

Fairness and Randomness in Machine Learning: Statistical Independence and Relativization

RABANUS DERR* AND ROBERT C. WILLIAMSON

Abstract

Fair Machine Learning endeavors to prevent unfairness arising in the context of machine learning applications embedded in society. To this end, several mathematical fairness notions have been proposed. The most known and used notions turn out to be expressed in terms of statistical independence, which is taken to be a primitive and unambiguous notion. However, two choices remain (and are largely unexamined to date): what exactly is the meaning of statistical independence and what are the groups to which we ought to be fair? We answer both questions by leveraging Richard Von Mises' theory of probability, which starts with data, and then builds the machinery of probability from the ground up. In particular, his theory places a relative definition of randomness as statistical independence at the center of statistical modelling. Much in contrast to the classically used, absolute i.i.d.-randomness, which turns out to be “orthogonal” to his conception. We show how Von Mises' frequential modeling approach fits well to the problem of fair machine learning and show how his theory (suitably interpreted) demonstrates the equivalence between the contestability of the choice of groups in the fairness criterion and the contestability of the choice of relative randomness. We thus conclude that the problem of being fair in machine learning is precisely as hard as the problem of defining what is meant by being random. In both cases there is a consequential choice, yet there is no universal “right” choice possible.

KEYWORDS AND PHRASES: Fairness in machine learning, Richard von Mises, Randomness, Statistical independence.

1. INTRODUCTION

Under the name “Fair Machine Learning” researchers have attempted to tackle problems of injustice, fairness, discrimination arising in the context of machine learning applications embedded in society [5]. Despite the variety of definitions of fairness and proposed “fair algorithms,” we still lack a conceptual understanding of fairness in machine learning [75, 84]. What does it mean for predictions to be fair? How does the statistical frame influence fairness? Is there fair data and what would it look like? For instance more concretely, how does a population of individuals and their corresponding predictions look like if a provided definition of fairness is fulfilled?

We focus on a collection of widely used fairness notions which are based on statistical independence e.g., [16, 46, 22], but examine them from a new perspective. Surprisingly, debates concerning these notions have not questioned the role and meaning of statistical independence upon which they are based. As we shall argue, statistical independence is far from being a mathematical concept linked to one unique interpretation (see §4.2). This paper, in contrast to much of the literature on fairness in machine learning, e.g., [16, 33, 46, 22], investigates what many definitions of fairness take for granted: a well-defined and meaningful notion of statistical independence.

Another, less popular, strand of research investigates the role of randomness in machine learning [94]. The standard randomness notion, independently and often identically distributed data points, suffers the longstanding critique of being inadequate (cf. [89]). Again, statistical independence lies at the foundation of this, hitherto unrelated to fairness, concept. (We will justify below our use of “randomness” as used here.)

At the core of both our observations is the unreflective use of a convenient mathematical theory of probability. Kolmogorov's axiomatization of probability theory, developed 1933 in his book [55] (translated in [56]), dominates most research in machine learning. As Kolmogorov explicitly stated, his theory was designed as a purely axiomatic, mathematical theory detached from meaning and interpretation. In particular, Kolmogorov's statistical independence lacks such reference. However, the modeling nature of machine learning and the arising ethical complications within machine learning applied in society ask for semantics of probabilistic notions.

In this work, we focus on statistical independence. We leverage a theory of probability axiomatized by Von Mises [107] in order to obtain *meaningful* access to probabilistic notions. (In leaning on Von Mises we are directly following the explicit advice of Kolmogorov [56, page 3, footnote 4].) This theory construes probability theory as “sci-

*Corresponding author.

entific”¹ (as opposed to purely mathematical) with the aim to describe the world and provide interpretations and verifiability [107, pages 1 and 14]. Von Mises’ theory of probability provides a mathematical definition of statistical independence which describes statistical phenomena observable in the world. In particular, Von Mises’ statistical independence is mathematically, but not conceptually, related to Kolmogorov’s definition.

In this paper, we, to the best of our knowledge, are the first to apply Von Mises’ randomness to machine learning and to interpret randomness in machine learning in a Von Mises’ way. The paper is structured as follows:

In Section 2, we outline our statistical perspective on machine learning. We present the “independent and identically distributed”-assumption (i.i.d.-assumption) as one commonly used choice for modeling randomness. The further occurrence of statistical independence as fundamental ingredient of fairness notions in machine learning (§3) pushes us to the question: “How to interpret statistical independence in (fair) machine learning?” Remarkably, “Independence” governs many discussions around fairness in machine learning without getting to a concrete meaning of this term. Its deeper semantics remain untouched even in the considerably exhaustive book by Barocas et al. [5, Chapter 3, p. 13].

We first dissect Kolmogorov’s widely used definition of statistical independence in Section 4 before we propose another mathematical notion following Von Mises. Von Mises uses his notion of independence in order to define randomness. We contrast his definition and the i.i.d.-assumption in Section 5. This reveals a general typification of mathematical definitions of randomness which most importantly differ in the *absoluteness* respectively *relativity* to the problem under consideration (§5.2).

Finally, we leverage Von Mises’ definition of statistical independence to redefine three fairness notions from machine learning (§6). Against the background of Von Mises’ probability theory, we then link randomness and fairness both expressed as statistical independence (§7). Thereby, we reveal an unexpected hypothesis: randomness and fairness can be considered equivalent concepts in machine learning. Randomness becomes a *relative*, even an ethical choice. Fairness, however, turns out to be a modeling assumption about the data used in the machine learning system.

Due to the frequent use of the word “independence” with different meanings in this paper, we differentiate. By “independence” we mean an abstract concept of unrelatedness and non-influence [92]. We use it interchangeably with “statistical independence” which emphasizes the probabilistic

and statistical context. When referring to the later introduced, formal definitions of statistical independence following Kolmogorov or Von Mises we explicitly state this. Finally, we assign “Independence” (capital “I”) to one of the fairness criteria in machine learning which demands for statistical independence of predictions and sensitive attribute [5, 83]. Despite the abstract appearance of this work, we consider it as part of the project to make the very abstract notion of “independence” or “statistical independence” at least a bit more concrete — specifically, we describe independence in terms of samples, not abstractions such as countably additive probability measures. We ask (and try to assist) the reader to become aware of the implicit assumptions taken about the concept of “independence”.

2. A STATISTICAL PERSPECTIVE ON MACHINE LEARNING

Machine Learning ingests data and provides decisions or inferences. In this sense, at its core, it is statistics.² Adopting this perspective, we understand machine learning as “modeling data generative processes”. Statistics, respectively machine learning, asks for properties of data generating processes given a collection of data [109, p. ix], [4, p. 1].

We assume that a data generative process occurs somehow in the world. We are confronted with a collection of numbers, the data, produced by the process and acquired by measurement. A “model” is a mathematical description of such a data generating process. This description should allow us to make predictions in an algorithmic fashion. Thus, we require, *inter alia*, a mathematical description of data — “a model for data collection” [18, p. 207]. What is arguably *the* standard model of data is stated in [30, p. 11]:

We shall assume in this book that $(X_1, Y_1), \dots, (X_n, Y_n)$, the data, is a sequence of independent identically distributed (i.i.d.) random pairs with the same distribution[...].

Similar definitions can be found in many machine learning or statistics textbooks (e.g., [18, Def. 5.1.1]). Data (measurements from the world) is conceived of (mathematically) as a collection of random variables which share the same distribution and which are (statistically) independent to each other. Implicit in this definition is that the data indeed *has* a stable distribution. The assumed independence can be interpreted as presumption of randomness of the data. Each data point was “drawn independently”³ from all others, with the

²Machine learning could be viewed as classical statistics with a stronger focus on algorithmic realizations (cf. [91, p. 6] or [58]). While there are ML approaches that do not seem statistical in the classical sense (e.g., worst case online sequence prediction), our general description still holds.

³Observe that as a mathematical theory of data this already leaves a lot to be desired: you will not find in any text a precise and constructive explanation of this process of “drawn independently”. To be sure, the notion of “statistical independence” is well defined (see further below), as is a collection of random variables. But not the mysterious process

¹We further use the term “scientific” in order to describe a theory modeling a phenomenon in the world providing interpretations and verifiability in the sense of Von Mises (spelled out more concretely in Section 4.3). Any more detailed discussion, e.g., along the lines of [76], is out of the scope of this paper.

obvious interpretation that each data point does not give a hint about the value of any other.

The i.i.d. assumption is two-fold: 1) the assumption of identical distributions of a sample, and 2) the mutual independence of points in the sample. The second assumption alone captures and pertains to randomness [50, Section 3]. However, since the use of i.i.d. is more common, we refer to this more specific assumption.

Even though many results in statistics and machine learning rely on the i.i.d. assumption (e.g., law of large numbers and central limit theorem in statistics [18], generalization bounds in statistical learning theory [91] and computationally feasible expressions in probabilistic machine learning [30]), it has always been subject of fundamental critique.⁴ Other randomness definitions are rarely applied, but exceptions exist [108, 96].

In summary, statistical conclusions often rely on the i.i.d.-description of data. This description embraces a model of randomness making randomness a substantial assumption about the data in statistics and machine learning. Interestingly, statistical independence lies at the foundation of another, hitherto unrelated, concept: many fairness criteria in fair machine learning are expressed in terms of statistical independence.

3. FAIR MACHINE LEARNING RELIES ON STATISTICAL INDEPENDENCE

With the broad use of machine learning algorithms in many socially relevant domains, e.g., recidivism risk prediction or algorithmic hiring, machine learning algorithms turned out to be part of discriminatory practices [22, 80]. These revelations were accompanied by the rise of an entire research field, called “fair machine learning” (cf. [16, 22, 5]). We do not attempt to summarize this large literature here. Instead, we simply take a snapshot of the most widely known fairness criteria in machine learning [5, p. 45].

3.1 Three Fairness Criteria in Machine Learning

The three so called observational fairness criteria, which are expressed in terms of statistical independence, encompass a large part of fair machine learning literature:⁵

of “drawing from”. The closest we can come to such a description of mechanism is the indirect abstract version: the data are created (by the world) in a manner that the (mathematical theorem) of the law of large numbers holds, and that their empirical distribution converges to the distribution which was given in the first place.

⁴Nicely summarized by Glenn Shafer in a comment on [40] “The i.i.d. case has also been central to statistics ever since Jacob Bernoulli proved the law of large numbers at the end of the 17th century, but its inadequacy was always obvious.” [89].

⁵Obviously, there are other fairness notions in machine learning which we have not listed above and which are not expressed as statistical independence, e.g., [33, 53, 111].

Independence demands that the predictions \hat{Y} are statistically independent of group membership S in socially salient, morally relevant groups [16, 52, 33].

Separation is formalized as $\hat{Y} \perp S|Y$, i.e., the prediction \hat{Y} is conditionally independent of the sensitive attribute S given the true label Y [46].

Sufficiency is fulfilled if and only if $Y \perp S|\hat{Y}$ [54].

For the sake of distinguishing between the fairness criteria “Independence” and statistical independence, we henceforth mark all fairness criteria by a leading capital letter. Each of the notions appear in a variety of ways and under different names [5, p. 45ff]. From the perspective of ethics, the fairness criteria have been substantiated via loss egalitarianism [11, 111], absence of discrimination [11], affirmative action [9, 83] or equality of opportunity [46].

Certainly, statistical independence is not equivalent to fairness in general (a constellation of concepts sharing a common name, and perhaps little else agreed by all) [33, 46, 83]. The nature of fairness has been discussed for decades, e.g., in political philosophy [82], moral philosophy [61] and actuarial science [1]. The “essentially contested” nature of fairness suggests that no universal, statistical criterion of fairness exists [39]. How fairness should be defined is a context-specific decision [84, 47].

