1}{\kappa-1}} = \kappa^{\frac{2}{1-\kappa}} \mathcal{E}_T^{\frac{1}{\kappa-1}} \subseteq \kappa^{\frac{2}{1-\kappa}} \mathcal{P}_T^{\frac{1}{1-\kappa}}. \quad \square$$

Corollary 11. *If $\kappa \in (0, 1)$, T is a probability measure on $\Omega \times \Theta$, and Y is its marginal distribution on Ω ,*

$$\kappa^{-1} \text{proj}_\Omega^{\text{sup}} \mathcal{P}_T^{1-\kappa} \subseteq \mathcal{P}_Y \subseteq \kappa^{\frac{1}{1-\kappa}} \text{proj}_\Omega^{\text{sup}} \mathcal{P}_T^{\frac{1}{1-\kappa}},$$

where $\text{proj}_\Omega^{\text{sup}}$ is defined similarly to $\text{proj}_\Omega^{\text{inf}}$ (with sup in place of inf).

Proof. As in the proof of Corollary 10, we have

$$\kappa^{-1} \text{proj}_\Omega^{\text{sup}} \mathcal{P}_T^{1-\kappa} \subseteq \text{proj}_\Omega^{\text{sup}} \mathcal{E}_T^{-1} = \mathcal{E}_Y^{-1} \subseteq \mathcal{P}_Y$$

and

$$\mathcal{P}_Y \subseteq \kappa^{\frac{1}{1-\kappa}} \mathcal{E}_Y^{\frac{1}{\kappa-1}} = \kappa^{\frac{1}{1-\kappa}} \text{proj}_\Omega^{\text{sup}} \mathcal{E}_T^{\frac{1}{\kappa-1}} \subseteq \kappa^{\frac{1}{1-\kappa}} \text{proj}_\Omega^{\text{sup}} \mathcal{P}_T^{\frac{1}{1-\kappa}}. \quad \square$$

5 Parametric Bayesian models

Now we generalize the notion of a Bayesian model to that of a parametric Bayesian (or *para-Bayesian*) model. This is a pair consisting of a statistical model $(P_\theta \mid \theta \in \Theta)$ on a sample space Ω and a statistical model $(Q_\pi \mid \pi \in \Pi)$ on the sample space Θ (so that the sample space of the second statistical model is the parameter space of the first statistical model). Intuitively, a para-Bayesian model is the counterpart of a Bayesian model in the situation of uncertainty about the prior: now the prior is a parametric family of probability measures rather than one probability measure.

The following definitions are straightforward generalizations of the definitions for the Bayesian case. The joint statistical model $T = (T_\pi \mid \pi \in \Pi)$ on the measurable space $\Omega \times \Theta$ is defined by

$$T_\pi(A \times B) := \int_B P_\theta(A) Q_\pi(d\theta), \quad (12)$$

for all measurable $A \subseteq \Omega$ and $B \subseteq \Theta$. For each $\pi \in \Pi$, Y_π is the marginal distribution of T_π on Ω : for any measurable $A \subseteq \Omega$, $Y_\pi(A) := T_\pi(A \times \Theta)$. The *product* $\mathcal{E}_P \mathcal{E}_Q$ of \mathcal{E}_P and \mathcal{E}_Q is still defined as the class of all measurable functions $f : \Omega \times \Theta \rightarrow [0, \infty]$ such that, for some $g \in \mathcal{E}_P$ and $h \in \mathcal{E}_Q$, we have the equality in (9) T_π -a.s., for all $\pi \in \Pi$.

Remark 12. Another representation of para-Bayesian models is as a sufficient statistic, as elaborated in [12]:

- For the para-Bayesian model (P, Q) , the statistic $(\theta, \omega) \in (\Theta \times \Omega) \mapsto \theta$ is a sufficient statistic in the statistical model (T_π) on the product space $\Theta \times \Omega$.
- If θ is a sufficient statistic for a statistical model (T_π) on a sample space Ω , then (P, Q) is a para-Bayesian model, where Q is the distribution of θ , and P_θ are (fixed versions of) the conditional distributions given θ .

Remark 13. Yet another way to represent a para-Bayesian model (P, Q) is a Markov family with time horizon 3:

- the initial state space is Π , the middle one is Θ , and the final one is Ω ;
- there is no initial probability measure on Π , the statistical model (Q_π) is the first Markov kernel, and the statistical model (P_θ) is the second Markov kernel.