Nevertheless, in order to incorporate fairness notions into algorithmic tools we require mathematical formalisations of fairness definitions. The three named criteria dominate most of the practical fair machine learning tools [5, p. 45], presumably because their simple definitions make it easy to incorporate them in learning procedures in a pre-, in- or post-processing way [5, Chapter 3, p. 20]. Regarding both the reductionist definition of fairness and the pragmatic justification, we emphasize that *our argumentation is solely with respect to the fairness criteria named above*.

The three fairness criteria are described as *group* fairness notions since each of the definitions is intrinsically relativized with respect to sensitive groups. The definition of sensitive groups substantially influences the notion of fairness. For instance, via custom categorization one can provide fairness by group-design (see [67, Section H.3] for a detailed discussion of the question of choice of groups). In addition, the meaning of groups in a societal context influences the choice of groups as elaborated in [49], and as explored in a long line of work in social psychology [17, 97, 48, 66, 19]. We contribute to the debate by drawing a connection between the choice of groups and the choice of randomness in §7.1.1.

3.2 Independence in Mathematics and the World

Behind the formalization of fairness as statistical independence, there is an apparently rigid, mathematical definition of statistical independence. The fairness criteria presume that we have the machinery of probability theory at

our disposal and the relationship of mathematics and the world is clear and unambiguous. However, as we elaborate further in the following, there is no single notion of “the” mathematical theory of probability [36]. Furthermore, it is not clear what it means to be statistically independent when talking of measurements in the world. Respectively, it is not obvious that the standard formulation of statistical independence is the right one to use.

The current definitions of fairness in machine learning fail and even hurt in practice, because debates on fairness notions take for granted a commonly agreed meaning of statistical independence. This common ground fatally does not exist given just the standard probability theory. Hence, given a specified scenario, e.g., hiring algorithm for public service in Kenya, fair machine learning research should enable founded debates on the meaning and purposefulness of deploying a specific notion of fairness, in particular those which require independence statements because their are algorithmically attractive. If the debate in Kenya would rely on everybody’s – who are involved in the deployment process – intuitive understanding of statistical independence, it would result in a fatal misalignment of reasonings.

We perceive this work as part of the project to emphasize the substantive character formal notions of fairness should have. In other words, fairness is a societal and ethical concept that does not allow for the separation of machine learning as technical tool on one side and the idea of fairness in society on the other side. Debates on fairness have to consider machine learning *in* society.

Hence, if we desire statistical independence to capture a fairness notion applicable to the world, we ought to understand what the mathematical formulae signify in the world. Thus, in addition to the debate about the fairness criteria, a debate on the interpretation of statistical concepts in ethical context is required.

In this work, we contribute to the understanding by scrutinizing the standard definition of statistical independence. Motivated by the occurrence of statistical independence as fundamental ingredient in randomness as well as fairness in machine learning, we first detail the standard account due to Kolmogorov. What is statistical independence? How does statistical independence relate to an independence in the world?

4. STATISTICAL INDEPENDENCE REVISITED

To make any sense of phenomena in our world, we need to ignore large parts of it in order to avoid being overwhelmed. Hence, we usually assume or presume the phenomena of interest depends only on a few factors and to be *independent* of everything else [72]. Thus the concept of independence is inherent in a variety of subjects ranging from causal reasoning [93], to logic [42], accounting [24], public law [69] and many more. Independence, as we understand it, grasps the

concept of incapability of an entity to be expressed, derived or deduced by something else [92].

4.1 From Independence to Statistical Independence

Of special interest to us is the concept of independence in probability theories and statistics [60][36, Section IIF, IIIG and VH]. Independence in a probabilistic context should somehow capture the unrelatedness between the occurrence of events, as has been understood for centuries:

Two Events are independent, when they have no connexion [sic] one with the other, and that the happening of one neither forwards nor obstructs the happening of the other. [28, Introduction, p. 6].

Modern probability theory loosely follows this intuition as we see in the following.

4.2 Statistical Independence As We Know It Lacks Semantics

Since the axiomatization of probability theory developed by [55] (translated in [56]), it displaced many other approaches. Mathematically, Kolmogorov’s measure-theoretic axiomatization dominates all other mathematical formalizations to probability and related concepts.⁶ In particular, his definition of statistical independence developed a well-accepted, ubiquitous notion. In a simple form it is given by:⁷

Definition 1 (Simplified Statistical Independence following Kolmogorov). Two events A, B are statistically independent iff

$$P(A \cap B) = P(A)P(B).$$

Independence plays a central role in Kolmogorov’s probability theory: “measure theory ends and probability begins with the definition of independence” [32, p. 37] (quoted in [98]). However, Kolmogorov’s definition of independence is subtle and requires closer investigation.

We employ a small toy example in order to convey the semantic emptiness of Kolmogorov’s definition: consider the experiment of throwing a die. The events under observation are $A = \{1, 2\}$, seeing one or two pips, respectively $B = \{2, 3\}$, seeing two or three pips. If the die were fair, so each face has equal probability $\frac{1}{6}$ to show up, the events A and B would turn out to not be independent. In contrast, if the die were loaded in a very special way $p_2 = p_3 = \frac{1}{2}$, $p_1 = p_4 = p_5 = p_6 = 0$, where p_i refers to the probability of seeing i pips, the events A and B would be independent following Kolmogorov’s definition.

Thus, statistical independence, even though defined over events, manifests in the correspondence of how probabilities are mapped to events and why. The definition focuses on events. But, the crucial ingredient is the probability. Thus,

⁶Exceptions exist, e.g., in quantum theory [44] or in statistics [108].

⁷For a rigorous definition see Appendix 5.

there is no unique interpretation of statistical independence. A more detailed interpretation and meaning heavily depends on the interpretation of probability in the first place.

Given this observation, we may ask for a notion of independence in the world, which is somehow captured by Kolmogorov’s definition. Kolmogorov himself underlined the avowedly mathematical, axiomatic nature of probability. His theory is in principle detached from any meaning of probabilistic concepts such as statistical independence in the world [56, p. 1]. He even questioned the validity of his axioms as reasonable descriptions of the world [56, p. 17].

However, one can possibly construct a notion of independence in world which is captured by Kolmogorov’s definition. If one assumes one’s calculations about one’s beliefs on the happening of events are governed by the mathematical rules laid out by Kolmogorov, then Kolmogorov’s definition of statistical independence captures one’s (in-the-world) understanding of an independence of beliefs on the happening of events. This sketch of a purely subjectivist account to statistical independence neglects a justification for the choice of mathematical formulation and skips over any reference to an objective world. In conclusion, Kolmogorov’s independence might capture a worldly concept. But, this independence in the world is not uniquely attached to Kolmogorov’s definition.

There is a third major irritation arising from Kolmogorov’s definition. As observed already above, Kolmogorov treats events as the entities of independence. Though, against the background of De Moivre [28]’s intuition on statistical independence, we wonder about this focus. Statistical independence, as De Moivre [28] already emphasized, refers to altering “the happening of the event,” but not the event itself. It is not the independence between the shown numbers of the die (whatever this means), but the independence of the processes how the number showed up (loading the die, throwing the die, etc.) which are captured by statistical independence.

This critique is not new. Von Mises already criticized the measure theoretical definition by Kolmogorov in his book [107, pp. 36–39]. In summary, he argued that there is no interpretation of statistical independence of “single events.” The unrelatedness, which the probabilistic notion of statistical independence is trying to capture, locates in the process of reoccurring events, but not single events themselves. More recently, Von Collani [104] argued in a similar way. It is probability which brings the definition of statistical independence to life and it is the question what probability means and why we use it which links the mathematical definition to a concept in the real world.

In machine learning and statistics it is often presumed that the mathematical definition of independence captures a worldly concept. As we argued in this section, this link is far from being well-defined. However, if we consider machine learning as worldly data modeling, then the natural question arises: what do we model when we leverage Kolmogorov’s statistical independence? What do we mean by

independence of events? In order to circumvent these questions, we propose to look into another mathematical theory of probability. This theory was led by the idea of modeling statistical phenomena in the world.

4.3 Statistical Independence and a Probability Theory with Inherent Semantics

Around 15 years before Kolmogorov’s *Grundbegriffe der Wahrscheinlichkeitsrechnung* [55] (translated as [56]), Von Mises proposed an earlier axiomatization of probability theory [105]. His less known theory approached the problem of a mathematical theory of probability through the lens of physics. Von Mises aimed for a “mathematical theory of repetitive events.”

This aim included the emphasis on the link between real-world and mathematical idealization. In particular, he offers interpretability and verifiability of his theory [107, p. 1 and 14].⁸ For interpretation he defined probabilities in a frequency-based way (see Definition 2). This inherently reflects the repetitive nature of the phenomena under description. By verifiability he referred to the ability to *approximately* verify the probabilistic statements made about the world [107, p. 45].

In summary, Von Mises’ theory, in our conception of machine learning, starts the “modeling of data generating processes” on an even more fundamental level than it is currently done via the use of Kolmogorov’s axiomatization. His aim for a mathematical description of data-generating processes (the sequence of repetitive events) aligns to our perspective on machine learning as laid out earlier (cf. Section 2). With Von Mises we obtain access to *meaningful* foundations for statistical concepts in machine learning. In particular, we redefine and reinterpret statistical independence in a Von Misesian way. This suggests new perspectives on the problem of fair machine learning and the concepts of fairness and randomness themselves.

For the further discussion, we summarize the major ingredients of Von Mises’ theory. Fortunately, it turns out that Von Mises’ notion of statistical independence, central to our discussion, is mathematically analogous to the well-known Kolmogorovian definition. Thus, one’s intuition on statistical independence is refined but its mathematical applicability remains.

4.4 Von Mises’ Theory of Probability and Randomness in a Nutshell

Von Mises’ axiomatization of probability theory is based on random sequences of events and the interpretation of

⁸Von Mises’ discussion pre-dates most of the modern work done in the philosophy of science. Popper [76], for instance, was inspired by Von Mises’ conception of randomness and probability. While Von Mises’ definition of verifiability and interpretability is arguably somewhat ill-conceived, nevertheless, he works, in contrast to Kolmogorov’s theory which is of purely syntactical nature, on a semantical project, where the connection of the real world and the mathematical formulation is of vital importance.

probability as the limiting frequency that an event occurs in such a sequence [105]. These random sequences, called *collectives*, are the main ingredients of his theory. For the sake of simplicity, we stick to binary collectives. Thus, collectives are 0-1-sequences with certain randomness properties which define probabilities for the labels 0 and 1. Nevertheless, it is possible to define collectives respectively probabilities on richer label sets. Collectives can be extended up to a continuum [107, II.B].

For notational economy, we note that a sequence taking values in $\{0, 1\}$, $(s_i)_{i \in \mathbb{N}}$ can be identified with a *function* $s: \mathbb{N} \rightarrow \{0, 1\}$.

Definition 2 (Collective [107, p. 12]). Let \mathcal{S} be a set of sequences $s: \mathbb{N} \rightarrow \{0, 1\}$ with $s(j) = 1$ for infinitely many $j \in \mathbb{N}$. In mathematical terms, *collectives* with respect to \mathcal{S} are sequences $x: \mathbb{N} \rightarrow \{0, 1\}$ for which the following two conditions hold.

1. The limit of relative frequencies of 1s,

$$\lim_{n \rightarrow \infty} \frac{|\{i \in \mathbb{N}: x(i) = 1, 1 \leq i \leq n\}|}{n}$$

exists.⁹ If it exists, then the limit of relative frequencies of 0s exists, too. We define p_1 , respectively $p_0 = 1 - p_1$, to be its value.