Theorem 14. *If (P, Q) is a para-Bayesian model with the joint statistical model T (as defined by (12)), we have (10).*

Proof. The inclusion “ \supseteq ” in (10) follows from the definition of T : if $g \in \mathcal{E}_P$ and $h \in \mathcal{E}_Q$, we have, for all $\pi \in \Pi$,

$$\int_{\Omega \times \Theta} g(\omega; \theta) h(\theta) T_\pi(d\omega, d\theta) = \int_{\Theta} \int_{\Omega} g(\omega; \theta) P_\theta(d\omega) h(\theta) Q_\pi(d\theta)$$

$$\leq \int_{\Theta} h(\theta) Q_{\pi}(d\theta) \leq 1.$$

To check the inclusion “ \subseteq ” in (10), let $f \in \mathcal{E}_T$. Define $h : \Theta \rightarrow [0, \infty]$ and $g : \Omega \times \Theta \rightarrow [0, \infty]$ by

$$\begin{aligned} h(\theta) &:= \int f(\omega, \theta) P_{\theta}(d\omega) \\ g(\omega; \theta) &:= f(\omega, \theta)/h(\theta) \end{aligned}$$

(setting, e.g., $0/0 := 0$ in the last fraction). Since by definition, $f(\omega, \theta) = g(\omega; \theta)h(\theta)$ T_{π} -a.s., it suffices to check that $h \in \mathcal{E}_Q$ and $g \in \mathcal{E}_P$. The inclusion $h \in \mathcal{E}_Q$ follows from the fact that, for any $\pi \in \Pi$,

$$\int_{\Theta} h(\theta) Q_{\pi}(d\theta) = \int_{\Theta} \int_{\Omega} f(\omega, \theta) P_{\theta}(d\omega) Q_{\pi}(d\theta) = \int_{\Omega \times \Theta} f(\omega, \theta) T_{\pi}(d\omega, d\theta) \leq 1.$$

And the inclusion $g \in \bar{\mathcal{E}}_P$ follows from the fact that, for any $\theta \in \Theta$,

$$\int g(\omega; \theta) P_{\theta}(d\omega) = \int \frac{f(\omega, \theta)}{h(\theta)} P_{\theta}(d\omega) = \frac{\int f(\omega, \theta) P_{\theta}(d\omega)}{h(\theta)} = \frac{h(\theta)}{h(\theta)} \leq 1$$

(we have ≤ 1 rather than $= 1$ because of the possibility $h(\theta) = 0$). \square

6 IID vs exchangeability: general case

De Finetti’s theorem (see, e.g., [19, Theorem 1.49]) establishes a close connection between IID and exchangeability for infinite sequences in \mathbf{Z}^{∞} , where \mathbf{Z} is a Borel measurable space: namely, the exchangeable probability measures are the convex mixtures of the IID probability measures (in particular, their upper envelopes, and therefore, e- and p-functions, coincide). This section discusses a somewhat less close connection in the case of sequences of a fixed finite length.

Fix $N \in \{1, 2, \dots\}$ (time horizon), and let $\Omega := \mathbf{Z}^N$ be the set of all sequences of elements of \mathbf{Z} (a measurable space, not necessarily Borel) of length N . An *IID probability measure* on Ω is a measure of the type Q^N , where Q is a probability measure on \mathbf{Z} . The *configuration* $\text{conf}(\omega)$ of a sequence $\omega \in \Omega$ is the multiset of all elements of ω , and a *configuration measure* is the pushforward of an IID probability measure on Ω under the mapping conf . Therefore, a configuration measure is a measure on the set of all multisets in \mathbf{Z} of size N (with the natural quotient σ -algebra).

Let \mathcal{E}_{iid} be the class of all e-functions with respect to the family of all IID probability measures on Ω and $\mathcal{E}_{\text{conf}}$ be the class of all e-functions with respect to the family of all configuration probability measures. Let $\mathcal{E}_{\text{exch}}$ be the class of all e-functions with respect to the family of all exchangeable probability measures on Ω ; remember that a probability measure P on Ω is *exchangeable* if, for any permutation $\pi : \{1, \dots, N\} \rightarrow \{1, \dots, N\}$ and any measurable set $E \subseteq \mathbf{Z}^N$,

$$P\{(z_1, \dots, z_N) \mid (z_{\pi(1)}, \dots, z_{\pi(N)}) \in E\} = P(E).$$

The *product* $\mathcal{E}_{\text{exch}}\mathcal{E}_{\text{conf}}$ of $\mathcal{E}_{\text{exch}}$ and $\mathcal{E}_{\text{conf}}$ is the set of all measurable functions $f : \Omega \rightarrow [0, \infty]$ such that, for some $g \in \mathcal{E}_{\text{exch}}$ and $h \in \mathcal{E}_{\text{conf}}$,

$$f(\omega) = g(\omega)h(\text{conf}(\omega))$$

holds for almost all $\omega \in \Omega$ (under any IID probability measure).