2. For all $s \in \mathcal{S}$,

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{|\{i \in \mathbb{N}: x(i) = 1 \text{ and } s(i) = 1, 1 \leq i \leq n\}|}{|\{j \in \mathbb{N}: s(j) = 1, 1 \leq j \leq n\}|} \\ = \lim_{n \rightarrow \infty} \frac{|\{i \in \mathbb{N}: x(i) = 1, 1 \leq i \leq n\}|}{n} = p_1. \end{aligned}$$

We call p_0 the *probability* of label 0. Conversely, p_1 is the probability of label 1.¹⁰ The existence of the limit (Condition 1) is a non-vacuous condition. One can easily construct sequences whose frequencies do not converge [37].

The sequences $s \in \mathcal{S}$ are called *selection rules*. A selection rule selects the j th element of x whenever $s(j) = 1$.¹¹ Informally, a collective (w.r.t \mathcal{S}) is a sequence which has invariant frequency limits with respect to all selection rules in \mathcal{S} . We call any selection rule which does not change the frequency limit of a collective *admissible*. This invariance property of collectives is often called “law of excluded gambling strategy” [107]. When thinking of a sequence of coin

tosses, a gambler is not able to gain an advantage by just considering specific selected coin tosses. The probability of seeing “heads” or “tails” remains unchanged.

Von Mises introduced the “law of excluded gambling strategy” with the goal to define randomness of a collective [107, p. 8]. A collective is called *random* with respect to \mathcal{S} . Consequently, Von Mises integrated randomness and probability into one theory. In fact, admissibility of selection rules is equivalent to statistical independence in the sense of Von Mises. But it is defined with respect to collectives instead of selection rules.

Definition 3 (Von Mises’ Definition of Statistical Independence of Collectives [107, p. 30, Def. 2]). A collective x with respect to \mathcal{S}_x is called statistically independent to the collective y with respect to \mathcal{S}_y iff the following limits exist and

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{|\{i \in \mathbb{N}: x(i) = 1 \text{ and } y(i) = 1, 1 \leq i \leq n\}|}{|\{j \in \mathbb{N}: y(j) = 1, 1 \leq j \leq n\}|} \\ = \lim_{n \rightarrow \infty} \frac{|\{i \in \mathbb{N}: x(i) = 1, 1 \leq i \leq n\}|}{n} = p_1, \end{aligned}$$

where $\frac{0}{0} := 0$. When two collectives are independent of each other we write

$$x \perp y.$$

In comparison to admissibility, the collective y adopts the role of a selection rule. It *is* in fact an admissible selection rule with the difference that a potentially finite number of elements in x are selected (cf. [41, p. 120f]). Conversely, Von Mises’ randomness is statistical independence with respect to sequences with infinitely many ones and potentially no frequency limit. (For a general comparison between Kolmogorov’s and Von Mises’ theory of probability see Appendix A and Table 2.)

4.5 Kolmogorov’s Independence versus Von Mises’ Independence

What is the relationship between Kolmogorov’s and Von Mises’ definition of statistical independence? On a conceptual level, the critique posed earlier, which questioned the meaning of statistical independence between events following Kolmogorov, gets resolved.

Von Mises adopted a strong frequential perspective on probabilities which clarifies the mapping from real world to mathematical definition. He idealized repetitive observations by infinite sequences and defined probabilities as limiting frequencies.¹² Von Mises’ independence states that there is no difference in counting the frequency of occurrences of an event in the entirety of the sequences or in a

⁹The mathematically inclined reader might notice the requirement for an order structure to obtain a definition of limit here. We use x as sequences with the standard order structure on the natural numbers. However, x can be generalized to be a net on more arbitrary base sets [51].

¹⁰Von Mises called p_0 chance as long as x is a sequence. When x is a collective, he called it probability.

¹¹We sweep under the carpet a substantial difference in Von Mises’ definition of selection rules and our definition. Von Mises allowed the selection rules to “see” the first n elements of the collective when deciding whether to choose the $n + 1$ th element [107, p. 9]. Our definition is more restrictive. We require the selected position to be determined before “seeing” the *entire* sequence. We focus on the *ex ante* nature of Von Mises randomness but neglect his recursive formalism.

¹²Without the idealization, again the mathematical description would miss a link to a worldly phenomenon. The idealization in terms of infinite sequences is substantial. In fact, the legitimacy of this idealization is the subject of another debate [45]. Nevertheless, the idealization taken in Von Mises’ framework is explicitly and transparently stated. Kolmogorov’s axioms do not possess such a statement.

subselected sequence. His independence forbids any statistical interference between processes described as sequences. No statistical patterns can be derived from one sequence by leveraging the other. Von Mises’ definition formalizes the concept of statistical independence between processes of re-occurring events.

In contrast to Kolmogorov, Von Mises’ definition does not evoke the conceptual obscurity. His focus on idealized sequences of repetitive events restricts his definition of statistical independence to specific applications with the gain of clarity in the goal of the mathematical description. Von Mises’ definition of independence makes statistical independence more concrete than Kolmogorov’s definition does.

On a more formal level, Kolmogorov defined statistical independence via the factorization of measure (cf. Definition 5), whereas Von Mises defined statistical independence via conditionalization of measures. The invariance of the frequency limit of a collective with regard to the subselection via another collective can be interpreted as the invariance of a probability of an event with regard to the conditioning on another event, i.e., “selecting with respect to” is “conditioning on” (cf. Theorem 1 and Theorem 2).

Mathematically, it turns out that Kolmogorov’s definition and Von Mises’ definition are both special cases (modulo the measure zero problem in conditioning) of a more general form of measure-theoretic statistical independence. A selection rule with converging frequency limit is admissible (respectively, statistically independent), to a collective if and only if the two are statistically independent of each other in the sense of Kolmogorov, when generalized to finitely additive probability spaces (see Appendix A for a formal statement of this claim). Thus, we can replace the known definition of statistical independence by Kolmogorov with the definition by Von Mises. Thereby, we give a specific meaning to statistical independence.

We have been motivated to dissect the notion of statistical independence for its central role in fair machine learning. Von Mises’ definition drew us closer to a more transparent mathematical formalization of statistical independence for fairness notions in machine learning. However, our discussion of Von Mises’ theory skipped over a substantial part of his work so far. Von Mises included a definition of randomness in his theory of probability. Much in contrast to Kolmogorov: There is no definition of “randomness” in Kolmogorov’s *Grundbegriffe der Wahrscheinlichkeitsrechnung* [55] (translated as [56]). Even more interestingly, Von Mises’ definition of randomness is stated in terms of statistical independence. The reader might notice that in Section 2 we already stumbled upon a heavily used notion of randomness in machine learning, which is expressed as statistical independence (i.i.d.). How do i.i.d. and Von Mises’ randomness relate to each other? How does the close connection between statistical independence and randomness complement our picture of the three fairness criteria from machine learning?

5. RANDOMNESS AS STATISTICAL INDEPENDENCE

The nature and definition of randomness seems as “random” as the term itself [35, 73, 68, 103, 8]. Usually, a very broad distinction between two approaches to randomness is made: process randomness versus outcome randomness [35]. In this work, we focus on outcome randomness and more specifically the role of randomness in statistics and machine learning.

Randomness is a modeling assumption in statistics (cf. Section 2). Upon looking into statistics and machine learning textbooks one often finds the assumption of independent and identically distributed (i.i.d.) data points as the expression of randomness [18, p. 207], [30, p. 4].

We adopt Von Mises’ differing account of randomness. The expression of randomness *relative* to the problem at hand, particularly in settings with data models such as statistics, turns out to be substantial.

5.1 Orthogonal Perspectives on Randomness as Independence in Machine Learning and Statistics

Von Mises defined a random sequence as a sequence which is statistically independent to a (pre-specified) set of selection rules respectively other sequences. In contrast, an i.i.d.-sequence consists of elements each statistically independent to all others.

Both definitions are stated in terms of statistical independence. But, the relationship of independence and randomness in terms of i.i.d. and in Von Mises’ theory differ substantially. Von Mises’ randomness is stated relative with respect to a set of selection rules. Furthermore, it is stated between sequences, respectively collectives. Whereas, in an i.i.d. sequence randomness is expressed between random variables. The randomness definitions are in an abstract sense “orthogonal.” We consider a concrete example for better understanding.

Horizontal Randomness. Let $\Omega = \mathbb{N}$ be a penguin colony. Let s, f be two attributes of a penguin, namely sex and whether a penguin has the penguin flu or not. Mathematically: $s: \Omega \rightarrow \{0, 1\}$, $f: \Omega \rightarrow \{0, 1\}$. So, penguins are individuals $n \in \Omega$ which we do not know individually, but we know some attributes of them. Suppose we are given a sequence $f(1), f(2), f(3), \dots$ of flu values with existing frequency limit. This allows us to state randomness of f with respect to the corresponding sequence of sex values s , containing infinitely many ones and having a frequency limit, by: the sequence of sex values s is admissible on f . Respectively, s and f are statistically independent of each other. In the context of colony Ω a penguin having flu is random with respect to the sex of the penguin.

Table 1. Typification of horizontal and vertical randomness (“RV” = “random variable”).

	Horizontal Randomness	Vertical Randomness
Data points are modelled as:	Evaluations of RVs	RVs
Mathematical definition of randomness of:	Sequences	Sequences of RVs
Explicit relativization:	Yes	No

Vertical Randomness. This is different to the i.i.d.-setting in which each penguin $i \in \mathbb{N}$ obtains its own random variable $F_i: \Omega \rightarrow \{0, 1\}$ on some probability space (Ω, \mathcal{F}, P) . Here, F_i encodes whether penguin i has the penguin flu or not. The sequence F_1, F_2, F_3, \dots somehow represents the colony. The included random variables share their distribution and are statistically independent to each other. The attribute flu is not random with respect to the attribute sex here, but the penguins are random with respect to each other. The random variables are (often implicitly) defined on a standard probability space on Ω . The set Ω here does *not* model the colony. It shrivels to an abstract source of randomness and probability.

The choice of perspective, horizontal or vertical, on randomness expressed as statistical independence is a question of the data model. The two types of randomness definitions are distinct in a number of ways. For a summary see Table 1. Most importantly, horizontal randomness is inherently expressed *with respect to* some mathematical object. Vertical randomness lacks this explicit relativization. This typification of horizontal and vertical, mathematical definitions of randomness is actually more broadly applicable.

To the set of vertical randomness notions one can add: exchangeability [27], α -mixing, β -mixing [94] and possibly many more. The set of horizontal randomness notions is spanned up by an entire branch of computer science and mathematics: algorithmic randomness.

Algorithmic randomness poses the question whether a sequence is random or not. This question arose in [105] within the attempt to axiomatize probability theory [10, p. 3]. In algorithmic randomness further definitions of random sequences have been proposed. For the sake of simplicity the considered sequences consist only of zeros and ones.

Four intuitions for random sequences crystallized [77, p. 280ff]: typicality, incompressibility, unpredictability and independence (see Appendix C). For our purposes, the key point to note is that a random sequence is typical, incompressible, unpredictable or independent with respect to “something” (they are all relativised in some way). Each of these intuitions has been expressed in various mathematical terms. In particular, formalizations of the same intuitions are not necessarily equivalent, and formalizations of different intuitions sometimes coincide or are logically related (see Appendix D).¹³ We mainly stick to the intuition of indepen-

dence in this paper. A random sequence is independent of “some” other sequences [105, 23].