Corollary 15. *It is true that*

$$\mathcal{E}_{\text{iid}} = \mathcal{E}_{\text{exch}}\mathcal{E}_{\text{conf}}. \quad (13)$$

Proof. It suffices to apply Theorem 14 in the situation where Θ is the set of all configurations, P_θ is the probability measure on \mathbf{Z}^N concentrated on the set of all sequences with the configuration θ and uniform on that set (we can order θ arbitrarily, and then P_θ assigns weight $1/N!$ to each permutation of that ordering), Π is the set of all IID probability measures on Ω , and Q_π is the pushforward of $\pi \in \Pi$ with respect to the mapping conf . \square

Corollary 15 is the non-algorithmic analogue of Theorem 3 of [29], given without a proof.

The next theorem gives the ranges of \mathcal{E}_{iid} , $\mathcal{E}_{\text{exch}}$, and $\mathcal{E}_{\text{conf}}$. For any set of functions \mathcal{F} we set

$$\sup \mathcal{F} := \sup_{f \in \mathcal{F}} \sup f;$$

i.e., $\sup \mathcal{F}$ is the supremum of the values attained by the functions in \mathcal{F} . Remember that the length N of the sequences considered in this section is fixed.

Theorem 16. *Suppose $|\mathbf{Z}| \geq N$. Then*

$$\sup \mathcal{E}_{\text{iid}} = N^N, \quad (14)$$

$$\sup \mathcal{E}_{\text{exch}} = N! \sim (2\pi N)^{1/2}(N/e)^N, \quad (15)$$

$$\sup \mathcal{E}_{\text{conf}} = N^N/N! \sim (2\pi N)^{-1/2}e^N, \quad (16)$$

where the two “ \sim ” refer to the asymptotics as $N \rightarrow \infty$.

Theorem 16 shows that (13) remain true when we put \sup in front of each of the three function classes. The counterpart of (16) in the algorithmic theory of randomness is Theorem 4 in [29].

A crude interpretation of Corollary 15 and Theorem 16 is that the condition of being an IID sequence can be split into two components: exchangeability and having an iid configuration; the first component is more important.

Proof of Theorem 16. For any $\omega \in \Omega$, $\sup_{f \in \mathcal{E}_{\text{iid}}} f(\omega)$ is attained at f that takes a non-zero value only at ω . Therefore,

$$\sup \mathcal{E}_{\text{iid}} = \sup_{\omega \in \Omega} \frac{1}{\sup_P P(\{\omega\})},$$

P ranging over the IID probability measures. The supremum will not change if P ranges over the probability measures on \mathbf{Z} concentrated on the elements of the sequence ω , which we will assume. Consider an ω consisting of n distinct elements of \mathbf{Z} . Order these distinct elements, and let m_i , $i = 1, \dots, n$, be the number of times the i th of these elements occurs in ω . Using the maximum likelihood estimate for the multinomial model, we can see that

$$\frac{1}{\sup_P P(\{\omega\})} = (N/m_1)^{m_1} \dots (N/m_n)^{m_n} = \frac{N^N}{m_1^{m_1} \dots m_n^{m_n}}.$$

The supremum of the last expression is attained when $n = N$ and $m_1 = \dots = m_n = 1$, and it is equal to N^N . This completes the proof of (14).

A similar argument also works for (16). We have

$$\sup \mathcal{E}_{\text{conf}} = \sup_{m_1, \dots, m_n} \frac{(N/m_1)^{m_1} \dots (N/m_n)^{m_n}}{\binom{N}{m_1, \dots, m_n}} = \frac{N^N/N!}{(m_1^{m_1}/m_1!) \dots (m_n^{m_n}/m_n!)} \quad (17)$$

Since $(m^m)/m! \geq 1$ for all $m \in \{1, 2, \dots\}$ and $(m^m)/m! > 1$ for all $m \in \{2, 3, \dots\}$, the second supremum in (17) is also attained when $n = N$ and $m_1 = \dots = m_n = 1$, which completes the proof of (16).

As for (15), it suffices to notice that

$$\sup_{m_1, \dots, m_n} \binom{N}{m_1, \dots, m_n} = \sup_{m_1, \dots, m_n} \frac{N!}{m_1! \dots m_n!}$$

is attained at $n = N$ and $m_1 = \dots = m_n = 1$.