5.2 Relative Randomness Instead of Absolute, Universal Randomness

The definition of randomness for sequences is inherently relative. Even though, the notion is relative with respect to “something,” most of the effort has been spent on finding *the* set of statistically independent sequences defining randomness [23, 68].¹⁴

Naively, one could attempt to define a random sequence as: a sequence is random if and only if it is independent with respect to *all* sequences. However, this approach is doomed to fail. There is no sequence fulfilling this condition except for trivial ones such as endless repetitions of zeros or ones (see Kamke’s critique of von Mises’ notion of randomness [100]).

So instead, research focused on computability expressed in various ways (because it was felt by those investigating these matters that computability was somehow given, or more primitive, and thus a natural way to resolve the relativity of the notion of randomness). Intuitively, randomness is considered the antithesis of computability [77, p. 288]: something which is computable is not random. Something which is random is not computable. If we then informally update the definition above we obtain: a sequence is random if and only if it is independent with respect to all computable sequences [23].¹⁵ Analogous to *the* definition of computability [77, p. 165], this is taken as an argument for the existence of *the* definition of randomness [77, p. 287].

In our work, we argue towards a relativized conception of randomness in line with work by [77], [50] and [107].¹⁶ A *relative* definition of randomness is a definition of randomness which is *relative* with respect to the problem under consideration.¹⁷ In contrast, an *absolute and universal* definition

¹⁴The analogous observation holds for all four intuitions (see Appendix C and [68]).

¹⁵To be precise, a random sequence following Church [23] is independent to all *partially computable* selection rules following Von Mises (see Footnote 11).

¹⁶Humphreys [50] presented randomness as relativized to a probabilistic hypothesis or reference class. Porter [77, p. 169] even postulated the “No-Thesis”-Thesis: Any notion of randomness neither defines a well-defined collection of random sequences nor captures all mathematical conceptions of randomness; confer the logic of “essentially contested” concepts [39], which, presumably, are unavoidably contested for the same reason.

¹⁷In machine learning, a problem under consideration is, for instance, animal classification via neural networks.

¹³For an overview of algorithmic randomness see [99, 68, 103, 31].

of randomness would preserve its validity in all problems. It presupposes the existence of *the* randomness.

Relative randomness with respect to the problem which we want to describe aligns to Von Mises’ theory of probability and randomness. Von Mises emphasized the *ex ante* choice of randomness [106, p. 89] *relative* to the problem at hand [107, p. 12]. First, one formalizes randomness with respect to the underlying problem, then one can consider a sequence to be random or not. Otherwise, if we are given a sequence, it is easy to construct a set of selection rules, such that the sequence is random with respect to this set.¹⁸ This, however, undermines the concept of randomness, which should capture the pre-existing typicality, incompressibility, unpredictability or independence of a sequence (cf. [108, p. 321]). Von Mises’ randomness intrinsically possesses a modeling character, similar to our needs in machine learning and statistics.

Given its role as modeling assumption in statistics, randomness lacks substantial justification to be expressed in any *absolute and universal* manner in this context. Neither are there reasons why computability¹⁹ is the only mathematical, expressive way to encode one of the four intuitions of randomness. The i.i.d. assumption, an *absolute and universal* definition of randomness, does not fit this purpose. To appropriately model data we require adjustable notions of randomness. Otherwise, we restrict our modeling choice without reason or gain.²⁰

Equipped with the interpretation of statistical independence as randomness we now return to our motivation for investigating statistical independence. ML-friendly fairness criteria are built upon statistical independence. In contrast to Kolmogorov, Von Mises’ statistical independence transparently refers to a concept of independence in the real world. To clarify the meaning of fairness expressed as statistical independence, we directly apply Von Mises’ independence to the fairness criteria listed in Section 3 in the following.

6. VON MISES’ FAIRNESS

With Von Mises’ definition of statistical independence we have a notion at our disposal which is conceptually focused on a more “scientific” perspective (i.e., making claims about the world) of statistical concepts. Since it is mathematically related to Kolmogorov’s standard account of statistical independence, Kolmogorov’s definition can, at many places, be easily replaced by Von Mises’ definition.

¹⁸This idea is transferable to other intuitions of randomness, for instance [102].

¹⁹Computability is often taken (at least by computer scientists) as a purely mathematical notion, detached from the world. An alternate view, close in spirit to Von Mises, is that computation is part of physics, and thus needs to be viewed in a *scientific*, and not merely *mathematical* manner [29].

²⁰This is not entirely true, as specific computability notions of randomness and the i.i.d.-assumption deliver convergence results for random sequences, which can be used to guarantee low estimation error in the long run.

Let us denote the three presented fairness criteria in a Von Mises’ way (cf. Section 4.5).

Definition 4 (Fairness as Statistical Independence). A collective $x: \mathbb{N} \rightarrow \{0, 1\}$ (with respect to a family of selection rules \mathcal{S}) is *fair with respect to a set of sensitive groups* $\mathcal{G} = \{s^j: \mathbb{N} \rightarrow \{0, 1\} | j \in J\}$ if

$$x \perp s^j \quad \forall j \in J$$

The 0-1-sequences s^j determine for each individual i whether it is part of the group or not (according to whether $s^j(i) = 1$ or $s^j(i) = 0$). We call these groups “sensitive,” as these are the groups which are of moral and ethical concern. In philosophical literature these groups are often called “socially salient groups” [2, 64].²¹ We see that the connection between Von Misesan independence and fairness arises from the observation that the set of sensitive groups \mathcal{G} is a family of selection rules, so that if $\mathcal{G} \subseteq \mathcal{S}$, then indeed the collective x will be fair for \mathcal{G} .

Following Von Mises’ interpretation of independence, the given definition reads as follows: we assume we are in the idealized setting of infinitely many individuals with values x_i , e.g., binary predictions. The predictions are fair if and only if there is no difference in counting the frequency of 1-predictions in the entirety or in the sensitive group. (For an illustration see Appendix 6.) A proper conceptualization of fairness requires such immediate semantics, but a purely mathematical theory of probability cannot offer these (see Section 4.2).

Each of the three fairness criteria is captured in Definition 4; the choice of fairness criterion manifests in the collective under consideration:

Independence The collective $x: \mathbb{N} \rightarrow \{0, 1\}$ consists of predictions; i.e., $\{0, 1\}$ is the set of predictions.

Separation The collective $x: \mathbb{N} \rightarrow \{0, 1\}$ is obtained via the subselection of predictions based on the sequence of true labels corresponding to the predictions.²²

Sufficiency The true labels are subselected by predictions.

To enable intuitive access to the Von Mises’ notions of fairness we provide a toy example in Appendix B. The three fairness criteria Independence, Separation and Sufficiency encompass a large part of fair machine learning [5, p. 45]. Von Mises’ statistical independence gives a consistent interpretation to all of them. In fact, Von Mises’ independence opens the door to further investigations. To this end, we recapitulate the strong linkage between statistical independence and randomness in Von Mises’ theory.

7. THE ETHICAL IMPLICATIONS OF MODELING ASSUMPTIONS

Machine learning methods try to model data in complex ways. Derived statements, such as predictions, then potentially get applied in society. In these cases one is obliged to

²¹Intersections of sensitive groups are not necessarily independent.

²²“Selecting with respect to” is “conditioning on.”

ask which ought-state the machine learning model should reflect [75, 83]. To enable a justified choice, statistical concepts in machine learning require relations to the real world. Furthermore, modeling even requires understanding of the entanglement of societal and statistical concepts.

We proposed one specific *meaningful* definition of statistical independence which can be directly applied to the three observational fairness criteria from fair machine learning. In addition, this Von Mises’ independence is key to a relativized notion of randomness. Pulling these threads together, we are now able to establish the following link: *Randomness is fairness. Fairness is randomness.*

7.1 Randomness is Fairness. Fairness is Randomness

The concepts fairness and randomness frequently appear jointly: [14] argues that a random allocation of goods is fair under certain conditions. Literature on sortition argues for just representation of society by random selection of people [74, 95].²³ Bennett [7, p. 633] even states that randomness encompasses fairness.

With Von Mises’ axiom 2 and Definition 4 we can now tighten the conceptual relationship of fairness and randomness. The proposition directly follows from the definition of randomness respectively fairness in the sense of Von Mises.

Proposition 1 (Randomness is fairness. Fairness is randomness.). *Let x be a collective with respect to \emptyset (the empty set). It is fair with respect to a set of sensitive groups (0-1-sequences) $\{s^j\}_{j \in J}$, if and only if it is random with respect to $\{s^j\}_{j \in J}$.*

The given proposition establishes a helpful link. It gives insights into both of the concepts. In particular, it substantiates the relativized conception of randomness in machine learning as it presents randomness as an ethical choice.

7.1.1 Randomness as Ethical Choice

Randomness in machine learning is a modeling assumption (Section 2). Fairness is an ethical choice.²⁴ In light of Proposition 1 randomness gets an ethical choice and fairness a modeling assumption. We now further detail this perspective.

We assume that we are given a fixed set of selection rules, which defines “the” randomness. As far-fetched as this may sound, if we, for example, accept the so called Martin-Löf randomness as *absolute and universal* definition, then we exactly do this and fix the set of selection rules to the partial computable ones (see Appendix D.1). A sequence which is random with respect to this specified set of selection rules is fair with respect to the groups defined by the selection rules. Rephrased in terms of Martin-Löf randomness: a Martin-Löf

random sequence is fair with respect to all partial computable groups. Only non-partial-computable groups (respectively sequences) can be discriminated against in this setting. If we interpret statistical independence as fairness (Section 3), then fairness is as *absolute and universal* as randomness here. Where did the “essentially contested” nature of fairness [39] leave the picture?

The set of admissible selection rules specifies the choice of sensitive groups, which indeed is a fraught and contestable choice [67, Section H.3]. Thus each selection rule gets ethically loaded. Furthermore, the choice of collective, which we consider as random, fixes the fairness criterion. In summary, the determination of randomness is analogous to the determination of fairness.

However one defines randomness, it is an ethical choice. For symmetry reasons one can equivalently state in machine learning: fairness is a modeling assumption. The randomness assumption has an ethical, moral and potentially legal implication. We need non-mathematical, contextual arguments to each problem at hand which justify the adjustable and explicit randomness assumptions.

Given that randomness is an ethical choice, an absolute, universal conception of randomness counteracts any ethical debate in machine learning. Discussions about sexism, racism and other kinds of discrimination and injustice persist over time without ever arrogating the discovery of “the” fairness [39]. But if “the” randomness as statistical independence would exist, then “the” fairness as statistical independence would be an accessible notion. For illustration, we reconsider Martin-Löf randomness. A Martin-Löf random sequence is independent, respectively fair, to the set of all partial computable selection rules. But, it is completely unclear what the ethical meaning of partial computable groups is. And, it remains unsolved whether the groups given by gender are partial computable, when we desire to be fair with respect to them. We conceive Proposition 1 as further counterargument to an absolute, universal definition of randomness. Randomness is, like fairness, better interpreted as a *relative* notion.

Further concluding, the equivalence of randomness and fairness highlights the deficiency of fairness notions in machine learning. The equivalence only holds due to the very reductionist perspective on fairness in fair machine learning. Despite their regular co-occurrence [14, 74], [7, p. 633], fairness and randomness are more multi-faceted and non-overlapping concepts as illustrated in Proposition 1.

7.1.2 Fairy Tales of Fairness: “Perfectly Fair” Data

With the relationship between fairness and randomness in mind, we now turn towards random data as primitive. Discussions in fair machine learning sometimes seemingly presume the existence of “perfectly fair” data (e.g., as highlighted in [83, p. 134]), as if fair machine learning merely tackles the cases where “perfectly fair” data is not available.