The asymptotic equivalences in (15) and (16) follow from Stirling's formula. \square

Since the suprema in Theorem 16 are attained at functions that zero everywhere except one point, we have the following corollary.

Corollary 17. *If $|\mathbf{Z}| \geq N$,*

$$\begin{aligned} \inf \mathcal{P}_{\text{id}} &= N^{-N}, \\ \inf \mathcal{P}_{\text{exch}} &= 1/N! \sim (2\pi N)^{-1/2} (e/N)^N, \\ \inf \mathcal{P}_{\text{conf}} &= N!/N^N \sim (2\pi N)^{1/2} e^{-N}. \end{aligned}$$

7 IID vs exchangeability: Bernoulli sequences

In this section we apply the definitions and results of the previous sections to the problem of defining Bernoulli sequences. Kolmogorov's main publications on this topic are [8] and [9]. The results of this section will be algorithm-free versions of the results in [25] (also described in V'yugin's review [32], Sections 11–13).

The definitions of the previous section simplify as follows. Now $\Omega := \{0, 1\}^N$ is the set of all binary sequences of length N . Let $\mathcal{E}_{\text{Bern}}$ be the class of all e-functions with respect to the family of all Bernoulli IID probability measures on Ω (this is a special case of \mathcal{E}_{iid}) and \mathcal{E}_{bin} be the class of all e-functions with respect to the family of all binomial probability measures on $\{0, \dots, N\}$ (this is a special case of $\mathcal{E}_{\text{conf}}$); remember that the Bernoulli measure B_p with parameter $p \in [0, 1]$ is defined by $B_p(\{\omega\}) := p^k(1-p)^{N-k}$, where $k := +\omega$ is the number of 1s in ω , and the binomial measure bin_p with parameter $p \in [0, 1]$ is defined by $\text{bin}_p(\{k\}) := \binom{N}{k} p^k(1-p)^{N-k}$. (The notation $+\omega$ for the number k of 1s in ω is motivated by k being the sum of the elements of ω .)

We continue to use the notation $\mathcal{E}_{\text{exch}}$ for the class of all e-functions with respect to the family of all exchangeable probability measures on Ω ; a probability measure P on Ω is exchangeable if and only if $P(\{\omega\})$ depends on ω only via $+\omega$. It is clear that a function $f : \Omega \rightarrow [0, \infty]$ is in $\mathcal{E}_{\text{exch}}$ if and only if, for each $k \in \{0, \dots, N\}$,

$$\binom{N}{k}^{-1} \sum_{\omega \in \Omega: +\omega=k} f(\omega) \leq 1.$$

The product $\mathcal{E}_{\text{exch}}\mathcal{E}_{\text{bin}}$ of $\mathcal{E}_{\text{exch}}$ and \mathcal{E}_{bin} is the set of all functions $\omega \in \Omega \mapsto g(\omega)h(+\omega)$ for $g \in \mathcal{E}_{\text{exch}}$ and $h \in \mathcal{E}_{\text{bin}}$. The following is a special case of Corollary 15.

Corollary 18. *It is true that*

$$\mathcal{E}_{\text{Bern}} = \mathcal{E}_{\text{exch}}\mathcal{E}_{\text{bin}}.$$

The intuition behind Corollary 18 is that a sequence $\omega \in \Omega$ is Bernoulli if and only if it is exchangeable and the number of 1s in it is binomial. The analogue of Corollary 18 in the algorithmic theory of randomness is Theorem 1 in [25], which says, using the terminology of that paper, that the Bernoulliness deficiency of ω equals the binomiality deficiency of $+\omega$ plus the conditional randomness deficiency of ω in the set of all sequences in $\{0, 1\}^N$ with $+\omega$ 1s given the binomiality deficiency of $+\omega$. Corollary 18 is simpler since it does not involve any analogue of the condition “given the binomiality deficiency of $+\omega$ ”.

Remark 19. Kolmogorov’s definition of Bernoulli sequences is via exchangeability. We can regard this definition as an approximation to definitions taking into account the binomiality of the number of 1s. In the paper [8] Kolmogorov uses the word “approximately” when introducing his notion of Bernoulliness (p. 663, lines 5–6 after the 4th displayed equation). However, it would be wrong to assume that here he acknowledges disregarding the requirement that the number of 1s should be binomial; this is not what he meant when he used the word “approximately” [11].