²³Representativity here can be interpreted as typicality, thus one of the four intuitions for randomness.

²⁴More specifically, the choice of operationalized fairness, one of the fairness criteria, and the choice of groups.

We interpret “perfectly fair” data as a collective with respect to all possible selection rules. The data does not depend on any (sensitive) group at all. In other words, “perfectly fair” data is “totally random” data. As we saw in Section 5.2 this is self-contradicting except of the trivial constant case. “Perfectly fair” data does not exist or is statistically useless.

7.2 Demanding Fairness Is Randomization: Fair Predictors Are Randomizers

In practice, it is often unreasonable to *assume* random or fair data as in Proposition 1. Instead one *demand*s for fairness respectively randomness of predictions. In these settings, fair machine learning techniques are deployed to exhibit ex post fulfillment of fairness criteria.

We assume for the following discussion that the collective x consists of predictions, as in the fairness criteria Independence or Separation. Fair machine learning techniques enforce statistical independence of predictions and sensitive attributes. Rephrased, fair machine learning techniques actually introduce randomness post-hoc into the predictions. Thus, fair machine learning techniques can potentially be interpreted as randomization techniques.

7.2.1 Fairness-Accuracy Trade-Off — Another Perspective

We noticed that fair predictions are random predictions with respect to the sensitive attribute. In contrast, accurate predictions exploit all dependencies between given attributes and predictive goal, including the sensitive attributes. Thus, in fair machine learning morally wrongful discriminative potential of sensitive attributes is thrown away by purpose. On these grounds, it is not surprising that an increase in fairness respectively randomness (usually) goes hand in hand with a decrease in accuracy [110]. Randomization of predictions leads to the so called fairness-accuracy trade-off.

Concluding, via Von Mises’ axiomatization we established: *Randomness is fairness. Fairness is randomness.* Exploiting this new perspective, we unlock another perspective on fair predictors as randomizers, demonstrate the nonexistence of “perfectly fair” data and treat randomness as an ethical choice, which can be neither universal nor total. In particular, the “essentially contested” nature of fairness is tied to the “essentially *relative*” nature of randomness.

8. CONCLUSION

Fair machine learning attained an increasing interest in the last years. However, its conceptual maturity lags behind. In particular, the interplay between data, its mathematical representation and their relation to fairness is encompassed by a veil of nescience. In this paper, we contribute towards a better understanding of randomness and fairness in machine learning.

We started from the most commonly used definition of statistical independence and questioned its representation

due to a lack of semantics. Generally, we observe that in machine learning, as in statistics, probability and its related concepts should be interpreted as modeling assumptions about the world (of data). Von Mises aimed for exactly this “scientific” perspective on probability theory. We lean on his statistical independence, which clarifies the relation to the real world, and his definition of randomness, which is *relative* and orthogonal to the i.i.d. assumption, but similarly expressed as statistical independence. Then by the three fairness criteria in machine learning we obtain a further interpretation of independence, which we finally exploit to argue for a *relative* conception of randomness, randomness as an ethical choice in machine learning and fair predictors as randomizers.

8.1 Future Work: Approximate Randomness and Fairness, Randomness as Fairness via Calibration

Despite future conclusions in-between the topics fairness and randomness in other research subjects as machine learning, we claim that a significant dimension is missing in the present discussion. Practitioners usually deal with approximate versions of randomness, statistical independence or fairness. Yet, approximation spans another dimension of choice beset with pitfalls [79, 63]. Several questions ranging from the choice of approximation to the interference of concepts arise. Future work should detail the implications of this choice.

Second, we conjecture that “*Randomness is Fairness. Fairness is randomness.*” can be substantiated via the intuition of unpredictability. Starting from [90] definition of unpredictability randomness, which is closely related to the calibration idea presented in [25], we can bridge to fairness as calibration as given in [22]. A recent work by Cynthia Dwork and collaborators in fact show a formal link between pseudo-randomness and fairness as calibration [34]. This work, however, still misses a more thorough discussion of the concepts of individual versus group fairness in machine learning [12]. As a subproblem, which is contained therein, the categorization into (sensitive) groups in fair machine learning deserves its own work.

Third, regarding a more thorough definition of statistical independence within the fairness criteria, we are convinced that a subjectivist interpretation of probability might reveal yet another perspective on the problem. We assume that the interplay between different interpretations of probability and ethical concepts such as fairness still leaves room for many important investigations.

Fourth, there are certainly more frameworks to give an interpretation and concretization to current notions in (fair) machine learning (cf. [41, 36]).

Last but not least, we already referred to sortition literature and random allocation. The somewhat different relation between fairness and randomness in this literature leads us

to speculate that further fruitful discussions between the two concepts may develop.

In the jungle of statistical concepts such as probability, uncertainty, randomness, independence etc. further relations to social and ethical concepts wait to be brought to light. And machine learning research should care:

The arguments that justify inference from a sample to a population should explicitly refer to the variety of non-mathematical considerations involved. [6, p. 11]

APPENDIX A. GENERALIZED VON MISESEAN PROBABILITY THEORY

In this appendix we outline a theory of probability subsuming that of Kolmogorov and Von Mises.

A.1 Kolmogorov's Notion of Independence

Kolmogorov axiomatized probability theory in his book [56] in a measure theoretical way. He defined a probability space (Ω, \mathcal{F}, P) as a measure space with base set Ω , σ -algebra \mathcal{F} and normalized measure P . Events are elements of \mathcal{F} , i.e., subsets of Ω , which obtain a probability via P . Statistical independence is defined as a specific assignment of probability to an intersection event.

Definition 5 (Kolmogorov's Definition of Statistical Independence of Events [56, p. 9, Def. 1]). Let (Ω, \mathcal{F}, P) be a probability space. Two events $A, B \in \mathcal{F}$ are called statistically independent iff

$$P(A \cap B) = P(A)P(B).$$

As we highlighted above, Kolmogorov's axiomatization is, despite its success, not the only mathematical theory of probability. Specifically, one can weaken the structure of the probability space and still work with concepts such as statistical independence, expectation, conditioning (e.g., [43, 21, 70]).

A.2 Finitely Additive Probability Space

We introduce a weaker measure structure, which we call finitely additive probability space. Interestingly, this weaker structure includes the axiomatization of Kolmogorov and Von Mises as special cases. We define a finitely additive probability space modified from [81, Def. 2.1.1 (7)] as

Definition 6 (Finitely Additive Probability Space). The tuple (N, \mathcal{A}, ν) is called a *finitely additive probability space* for a base set N , a set of measurable sets $\mathcal{A} \subset \mathcal{P}(N)$ containing the empty set $\emptyset \in \mathcal{A}$ and a finitely additive probability measure $\nu : \mathcal{A} \rightarrow [0, 1]$ satisfying the following conditions:

- (1) $\nu(\emptyset) = 0$,
- (2) if $A_1, A_2, A_1 \cup A_2 \in \mathcal{A}$ and $A_1 \cap A_2 = \emptyset$ then $\nu(A_1 \cup A_2) = \nu(A_1) + \nu(A_2)$.

Observe that this definition does not impose any structural restrictions on the set of subsets \mathcal{A} .

Kolmogorov's probability space is certainly a specific finitely additive probability space in our sense, as every set σ -algebra contains the empty set and every countably additive probability is finitely additive.

Analogously, Von Mises implicitly uses a finitely additive probability space. This space is given by $(\mathbb{N}, \mathcal{A}_{\text{VM}}, \nu_{\text{VM}})$, where \mathbb{N} are the natural numbers and $\mathcal{A}_{\text{VM}}, \nu_{\text{VM}}$ are defined in the following.

First, we consider the finitely additive base measure

$$\nu_{\text{VM}}(A) := \lim_{n \rightarrow \infty} \frac{|A \cap \mathbb{N}_n|}{n},$$

where $A \subset \mathbb{N}$ and $\mathbb{N}_n = \{1, \dots, n\}$. ν_{VM} is called the “natural density” in the number theory literature [71, p. 256]. From this definition is not clear whether the given limit exists. Thus, we define the set of measurable sets by

$$\mathcal{A}_{\text{VM}} := \{A : \nu_{\text{VM}}(A) \text{ exists}\},$$

which is called the “density logic” in [78]. It is a pre-Dynkin-system [88].

We generalize Kolmogorov's definition of statistical independence of events to finitely additive probability spaces.

Definition 7 (Statistical Independence of Events on a Finitely Additive Probability Space). Let (N, \mathcal{A}, ν) be a finitely additive probability space. Two measurable sets $A, B \in \mathcal{A}$ are independent iff

1. $A \cap B \in \mathcal{A}$
2. $\nu(A \cap B) = \nu(A)\nu(B)$

Observe that the first condition is naturally fulfilled in Kolmogorov's σ -algebra. In the case of Von Mises' density logic, this constraint is strict. A pre-Dynkin-system is not closed under arbitrary intersections.

A.3 Von Mises' Admissibility and Kolmogorov's Independence Are Analogues

We now reconsider the definition of collectives and selection rules. Both are 0, 1-sequences on the natural numbers, with the restriction that selection rules contain infinitely many 1's and their frequency limit does not have to exist. Both sequences can be interpreted as indicator functions on the natural numbers. So for $x : \mathbb{N} \rightarrow \{0, 1\}$ and $s : \mathbb{N} \rightarrow \{0, 1\}$ we write $X = x^{-1}(1)$ and $S = s^{-1}(1)$ for the corresponding subsets of the natural numbers.

We want to show that Von Mises' admissibility condition is equivalent to the given definition of statistical independence on finitely additive probability spaces. Actually, the equivalence only holds for the slightly restricted case in which the selection rule itself possesses converging frequencies.

Theorem 1 (Admissibility in Von Mises setting implies Statistical Independence on Finitely Additive Probability Spaces). *Let x be a collective with respect to s . Suppose furthermore that s has a converging frequency limit. Then X and S , the indexed sets corresponding to the collective (respectively selection rule) on the finitely additive probability space $(\mathbb{N}, \mathcal{A}_{\text{vM}}, \nu_{\text{vM}})$ are statistically independent.*

Proof. We know that

$$\begin{aligned} & \lim_{n \rightarrow \infty} \frac{|\{i : x(i) = 1 \text{ and } s(i) = 1, 1 \leq i \leq n\}|}{|\{j : s(j) = 1, 1 \leq j \leq n\}|} \\ &= \lim_{n \rightarrow \infty} \frac{|\{i : x(i) = 1, 1 \leq i \leq n\}|}{n} \end{aligned}$$

which we can rewrite as

$$\lim_{n \rightarrow \infty} \frac{|X \cap S \cap \mathbb{N}_n|}{|S \cap \mathbb{N}_n|} = \lim_{n \rightarrow \infty} \frac{|X \cap \mathbb{N}_n|}{n}.$$

This gives

$$\begin{aligned} \nu_{\text{vM}}(X \cap S) &= \lim_{n \rightarrow \infty} \frac{|X \cap S \cap \mathbb{N}_n|}{n} \\ &= \lim_{n \rightarrow \infty} \frac{|X \cap S \cap \mathbb{N}_n|}{|S \cap \mathbb{N}_n|} \frac{|S \cap \mathbb{N}_n|}{n} \\ &= \lim_{n \rightarrow \infty} \frac{|X \cap S \cap \mathbb{N}_n|}{|S \cap \mathbb{N}_n|} \lim_{n \rightarrow \infty} \frac{|S \cap \mathbb{N}_n|}{n} \\ &= \lim_{n \rightarrow \infty} \frac{|X \cap \mathbb{N}_n|}{n} \lim_{n \rightarrow \infty} \frac{|S \cap \mathbb{N}_n|}{n} \\ &= \nu_{\text{vM}}(X) \nu_{\text{vM}}(S), \end{aligned}$$

by help of a standard result for the multiplication of sequence limits [62, Theorem 3.1.7]. \square

Theorem 2. *Let X and S be two statistically independent events on the finitely additive probability space $(\mathbb{N}, \mathcal{A}_{\text{vM}}, \nu_{\text{vM}})$ with $\nu_{\text{vM}}(S) > 0$, then the corresponding collective x , indicator function of X , has the admissible selection rule s , indicator function of S .*

Proof. It is given that

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{|X \cap S \cap \mathbb{N}_n|}{n} &= \nu_{\text{vM}}(X \cap S) \\ &= \nu_{\text{vM}}(X) \nu_{\text{vM}}(S) = \lim_{n \rightarrow \infty} \frac{|X \cap \mathbb{N}_n|}{n} \lim_{n \rightarrow \infty} \frac{|S \cap \mathbb{N}_n|}{n}. \end{aligned}$$

Furthermore $\nu_{\text{vM}}(S) > 0$ implies that the corresponding selection rule selects infinitely many elements and $\lim_{n \rightarrow \infty} \frac{n}{|S \cap \mathbb{N}_n|} = \frac{1}{\nu_{\text{vM}}(S)}$.

This implies

$$\begin{aligned} & \lim_{n \rightarrow \infty} \frac{|X \cap S \cap \mathbb{N}_n|}{|S \cap \mathbb{N}_n|} \\ &= \lim_{n \rightarrow \infty} \frac{|X \cap S \cap \mathbb{N}_n|}{n} \frac{n}{|S \cap \mathbb{N}_n|} \end{aligned}$$

$$\begin{aligned} &= \lim_{n \rightarrow \infty} \frac{|X \cap S \cap \mathbb{N}_n|}{n} \lim_{n \rightarrow \infty} \frac{n}{|S \cap \mathbb{N}_n|} \\ &= \lim_{n \rightarrow \infty} \frac{|X \cap \mathbb{N}_n|}{n} \lim_{n \rightarrow \infty} \frac{|S \cap \mathbb{N}_n|}{n} \lim_{n \rightarrow \infty} \frac{n}{|S \cap \mathbb{N}_n|} \\ &= \lim_{n \rightarrow \infty} \frac{|X \cap \mathbb{N}_n|}{n}. \end{aligned} \quad \square$$

We note some caveats regarding the preceding discussion.

1. We require the frequency limit for s to exist in Theorem 1. So the sequence S is not an entirely general selection rule.
2. The condition $\nu_{\text{vM}}(S) > 0$ in Theorem 2 ensures that we do not condition on measure zero events. Furthermore, it guarantees at this point, that the indicator function of S is a selection rule containing infinitely many ones. Besides that, its frequency limit exists.
3. Even though given here only for the binary case. The argumentation can be extended to continuum labeled collectives [107, II.B].
4. In our discussion, we focus on admissibility instead of statistical independence in the sense of Von Mises, since our argument finally focuses on the randomness interpretation of statistical independence. Admissibility and Von Mises' statistical independence are, despite the objects on which they are defined, equivalent.
5. Statistical independence in Von Mises' setting and Kolmogorov's setting are not equivalent. They are special cases of a more general form of statistical independence. So they are mathematically analogous under mild conditions. But we emphasize the conceptual difference. Von Mises refused to define statistical independence between mere events [107, p. 35]. Instead he demanded statistical independence to be defined between collectives underlining that independence is a concept about aggregates, the collectives, not single occurring events.

APPENDIX B. VON MISES' NOTIONS OF FAIRNESS IN PRACTICE

In practice we never observe infinite sequences, which are the building blocks of Von Mises' theory. Nevertheless, infinite sequences can be seen as idealizations of finite, observable sequences. We make the following crucial and debatable assumption here: *Given a sequence of sufficient length, the frequencies observed in this sequences are arbitrarily close to the limit of the frequency in the continuation of this sequence.*

This assumption is critical for at least three parts:

1. We assume convergence (cf. [37]).
2. In finite data we can only observe fractional frequencies, but not irrational ones (converging frequencies do fill the entire $[0,1]$).
3. The frequency limit of any infinite sequence is not governed by any frequency observed on a finite sequence.




Figure 1: Generating Distribution and Predictor of the Toy Example.




Figure 2: Groups in the Toy Example.

For further simplification, we constrain ourselves to 0,1-outcomes and 0,1-decisions. In the following, we call 1 as outcome or decision “positive”, comparable to a setting where the 1-label corresponds to “getting a loan” and 0 to “not getting a loan”.

We consider the following simple toy example: let $X = \mathbb{R}$, $Y = \{0, 1\}$. The data points of the 0 label are distributed according to a Gaussian distribution with mean -1 and standard deviation 1 . The data points of the positive labels are distributed according to a Gaussian distribution with mean 1 and standard deviation 0.5 . We consider the simple logistic probabilistic predictor $p(x) = \frac{e^{cx+a}}{1+e^{cx+a}}$ with $c = 2$ and $a = -0.2$ and $x \in X$. We consider two groups: group A consists of all data points with values between -1 and -2 , group B consists of all data points with values between smaller than -1.5 or greater than 1.5 . In this case, the input and sensitive groups are highly correlated. Figure 1 and Figure 2 illustrate the example.

Since the predictor is probabilistic we have to introduce a threshold $q \in [0, 1]$ which derives a 0,1-decision from the prediction. The following plots show how the frequency of positive outcomes respectively positive decisions depending on the notion of fairness changes for different decision-thresholds.

Independence We compute the frequency of positive decisions for each group and the entirety. Independence requires these frequencies to be equal. See Figure 3.

Separation We compute the frequency of positive decisions given that the outcomes were positive for each group and the entirety. Separation requires these frequencies to be equal. See Figure 4.

Sufficiency We compute the frequency of positive outcomes given that the decision were positive for each




Figure 3: Frequency of Positive Decisions per Group depending on the Decision-Threshold.

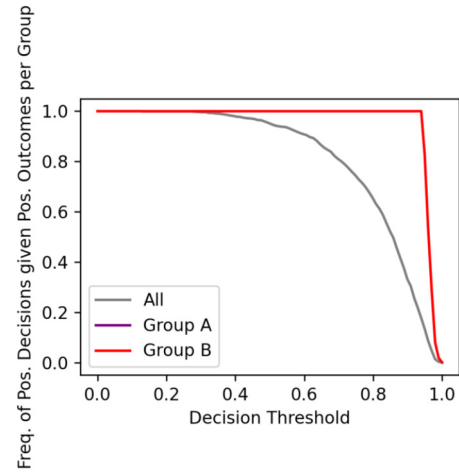


Figure 4: Frequency of Positive Decisions given Positive Outcomes per Group depending on the Decision-Threshold.




Figure 5: Frequency of Positive Outcomes given Positive Decisions per Group depending on the Decision-Threshold.

group and the entirety. Sufficiency requires these frequencies to be equal. See Figure 5.

What do we learn from this toy example:

1. Our suggested notions have an empirical counterpart. Even though, the assumption taken above is critical. Much in contrast to most literature in fair machine learning, however, we (can) specify the idealization assumption explicitly. The regular setup hides many of the concerns behind a curtain of acceptance.
2. We observe that there is nothing special about our definitions of fairness. The frequential approach is intuitive and simplifies communication of technical results about machine learning fairness in societal debates.
3. We did not allude to specific algorithms which guarantee accurate predictions under fairness constraints. All such existing algorithms can be deployed in order to provide fair predictions following our definitions. Our notions are reinterpretations, not entirely new setups.
4. There is still a further, important, debate missing in this work. As pointed out in the conclusion of the paper, further conceptual clarification has to account for approximations of fairness notions to further enable principled argumentation in fairness debates.

APPENDIX C. FOUR INTUITIONS OF RANDOMNESS

Typicality A sequence is called random if it shares “some” characteristics of any possible sequence. Martin-Löf and Schnorr formalized this idea via statistical tests [65, 87].

Incompressibility The information contained in a random sequences is (approximately) as large as the sequence itself. The sequences cannot be compressed by “some” procedure [57, 20, 86, 59].

Unpredictability In a random sequence one can know all foregoing elements without being able to predict by “some” procedure the next element [102, 68, 90, 38, 26].

Independence A random sequence is independent of “some” other sequences [105, 23].

APPENDIX D. RELATIONS BETWEEN RANDOMNESS DEFINITIONS

We briefly outline some of the known relationships between various notions of randomness.

On one hand, there are mathematical expressions capturing the same randomness intuitions meanwhile being mathematically distinct to each other.

For instance, the definition of a typical sequence following Martin-Löf [65] implies Schnorr’s definition of typicality [85, 87] but not vice versa [99, p. 143]. Cooman and De Bock’s imprecise unpredictability randomness [26] is strictly more expressive than Vovk and Shafer’s unpredictability randomness [90, Section 1.1], [26, Theorem 37].

On the other hand, there are mathematical expressions capturing differing randomness intuitions meanwhile being necessary, sufficient or even equivalent to each other. For instance, the Levin-Schnorr theorem, simultaneously proven by Levin [59] and Schnorr [86, 87], established an equivalence of typicality following Martin-Löf [65] and incompressibility following Levin [59] and Schnorr [86] based on the idea of [57] (e.g., see in [103, Theorem 5.3] and references therein). Cooman and de Bock expressed Schnorr [85, 87] and Martin-Löf randomness [65] in terms of an unpredictability approach [26]. Muchnik showed in [68] that all incompressible sequences, again following Levin and Schnorr’s prefix-complexity approach [86, 59] are unpredictable in his sense.

In Ville’s thesis [102] he attempted to generalize the idea of “excluding a gambling strategy” (see Section 4.4) via a game-theoretic approach. He showed [102, p. 76] that for any set of admissible selection rules \mathcal{S} he could construct a gambling strategy, more exactly a capital process associated to a gambling strategy, which captures the same randomness definition. The opposite direction is, however, impossible [102, p. 39]. But Ambos-Spies et al. [3] introduced a weaker form of [102]’s randomness as unpredictability which is equivalent to Church’s randomness [23] as independence [31, Section 12.3]. Finally, van Lambalgen observed an abstract interpretation of “random with respect to something” as “independent to” in [101].

D.1 A Prototypical, Absolute and Universal Notion of Randomness

One often referred *absolute and universal* definition of randomness is Martin-Löf’s typicality approach [15]. Martin-Löf’s randomness as typicality has been equivalently formalized in terms of incompressibility and unpredictability (see Section D). Furthermore, a sequence, which is Martin-Löf random [65], is statistically independent to all partial computable selection rules [23], [96, Theorem 11]. Contrarily, not every collective with respect to all partial computable selection rules is a Martin-Löf random sequence [13, p. 193]. So Martin-Löf’s definition is linked to all four intuitions.

APPENDIX E. KOLMOGOROV’S VERSUS VON MISES’ PROBABILITY THEORIES IN A TABLE

See Table 2.

APPENDIX F. PENGUIN COLONY EXAMPLE FOR FAIR COLLECTIVE

See Figure 6.

Table 2. Summary of main difference between Kolmogorov's and Von Mises' probability theories.