The reason for Kolmogorov’s definition of Bernoulliness being different from the definitions based on e-values and p-values is that $+\omega$ carries too much information about ω ; intuitively [26], $+\omega$ contains not only useful information

about the probability p of 1 but also noise. To reduce the amount of noise, we will use an imperfect estimator of p . Set

$$p(a) := \sin^2 \left(aN^{-1/2} \right), \quad a = 1, \dots, N^* - 1, \quad N^* := \left\lfloor \frac{\pi}{2} N^{1/2} \right\rfloor, \quad (18)$$

where $\lfloor \cdot \rfloor$ stands for integer part. Let $E : \{0, \dots, N\} \rightarrow [0, 1]$ be the estimator of p defined by $E(k) := p(a)$, where $p(a)$ is the element of the set (18) that is nearest to k/N among those satisfying $p(a) \leq k/N$; if such elements do not exist, set $E(k) := p(1)$.

Denote by \mathfrak{A} the partition of the set $\{0, \dots, N\}$ into the subsets $E^{-1}(E(k))$, where $k \in \{0, \dots, N\}$. For any $k \in \{0, \dots, N\}$, $\mathfrak{A}(k) := E^{-1}(E(k))$ denotes the element of the partition \mathfrak{A} containing k . Let \mathcal{E}_{sin} be the class of all e-functions with respect to the statistical model $\{U_k \mid k \in \{0, \dots, N\}\}$, U_k being the uniform probability measure on $\mathfrak{A}(k)$. (This is a Kolmogorov-type statistical model, consisting of uniform probability measures on finite sets; see, e.g., [30, Section 4].)

Theorem 20. *For some universal constant $c > 0$,*

$$c^{-1} \mathcal{E}_{\text{sin}} \subseteq \mathcal{E}_{\text{bin}} \subseteq c \mathcal{E}_{\text{sin}}.$$

The analogue of Theorem 20 in the algorithmic theory of randomness is Theorem 2 in [25], and the proof of Theorem 20 can be extracted from that of Theorem 2 in [25] (details omitted).

Remark 21. Paper [25] uses a net slightly different from (18); (18) was introduced in [26] and also used in [32].

To state corollaries in terms of p-values of Corollary 18 and Theorem 20, we will use the obvious notation $\mathcal{P}_{\text{Bern}}$, $\mathcal{P}_{\text{exch}}$, and \mathcal{P}_{bin} .

Corollary 22. *For each $\kappa \in (0, 1)$,*

$$\kappa^{-1} \mathcal{P}_{\text{Bern}}^{1-\kappa} \subseteq \mathcal{P}_{\text{exch}} \mathcal{P}_{\text{bin}} \subseteq \kappa^{\frac{2}{1-\kappa}} \mathcal{P}_{\text{Bern}}^{\frac{1}{1-\kappa}}. \quad (19)$$

Proof. Similarly to Corollary 10, the left inclusion of (19) follows from

$$\kappa^{-1} \mathcal{P}_{\text{Bern}}^{1-\kappa} \subseteq \mathcal{E}_{\text{Bern}}^{-1} = \mathcal{E}_{\text{exch}}^{-1} \mathcal{E}_{\text{bin}}^{-1} \subseteq \mathcal{P}_{\text{exch}} \mathcal{P}_{\text{bin}},$$

and the right inclusion of (19) follows from

$$\mathcal{P}_{\text{exch}} \mathcal{P}_{\text{bin}} \subseteq \kappa^{\frac{2}{1-\kappa}} (\mathcal{E}_{\text{exch}} \mathcal{E}_{\text{bin}})^{\frac{1}{\kappa-1}} = \kappa^{\frac{2}{1-\kappa}} \mathcal{E}_{\text{Bern}}^{\frac{1}{\kappa-1}} \subseteq \kappa^{\frac{2}{1-\kappa}} \mathcal{P}_{\text{Bern}}^{\frac{1}{1-\kappa}}. \quad \square$$

Corollary 23. *There is a universal constant $c > 0$ such that, for each $\kappa \in (0, 0.9)$,*

$$c\kappa^{-1} \mathcal{P}_{\text{sin}}^{1-\kappa} \subseteq \mathcal{P}_{\text{bin}} \subseteq c^{-1} \kappa^{\frac{1}{1-\kappa}} \mathcal{P}_{\text{sin}}^{\frac{1}{1-\kappa}}. \quad (20)$$