	Kolmogorov	Von Mises
fundamental structure	probability space (Ω, \mathcal{F}, P)	collectives $(x_i)_{i \in \mathbb{N}}$, implicitly $(\mathbb{N}, \mathcal{A}_{\text{VM}}, \nu_{\text{VM}})$
base set	(almost arbitrary) set Ω	\mathbb{N}
set of events	σ -algebra \mathcal{F}	Dynkin-system \mathcal{A}_{VM}
probability	finite positive measure	limit of frequency sequence
probability measure	countably additive	finitely additive ν_{VM}
randomness	no explicit mathematical definition	explicit mathematical definition
statistical independence	factorization of joint distribution	frequency limit doesn't change under subselection
data model	each data point a random variable X_i on (Ω, \mathcal{F}, P)	each data point an element in a collective $(x_i)_{i \in \mathbb{N}}$




Figure 6: Example for a subselection and fair collective. n_b denotes the number of black penguins among the first n -penguins. Blackness of penguins is distributed fairly with respect to sex if $p_b = p_{b|F}$.

ACKNOWLEDGEMENTS

Many thanks to Christian Fröhlich, Eric Raidl, Sebastian Zenzulka, Thomas Grote and Benedikt Höltingen for helpful discussions and feedback. Furthermore, the authors thank all participants of the Philosophy of Science meets Machine Learning Conference 2022 in Tübingen, for all helpful comments and debates.

The authors appreciate and thank the International Max Planck Research School for Intelligent System (IMPRS-IS) for supporting Rabanus Derr.

FUNDING

This work was funded in part by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy — EXC number 2064/1 — Project number 390727645; it was also supported by the German Federal Ministry of Education and Research (BMBF): Tübingen AI Center.

REFERENCES

- [1] ABRAHAM, K. S. Efficiency and fairness in insurance risk classification. *Virginia Law Review* **71**(3) 403–451 (1985).
- [2] ALTMAN, A. Discrimination. In (E. N. Zalta, ed.) *The Stanford Encyclopedia of Philosophy* (Summer 2020), 2020th ed. Stanford University (2020). <https://plato.stanford.edu/entries/discrimination/>
- [3] AMBOS-SPIES, K., MAYORDOMO, E., WANG, Y. and ZHENG, X. Resource-bounded balanced genericity, stochasticity and weak randomness. In *Annual Symposium on Theoretical Aspects of Computer Science* 61–74. Springer (1996). [MR1462086](https://doi.org/10.1007/BF01186208)
- [4] BANDYOPADHYAY, P. S. and FORSTER, M. R. Philosophy of statistics: An introduction. In (P. S. Bandyopadhyay and M. R. Forster, eds.) *Philosophy of Statistics* **7**. Elsevier (2011). <https://doi.org/10.1016/B978-0-444-51862-0.50001-0>. [MR3295937](https://doi.org/10.1016/B978-0-444-51862-0.50001-0)
- [5] SOLON BAROCAS, HARDT, M. and NARAYANAN, A. *Fairness and Machine Learning: Limitations and Opportunities* (2019). <http://www.fairmlbook.org>
- [6] BATTERSBY, M. The rhetoric of numbers: Statistical inference as argumentation. In *OSSA Conference Archive* (2003).
- [7] BENNETT, D. Defining randomness. In (P. S. Bandyopadhyay and M. R. Forster, eds.) *Philosophy of Statistics* **7**. Elsevier (2011). <https://doi.org/10.1016/B978-0-444-51862-0.50001-0>. [MR3295937](https://doi.org/10.1016/B978-0-444-51862-0.50001-0)
- [8] BERKOVITZ, J., FRIGG, R. and KRONZ, F. The ergodic hierarchy, randomness and hamiltonian chaos. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* **37**(4) 661–691 (2006). <https://doi.org/10.1016/j.shpsb.2006.02.003>. [MR2344115](https://doi.org/10.1016/j.shpsb.2006.02.003)
- [9] BIDDLE, D. *Adverse Impact and Test Validation: A Practitioner's Guide to Valid and Defensible Employment Testing*. Gower Publishing, Ltd. (2006).
- [10] BIENVENU, L., SHAFER, G. and SHEN, A. On the history of martingales in the study of randomness. *Electronic Journal for History of Probability and Statistics* **5**(1) 1–40 (2009). [MR2520666](https://doi.org/10.1016/j.shpsb.2006.02.003)
- [11] BINNS, R. Fairness in machine learning: Lessons from political philosophy. In *Conference on Fairness, Accountability and Transparency* 149–159. PMLR (2018).
- [12] BINNS, R. On the apparent conflict between individual and group fairness. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* 514–524 (2020).
- [13] BLUM, N. *Einführung in Formale Sprachen, Berechenbarkeit, Informations-und Lerntheorie*. Oldenbourg Wissenschaftsverlag (2009).
- [14] BROOME, J. Selecting people randomly. *Ethics* **95**(1) 38–55 (1984).
- [15] BUSS, S. and MINNES, M. Probabilistic algorithmic randomness. *The Journal of Symbolic Logic* **78**(2) 579–601 (2013). [MR3145197](https://doi.org/10.1016/j.shpsb.2006.02.003)
- [16] TOON, C. and VERWER, S. Three naive Bayes approaches for discrimination-free classification. *Data Mining and Knowledge Discovery* **21**(2) 277–292 (2010). <https://doi.org/10.1007/s10618-010-0190-x>. [MR2720507](https://doi.org/10.1007/s10618-010-0190-x)
- [17] CAMPBELL, D. T. Common fate, similarity, and other indices of the status of aggregates of persons as social entities. *Behavioral Science* **3**(1) 14–25 (1958).
- [18] CASELLA, G. and BERGER, R. L. *Statistical Inference* 2nd ed. Duxbury Advanced Series (2002). [MR1051420](https://doi.org/10.1016/j.shpsb.2006.02.003)
- [19] CASTANO, E., YZERBYT, V., PALADINO, M.-P. and SACCHI, S. I belong, therefore, I exist: Ingroup identification, ingroup entitativity, and ingroup bias. *Personality and Social Psychology Bulletin* **28**(2) 135–143 (2002).
- [20] CHAITIN, G. J. On the length of programs for computing finite binary sequences. *Journal of the ACM* **13**(4) 547–569 (1966). <https://doi.org/10.1145/321356.321363>. [MR0210520](https://doi.org/10.1145/321356.321363)

- [21] CHICHILNISKY, G. The foundations of statistics with black swans. *Mathematical Social Sciences* **59**(2) 184–192 (2010). <https://doi.org/10.1016/j.mathsocsci.2009.09.007>. MR2650318
- [22] CHOULDECHOVA, A. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big Data* **5**(2) 153–163 (2017).
- [23] CHURCH, A. On the concept of a random sequence. *Bulletin of the American Mathematical Society* **46**(2) 130–135 (1940). <https://doi.org/10.1090/S0002-9904-1940-07154-X>. MR0000911
- [24] COOK, D. C. The concept of independence in accounting. In *Federal Securities Law and Accounting 1933–1970: Selected Addresses* 198–222. Routledge (2020).
- [25] DAWID, P. On individual risk. *Synthese* **194**(9) 3445–3474 (2017). <https://doi.org/10.1007/s11229-015-0953-4>. MR3704899
- [26] DE COOMAN, G. and DE BOCK, J. Randomness is inherently imprecise. *International Journal of Approximate Reasoning* **141**. 28–68 (2022). <https://doi.org/10.1016/j.ijar.2021.06.018>. MR4364895
- [27] DE FINETTI, B. Funzione caratteristica di un fenomeno aleatorio. In *Atti del Congresso Internazionale dei Matematici: Bologna del 3 al 10 de settembre di 1928* 179–190 (1929).
- [28] DE MOIVRE, A. *The Doctrine of Chances: A Method of Calculating the Probabilities of Events in Play* 2nd ed. Frank Cass and Company Limited (1738/1967). MR0231695
- [29] DEUTSCH, D. *The Beginning of Infinity: Explanations that Transform the World*. Penguin (2011). MR2984795
- [30] DEVROYE, L., GYÖRFI, L. and LUGOSI, G. *A Probabilistic Theory of Pattern Recognition*. Springer (1996). <https://doi.org/10.1007/978-1-4612-0711-5>. MR1383093
- [31] DOWNEY, R. G. and HIRSCHFELDT, D. R. *Algorithmic Randomness and Complexity*. Springer (2010). <https://doi.org/10.1007/978-0-387-68441-3>. MR2732288
- [32] DURRETT, R. *Probability: Theory and Examples* 5th ed. Cambridge University Press (2019). <https://doi.org/10.1017/9781108591034>. MR3930614
- [33] DWORK, C., HARDT, M., PITASSI, T., REINGOLD, O. and ZEMEL, R. Fairness through awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference – ITCS 2012* 214–226 (2012). MR3388391
- [34] DWORK, C., LEE, D., LIN, H. and TANKALA, P. From pseudorandomness to multi-group fairness and back. In *The Thirty Sixth Annual Conference on Learning Theory* 3566–3614. PMLR (2023).
- [35] EAGLE, A. Chance versus Randomness. In (E. N. Zalta, ed.) *The Stanford Encyclopedia of Philosophy*, Metaphysics Research Lab, Stanford University, Spring 2021 edition, 2021.
- [36] FINE, T. L. *Theories of Probability: An Examination of Foundations*. Academic Press (1973). MR0433529
- [37] FRÖHLICH, C., DERR, R. and WILLIAMSON, R. C. Towards a strictly frequentist theory of imprecise probability. In (E. Miranda, I. Montes, E. Quaeghebeur and B. Vantaggi, eds.) *Proceedings of the Thirteenth International Symposium on Imprecise Probability: Theories and Applications. Proceedings of Machine Learning Research* **215** 230–240. PMLR (2023). MR4663313
- [38] FRONGILLO, R. M. and NOBEL, A. B. Memoryless sequences for general losses. *J. Mach. Learn. Res.* **21** 80 (2020). MR4119148
- [39] GALLIE, W. B. Essentially contested concepts. In *Proceedings of the Aristotelian society* **56** 167–198 (1955).
- [40] GAMMERMAN, A. and VOVK, V. Hedging predictions in machine learning. *The Computer Journal* **50**(2) 151–163 (2007).
- [41] GILLIES, D. *An Objective Theory of Probability*. Routledge (2010) (First published 1973).
- [42] GRÄDEL, E. and VÄÄNÄNEN, J. Dependence and independence. *Studia Logica* **101**(2) 399–410 (2013). <https://doi.org/10.1007/s11225-013-9479-2>. MR3038039
- [43] GUDDER, S. P. Quantum probability spaces. *Proceedings of the American Mathematical Society* **21**(2) 296–302 (1969). <https://doi.org/10.2307/2036988>. MR0243793
- [44] GUDDER, S. P. *Stochastic Methods in Quantum Mechanics*. Dover Publications Inc. (1979). MR0543489
- [45] HÁJEK, A. Fifteen arguments against hypothetical frequentism. *Erkenntnis* **70**(2) 211–235 (2009). <https://doi.org/10.1007/s10670-009-9154-1>. MR2481794
- [46] HARDT, M., PRICE, E. and SREBRO, N. Equality of opportunity in supervised learning. In *Advances in Neural Information Processing Systems* 3323–3331 (2016).
- [47] HERTWECK, C., HEITZ, C. and LOI, M. On the moral justification of statistical parity. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* 747–757 (2021).
- [48] HOGG, M. A. and ABRAMS, D. *Social Identifications: A Social Psychology of Intergroup Relations and Group Processes*. Routledge (1998).
- [49] HU, L. and KOHLER-HAUSMANN, I. What’s sex got to do with fair machine learning? arXiv preprint arXiv:2006.01770 (2020).
- [50] HUMPHREYS, P. W. Randomness, independence, and hypotheses. *Synthese* **36** 415–426 (1977). <https://doi.org/10.1007/BF00486105>. MR0517129
- [51] IVANENKO, V. I. *Decision Systems and Nonstochastic Randomness*. Springer (2010). <https://doi.org/10.1007/978-1-4419-5548-7>. MR2606231
- [52] KAMISHIMA, T., AKAHO, S., ASOH, H. and SAKUMA, J. Fairness-aware classifier with prejudice remover regularizer. In *Proceedings of the 2012th European Conference on Machine Learning and Knowledge Discovery in Databases-Volume Part II ECMLPKDD’12*: 35–50 (2012).
- [53] KILBERTUS, N., ROJAS CARULLA, M., PARASCANDOLO, G., HARDT, M., JANZING, D. and SCHÖLKOPF, B. Avoiding discrimination through causal reasoning. In *Advances in Neural Information Processing Systems* **30** (2017).
- [54] KLEINBERG, J., MULLAINATHAN, S. and RAGHAVAN, M. Inherent trade-offs in the fair determination of risk scores. In *Leibniz International Proceedings in Informatics, LIPIcs* **67** 1–22 (2017). MR3754967
- [55] KOLMOGOROV, A. N. *Grundbegriffe der Wahrscheinlichkeitsrechnung*. Springer (1933). MR0362415
- [56] KOLMOGOROV, A. N. *Foundations of the Theory of Probability: Second English Edition*. Chelsea Publishing Company (1956). MR0079843
- [57] KOLMOGOROV, A. N. Three approaches to the definition of the concept “quantity of information”. *Problemy Peredachi Informatsii* **1**(1) 3–11 (1965). MR0184801
- [58] LAFFERTY, J. and WASSERMAN, L. Challenges in statistical machine learning. *Statistica Sinica* **16**(2) 307 (2006). MR2267237
- [59] LEVIN, L. A. The concept of random sequence. *Soviet Mathematics Doklady* **14**. 1413–1416 (1973). MR0366096
- [60] LEVIN, L. A. A concept of independence with applications in various fields of mathematics. Technical Report MIT/LCS/TR-235, MIT, Laboratory for Computer Science, April 1980.
- [61] LIPPERT-RASMUSSEN, K. The badness of discrimination. *Ethical Theory and Moral Practice* **9**(2) 167–185 (2006).
- [62] VITALI, L. *Analysis 1 Lecture Notes 2013/2014*. University of Bristol, 2014.
- [63] LOHAUS, M., PERROT, M. and VON LUXBURG, U. Too relaxed to be fair. In *International Conference on Machine Learning* 6360–6369. PMLR (2020).
- [64] LOI, M., HERLITZ, A. and HEIDARI, H. A philosophical theory of fairness for prediction-based decisions. Technical report, Politecnico di Milano, 2019. <https://ssrn.com/abstract=3450300>
- [65] MARTIN-LÖF, P. The definition of random sequences. *Information and Control* **9**(6) 602–619 (1966). MR0223179
- [66] MCGARTY, C. *Categorization in Social Psychology*. Sage Publications (1999).
- [67] MENON, A. K. and WILLIAMSON, R. C. The cost of fairness in binary classification. In *Conference on Fairness, Accountability and Transparency* 107–118. PMLR (2018).