Proof. As in the previous proof, the left inclusion of (20) follows from

$$\kappa^{-1}\mathcal{P}_{\text{sin}}^{1-\kappa} \subseteq \mathcal{E}_{\text{sin}}^{-1} \subseteq c^{-1}\mathcal{E}_{\text{bin}}^{-1} \subseteq c^{-1}\mathcal{P}_{\text{bin}},$$

and the right inclusion from

$$\mathcal{P}_{\text{bin}} \subseteq \kappa^{\frac{1}{1-\kappa}}\mathcal{E}_{\text{bin}}^{\frac{1}{\kappa-1}} \subseteq c^{-1}\kappa^{\frac{1}{1-\kappa}}\mathcal{E}_{\text{sin}}^{\frac{1}{\kappa-1}} \subseteq c^{-1}\kappa^{\frac{1}{1-\kappa}}\mathcal{P}_{\text{sin}}^{\frac{1}{1-\kappa}},$$

where c stands for a positive universal constant. \square

In conclusion of this section, let us state the binary version of Theorem 16 and its corollary.

Theorem 24. *Suppose $N \in \{2, 4, \dots\}$ is an even number. Then*

$$\sup \mathcal{E}_{\text{Bern}} = 2^N, \quad (21)$$

$$\sup \mathcal{E}_{\text{exch}} = \binom{N}{N/2} \sim (\pi N/2)^{-1/2} 2^N, \quad (22)$$

$$\sup \mathcal{E}_{\text{bin}} = 2^N / \binom{N}{N/2} \sim (\pi N/2)^{1/2}, \quad (23)$$

where the two “ \sim ” again refer to the asymptotics as $N \rightarrow \infty$.

Proof. The argument is similar to that in the proof of Theorem 16. The supremum in (21) is attained at the function that takes value 2^N at the sequence $0 \dots 01 \dots 1$ ($N/2$ 0s followed by $N/2$ 1s) and is zero everywhere else. Replacing 2^N by $\binom{N}{N/2}$, we obtain a function attaining the supremum in (22). The supremum in (23) is attained at the function on $\{0, \dots, N\}$ that takes value $2^N / \binom{N}{N/2}$ at $N/2$ and is zero everywhere else. Finally, the asymptotic equivalences follow from Stirling’s formula. \square

We can see that $\sup \mathcal{E}_{\text{bin}}$ (given by (23)) is much smaller than $\sup \mathcal{E}_{\text{Bern}}$ and $\sup \mathcal{E}_{\text{exch}}$ (given by (21) and (22), respectively). This might be interpreted as exchangeability being the main component of Bernoulliness.

Corollary 25. *If N is an even number,*

$$\begin{aligned} \inf \mathcal{P}_{\text{Bern}} &= 2^{-N}, \\ \inf \mathcal{P}_{\text{exch}} &= 1 / \binom{N}{N/2} \sim (\pi N/2)^{1/2} 2^{-N}, \\ \inf \mathcal{P}_{\text{bin}} &= \binom{N}{N/2} 2^{-N} \sim (\pi N/2)^{-1/2}. \end{aligned}$$

8 Conclusion

In this section we discuss some directions of further research. A major advantage of the non-algorithmic approach to randomness proposed in this paper is the

absence of unspecified constants; in principle, all constants can be computed. The most obvious open problem is to find the best constant c in Theorem 20.

In Section 7 we discussed a possible implementation of Kolmogorov's idea of defining Bernoulli sequences. However, Kolmogorov's idea was part of a wider programme; e.g., in [9, Section 5] he sketches a way of applying a similar approach to Markov sequences. For other possible applications, see [30, Section 4] (most of these applications were mentioned by Kolmogorov in his papers and talks). Analogues of Corollary 18 in Section 7 can be established for these other applications (cf. [12] and Remark 12), but it is not obvious whether Theorem 20 can be extended in a similar way.

Acknowledgments

Thanks to Wouter Koolen, Vladimir V'yugin, Alex Gammerman, and Ilia Nouretdinov for useful discussions. This research was partially supported by Astra Zeneca and Stena Line.

Appendix: Non-algorithmic theory of complexity

The definitions of p-functions and e-functions given at the beginning of Section 2 can be applied, without changing a word, to any measure P on Ω , without the restriction $P(\Omega) = 1$. It is interesting that, for some P , such generalizations also have useful applications. In particular, the following generalization of Proposition 2 includes an algorithm-free version of standard relations (see, e.g., [14, Theorem 4]) between plain and prefix Kolmogorov complexities.