- [68] MUCHNIK, A. A., SEMENOV, A. L. and USPENSKY, V. A. Mathematical metaphysics of randomness. *Theoretical Computer Science* **207**(2) 263–317 (1998). [https://doi.org/10.1016/S0304-3975\(98\)00069-3](https://doi.org/10.1016/S0304-3975(98)00069-3). MR1643438
- [69] MURCHISON, B. C. The concept of independence in public law. *Emory Law Journal* **41**. 961 (1992).
- [70] NARENS, L. An introduction to lattice based probability theories. *Journal of Mathematical Psychology* **74**. 66–81 (2016). <https://doi.org/10.1016/j.jmp.2016.04.013>. MR3552130
- [71] NATHANSON, M. B. *Elementary Methods in Number Theory* **195**. Springer (2008). MR1732941
- [72] NAYLOR, A. W. On decomposition theory: generalized dependence. *IEEE Transactions on Systems, Man, and Cybernetics* **11**(10) 699–713 (1981). <https://doi.org/10.1109/TSMC.1981.4308590>. MR0641948
- [73] DONALD and ORNSTEIN, S. Ergodic theory, randomness, and “chaos”. *Science* **243**(4888) 182–187 (1989). <https://doi.org/10.1126/science.243.4888.182>. MR0981173
- [74] PARKER, J. M. *Randomness and legitimacy in selecting democratic representatives*. The University of Texas at Austin, 2011.
- [75] PASSI, S. and BAROCAS, S. Problem formulation and fairness. In *Proceedings of the Conference on Fairness, Accountability, and Transparency* 39–48 (2019).
- [76] POPPER, K. *The Logic of Scientific Discovery*. Routledge (2010). MR1195353
- [77] PORTER, C. P. *Mathematical and philosophical perspectives on algorithmic randomness*. PhD thesis, University of Notre Dame, 2012. MR3217940
- [78] PTÁK, P. Concrete quantum logics. *International Journal of Theoretical Physics* **39**(3) 827–837 (2000). <https://doi.org/10.1023/A:1003626929648>. MR1792201
- [79] PUTNAM, H. *The Meaning of the Concept of Probability in Application to Finite Sequences* (Routledge Revivals). Garland Publishing (1990).
- [80] RAGHAVAN, M., BAROCAS, S., KLEINBERG, J. and LEVY, K. Mitigating bias in algorithmic hiring: Evaluating claims and practices. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* 469–481 (2020).
- [81] BHASKARA RAO, K. P. S. and BHASKARA RAO, M. *Theory of Charges: A Study of Finitely Additive Measures*. Academic Press (1983). MR0751777
- [82] RAWLS, J. *A Theory of Justice*. The Belknap Press of Harvard University Press (1971).
- [83] RÁZ, T. Group fairness: Independence revisited. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* 129–137 (2021). MR4297512
- [84] SCANTAMBURLO, T. Non-empirical problems in fair machine learning. *Ethics and Information Technology* **23**(4) 703–712 (2021).
- [85] SCHNORR, C.-P. *Zufälligkeit und Wahrscheinlichkeit: eine algorithmische Begründung der Wahrscheinlichkeitstheorie*. *Lecture Notes in Mathematics* **218**. Springer (1971). MR0414225
- [86] SCHNORR, C.-P. The process complexity and effective random tests. In *Proceedings of the Fourth Annual ACM Symposium on Theory of Computing* 168–176 (1972). [https://doi.org/10.1016/S0022-0000\(73\)80030-3](https://doi.org/10.1016/S0022-0000(73)80030-3). MR0325366
- [87] SCHNORR, C.-P. A survey of the theory of random sequences. In *Basic Problems in Methodology and Linguistics* 193–211. Springer (1977). MR0517133
- [88] SCHURZ, G. and LEITGEB, H. Finitistic and frequentistic approximation of probability measures with or without σ -additivity. *Studia Logica* **89**(2) 257–283 (2008). <https://doi.org/10.1007/s11225-008-9128-3>. MR2429951
- [89] SHAFER, G. Discussion on hedging predictions in machine learning by A. Gammernan and V. Vovk. *The Computer Journal* **50**(2) 164–172 (2007). <https://doi.org/10.1093/comjnl/bxl066>
- [90] SHAFER, G. and VOVK, V. *Game-Theoretic Foundations for Probability and Finance* **455**. John Wiley & Sons (2019). <https://doi.org/10.1002/0471249696>. MR1852450
- [91] SHALEV-SHWARTZ, S. and BEN-DAVID, S. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press (2014). MR3277164
- [92] SIMPSON, J. and WEINE, E. “independence, n”. In *Oxford English Dictionary*, 2nd ed. Oxford University Press (1989).
- [93] SPOHN, W. Stochastic independence, causal independence, and shieldability. *Journal of Philosophical Logic* **9**(1) 73–99 (1980). <https://doi.org/10.1007/BF00258078>. MR0563250
- [94] STEINWART, I., HUSH, D. and SCOVEL, C. Learning from dependent observations. *Journal of Multivariate Analysis* **100**(1) 175–194 (2009). <https://doi.org/10.1016/j.jmva.2008.04.001>. MR2460486
- [95] STONE, P. *The Luck of the Draw: The Role of Lotteries in Decision-Making*. Oxford University Press (2011).
- [96] TADAKI, K. An operational characterization of the notion of probability by algorithmic randomness. In *Proceedings of the 37th Symposium on Information Theory and its Applications (SITA2014)* **5** 389–394 (2014).
- [97] TAJFEL, H. Social identity and intergroup behaviour. *Social Science Information* **13**(2) 65–93 (1974).
- [98] TAO, T. 275a, notes 2: Product measures and independence, October 2015. <https://terrytao.wordpress.com/2015/10/12/275a-notes-2-product-measures-and-independence/>
- [99] USPENSKII, V. A., SEMENOV, A. L. and SHEN, A. KH. Can an individual sequence of zeros and ones be random? *Russian Mathematical Surveys* **45**(1) 121 (1990). <https://doi.org/10.1070/RM1990v045n01ABEH002321>. MR1050929
- [100] VAN LAMBALGEN, M. Von Mises’ definition of random sequences reconsidered. *The Journal of Symbolic Logic* **52**(3) 725–755 (1987). <https://doi.org/10.2307/2274360>. MR0902987
- [101] VAN LAMBALGEN, M. The axiomatization of randomness. *The Journal of Symbolic Logic* **55**(3) 1143–1167 (1990). <https://doi.org/10.2307/2274480>. MR1071321
- [102] VILLE, J. *Étude critique de la notion de collectif*. Gauthier-Villars (1939). MR3533075
- [103] VOLCHAN, S. B. What is a random sequence? *The American Mathematical Monthly* **109**(1) 46–63 (2002). <https://doi.org/10.2307/2695767>. MR1903512
- [104] VON COLLANI, E. A note on the concept of independence. *Economic Quality Control* **21**(1) 155–164 (2006).
- [105] VON MISES, R. Grundlagen der Wahrscheinlichkeitsrechnung. *Mathematische Zeitschrift* **5**(1) 52–99 (1919). <https://doi.org/10.1007/BF01203155>. MR1544374
- [106] VON MISES, R. *Probability, Statistics, and Truth*. Dover Publications, Inc. (1981). MR0668875
- [107] VON MISES, R. and GEIRINGER, H. *Mathematical Theory of Probability and Statistics*. Academic Press (1964). MR0178486
- [108] VOVK, V. G. A logic of probability, with application to the foundations of statistics. *Journal of the Royal Statistical Society: Series B (Methodological)* **55**(2) 317–341 (1993). MR1224399
- [109] WASSERMAN, L. *All of Statistics: A Concise Course in Statistical Inference*. Springer (2004). <https://doi.org/10.1007/978-0-387-21736-9>. MR2055670
- [110] WICK, M., PANDA, S. and TRISTAN, J.-B. Unlocking fairness: a trade-off revisited. In *Proceedings of the 37th International Conference on Machine Learning* **32** (2020).
- [111] WILLIAMSON, R. C. and MENON, A. Fairness risk measures. In *International Conference on Machine Learning* 6786–6797. PMLR (2019).

Rabanus Derr. Maria-von-Linden-Str. 6, 72076 Tübingen, University of Tübingen and Tübingen AI Center, Germany. E-mail address: rabanus.derr@uni-tuebingen.de

Robert C. Williamson. Maria-von-Linden-Str. 6, 72076 Tübingen, University of Tübingen and Tübingen AI Center, Germany. E-mail address: bob.williamson@uni-tuebingen.de