Proposition 26. *Let $h : [0, \infty) \rightarrow [0, \infty)$ be a continuous function that is strictly decreasing over $\{h > 0\}$ and satisfies $\int h \leq 1$, where \int stands for the integration with respect to the Lebesgue measure. For any measure P ,*

$$h(\mathcal{P}_P) \subseteq \mathcal{E}_P \subseteq \mathcal{P}_P^{-1}. \quad (24)$$

Proof. The right inclusion in (24) still follows from the Markov inequality (2). As for the left inclusion, we have, for any $f \in \mathcal{P}_P$ and any h satisfying the conditions of the proposition,

$$\begin{aligned} \int h(f) \, dP &= \int_0^\infty P(h(f) \geq c) \, dc = \int_0^\infty P(f \leq h^{-1}(c)) \, dc \\ &\leq \int_0^\infty h^{-1}(c) \, dc = \int h^{-1} = \int h \leq 1, \end{aligned}$$

where the last equality follows from Fubini's theorem. \square

An example of a function h satisfying Proposition 26 is

$$h(c) := \begin{cases} \frac{\kappa}{2} c^{\kappa-1} & \text{if } c \leq 1 \\ \frac{\kappa}{2} c^{-\kappa-1} & \text{if } c \geq 1, \end{cases} \quad (25)$$

where $\kappa \in (0, 1)$ is a constant.

Let P be the counting measure on \mathbb{N} . An example of $f \in \mathcal{P}_P$ is $f := 2^{C+1}$, where C is plain Kolmogorov complexity; this function f is the smallest, to within a constant factor, upper semicomputable element of \mathcal{P}_P (see [23, Theorem 8]). An example of $m \in \mathcal{E}_P$ is the largest, to within a constant factor, lower semicomputable measure on \mathbb{N} (see [23, Section 4.2]). Proposition 26 applied to the function (25) gives

$$\frac{\kappa}{2}(2^{C+1})^{-\kappa-1} \leq^{\times} m \leq^{\times} (2^{C+1})^{-1},$$

where \leq^{\times} stands for inequality to within a constant factor. The last equation can be rewritten as

$$C \leq^+ K \leq^+ (1 + \kappa)C,$$

where K is prefix complexity and \leq^+ stands for inequality to within an additive constant. (Better inequalities can be obtained if we use h of a form similar to (3).)

Remark 27. The main reason [17] for using “i-values” instead of “e-values” in the late 1990s and early 2000s (see Remark 1) was the desire to cover the case of measures P that are not necessarily probability measures, such as counting measures, which makes “integral” more appropriate than “expectation”.

References

- [1] Nicolas Bourbaki. *Elements of Mathematics. Integration*. Springer, Berlin, 2004. In two volumes. The French originals published in 1952–1969.
- [2] Peter Gács. Uniform test of algorithmic randomness over a general space. *Theoretical Computer Science*, 341:91–137, 2005.
- [3] Alex Gammerman and Vladimir Vovk. Data labelling apparatus and method thereof, 2003. US Patent Application 0236578 A1. Available on the Internet (accessed in February 2020).
- [4] Sander Greenland. Valid P-values behave exactly as they should: some misleading criticisms of P-values and their resolution with S-values. *American Statistician*, 73(S1):106–114, 2019.
- [5] Peter Grünwald, Rianne de Heide, and Wouter M. Koolen. Safe testing. Technical Report arXiv:1906.07801 [math.ST], arXiv.org e-Print archive, June 2019.
- [6] Yuri Gurevich and Vladimir Vovk. Test statistics and p-values. *Proceedings of Machine Learning Research*, 105:89–104, 2019. COPA 2019.
- [7] Andrei N. Kolmogorov. Несколько теорем об алгоритмической энтропии и алгоритмическом количестве информации. Успехи математических наук, 23(2):201, 1968. Abstract of a talk before the Moscow Mathematical Society. Meeting of 31 October 1967.

- [8] Andrei N. Kolmogorov. Logical basis for information theory and probability theory. *IEEE Transactions on Information Theory*, IT-14:662–664, 1968. Russian original: К логическим основам теории информации и теории вероятностей, published in Проблемы передачи информации.
- [9] Andrei N. Kolmogorov. Combinatorial foundations of information theory and the calculus of probabilities. *Russian Mathematical Surveys*, 38:29–40, 1983. Russian original: Комбинаторные основания теории информации и исчисления вероятностей.
- [10] Andrei N. Kolmogorov. On logical foundations of probability theory. In Yu. V. Prokhorov and K. Itô, editors, *Probability Theory and Mathematical Statistics*, volume 1021 of *Lecture Notes in Mathematics*, pages 1–5. Springer, 1983. Talk at the Fourth USSR–Japan Symposium on Probability Theory and Mathematical Statistics (Tbilisi, August 1982) recorded by Alexander A. Novikov, Alexander K. Zvonkin, and Alexander Shen. Our quote follows *Selected Works of A. N. Kolmogorov*, volume II, Probability Theory and Mathematical Statistics, edited by A. N. Shirayev, Kluwer, Dordrecht, p. 518.
- [11] Andrei N. Kolmogorov. Personal communication, ca 1983.
- [12] Steffen L. Lauritzen. *Extremal Families and Systems of Sufficient Statistics*, volume 49 of *Lecture Notes in Statistics*. Springer, New York, 1988.
- [13] Leonid A. Levin. Uniform tests of randomness. *Soviet Mathematics Doklady*, 17:337–340, 1976. Russian original: Равномерные тесты случайности.
- [14] Leonid A. Levin. Various measures of complexity for finite objects (axiomatic description). *Soviet Mathematics Doklady*, 17:522–526, 1976. Russian original: О различных мерах сложности конечных объектов: аксиоматическое описание.
- [15] Ming Li and Paul Vitányi. *An Introduction to Kolmogorov Complexity and Its Applications*. Springer, fourth edition, 2019.
- [16] Per Martin-Löf. The definition of random sequences. *Information and Control*, 9:602–619, 1966.
- [17] Ilija Nourtdinov. Personal communication, February 2020.
- [18] Ilija Nourtdinov, Vladimir V’yugin, and Alex Gammerman. Transductive Confidence Machine is universal. In Ricard Gavaldà, Klaus P. Jantke, and Eiji Takimoto, editors, *Proceedings of the Fourteenth International Conference on Algorithmic Learning Theory*, volume 2842 of *Lecture Notes in Artificial Intelligence*, pages 283–297, Berlin, 2003. Springer.
- [19] Mark J. Schervish. *Theory of Statistics*. Springer, New York, 1995.

- [20] Glenn Shafer. The language of betting as a strategy for statistical and scientific communication. Technical Report arXiv:1903.06991 [math.ST], arXiv.org e-Print archive, March 2019.
- [21] Glenn Shafer. On the nineteenth-century origins of significance testing and p-hacking. The Game-Theoretic Probability and Finance project, <http://probabilityandfinance.com>, Working Paper 55, September 2019, first posted July 2019.
- [22] Glenn Shafer and Vladimir Vovk. *Game-Theoretic Foundations for Probability and Finance*. Wiley, Hoboken, NJ, 2019.
- [23] Alexander Shen, Vladimir A. Uspensky, and Nikolai Vereshchagin. *Kolmogorov Complexity and Algorithmic Randomness*. American Mathematical Society, Providence, RI, 2017.
- [24] Mikhail Ya. Souslin. Sur une définition des ensembles mesurables B sans nombres transfinis. *Comptes rendus hebdomadaires des séances de l'Académie des sciences*, 164:88–91, 1917.
- [25] Vladimir Vovk. On the concept of the Bernoulli property. *Russian Mathematical Surveys*, 41:247–248, 1986. Russian original: О понятии бернуллиевости. Another English translation with proofs: [27].
- [26] Vladimir Vovk. Learning about the parameter of the Bernoulli model. *Journal of Computer and System Sciences*, 55:96–104, 1997.
- [27] Vladimir Vovk. On the concept of Bernoulliness. Technical Report arXiv:1612.08859 [math.ST], arXiv.org e-Print archive, December 2016.
- [28] Vladimir Vovk and Alex Gammerman. Statistical applications of algorithmic randomness. In *Bulletin of the International Statistical Institute. The 52nd Session, Contributed Papers*, volume LVIII, book 3, pages 469–470, 1999.
- [29] Vladimir Vovk, Alex Gammerman, and Craig Saunders. Machine-learning applications of algorithmic randomness. In *Proceedings of the Sixteenth International Conference on Machine Learning*, pages 444–453, San Francisco, CA, 1999. Morgan Kaufmann.
- [30] Vladimir Vovk and Glenn Shafer. Kolmogorov’s contributions to the foundations of probability. *Problems of Information Transmission*, 39:21–31, 2003.
- [31] Vladimir Vovk and Vladimir V. V’yugin. On the empirical validity of the Bayesian method. *Journal of the Royal Statistical Society B*, 55:253–266, 1993.
- [32] Vladimir V. V’yugin. Algorithmic complexity and stochastic properties of finite binary sequences. *Computer Journal*, 42:294–317, 1999.