

ON THE COMPUTABILITY OF CONDITIONAL PROBABILITY

NATHANAEL L. ACKERMAN, CAMERON E. FREER, AND DANIEL M. ROY

ABSTRACT. As inductive inference and machine learning methods in computer science see continued success, researchers are aiming to describe even more complex probabilistic models and inference algorithms. What are the limits of mechanizing probabilistic inference? We investigate the computability of conditional probability, a fundamental notion in probability theory and a cornerstone of Bayesian statistics, and show that there are computable joint distributions with noncomputable conditional distributions, ruling out the prospect of general inference algorithms, even inefficient ones. Specifically, we construct a pair of computable random variables in the unit interval such that the conditional distribution of the first variable given the second encodes the halting problem. Nevertheless, probabilistic inference is possible in many common modeling settings, and we prove several results giving broadly applicable conditions under which conditional distributions are computable. In particular, conditional distributions become computable when measurements are corrupted by independent computable noise with a sufficiently smooth density.

1. INTRODUCTION

The use of probability to reason about uncertainty is key to modern science and engineering, and the operation of *conditioning*, used to perform Bayesian inductive reasoning in probabilistic models, directly raises many of its most important computational problems. Faced with probabilistic models of increasingly complex phenomena that stretch or exceed the limitations of existing representations and algorithms, researchers have proposed new representations and formal languages for describing joint distributions on large collections of random variables, and have developed new algorithms for performing automated probabilistic inference. What are the limits of this endeavor? Can we hope to automate probabilistic reasoning via a general inference algorithm that can compute conditional probabilities for an *arbitrary* computable joint distribution?

We demonstrate that there are computable joint distributions with noncomputable conditional distributions. Of course, the fact that generic algorithms cannot exist for computing conditional probabilities does not rule out the possibility that large classes of distributions may be amenable to automated inference. The challenge for mathematical theory is to explain the widespread success of probabilistic methods and characterize the circumstances when conditioning is possible. In this

Keywords. Computable probability theory, conditional probability, real computation.

2010 Mathematics Subject Classification. Primary: 03D78; Secondary: 62F15, 68T37, 60B05, 03F60, 65G50, 60A05, 60A10.

vein, we describe broadly applicable conditions under which conditional probabilities are computable.

1.1. Probabilistic programming. Within probabilistic Artificial Intelligence (AI) and machine learning, the study of formal languages and algorithms for describing and computing answers from probabilistic models is the subject of *probabilistic programming*. Probabilistic programming languages themselves build on modern programming languages and their facilities for recursion, abstraction, modularity, etc., to enable practitioners to define intricate, in some cases infinite-dimensional, models by implementing a generative process that produces an exact sample from the model’s joint distribution. (See, e.g., IBAL [Pfe01], λ_{\circ} [PPT08], Church [GMR⁺08], and HANSEI [KS09]. For related and earlier efforts, see, e.g., PHA [Poo91], Infer.NET [MWGK10], Markov Logic [RD06]. Probabilistic programming languages have been the focus of a long tradition of research within programming languages, model checking and formal methods.) In many of these languages, one can easily represent the higher-order stochastic processes (e.g., distributions on data structures, distributions on functions, and distributions on distributions) that are essential building blocks in modern nonparametric Bayesian statistics. In fact, the most expressive such languages are each capable of describing the same robust class as the others—the class of *computable distributions*, which delineates those from which a probabilistic Turing machine can sample to arbitrary accuracy.

Traditionally, inference algorithms for probabilistic models have been derived and implemented by hand. In contrast, probabilistic programming systems have introduced varying degrees of support for computing conditional distributions. Given the rate of progress toward broadening the scope of these algorithms, one might hope that there would eventually be a generic algorithm supporting the entire class of computable distributions.

Despite recent progress towards a general such algorithm, support for conditioning with respect to continuous random variables has remained ad-hoc and incomplete. Our results explain why this is necessarily the case.

1.2. Computable Distributions. In order to characterize the computational limits of probabilistic inference, we work within the framework of *computable probability theory*, which pertains to the computability of distributions, random variables, and probabilistic operations; and builds on the classical computability theory of deterministic functions. Just as the notion of a Turing machine allows one to prove results about discrete computations performed using an arbitrary (sufficiently powerful) programming language, the notion of a probabilistic Turing machine provides a basis for precisely describing the operations that probabilistic programming languages are capable of performing.

The tools for describing computability in this setting are drawn from the theory of *computable metric spaces*, within the subject of *computable analysis*. This theory gives us the ability to study distributions on arbitrary computable metric spaces, including, e.g., distributions on distributions. In Section 2 we present the necessary definitions and results from computable probability theory.

1.3. Conditional Probability. For an experiment with a discrete set of outcomes, computing conditional probabilities is, in principle, straightforward as it is simply a ratio of probabilities. However, in the case of conditioning on the value of a continuous random variable, this ratio is undefined. Furthermore, in modern Bayesian statistics, and especially the probabilistic programming setting, it is common to place distributions on higher-order objects, and so one is already in a situation where elementary notions of conditional probability are insufficient and more sophisticated measure-theoretic notions are necessary.

Kolmogorov [Kol33] gave an axiomatic characterization of conditional probabilities, but this definition provides no recipe for their calculation. Other issues also arise: In this setting, conditional probabilities are formalized as measurable functions that are defined only up to measure zero sets. Therefore, without additional assumptions, a conditional probability is not necessarily well-defined for any particular value of the conditioning random variable. This has long been understood as a challenge for statistical applications, in which one wants to evaluate conditional probabilities given particular values for observed random variables. In this paper, we are therefore especially interested in situations where it makes sense to ask for the conditional distribution given a particular point. In particular, we focus on the case when conditional distributions are everywhere or almost everywhere continuous, and thus can be given a unique definition for individual points in the support of the underlying measure. As we will argue, this is necessary if there is to be any hope of conditioning being computable.

Under certain conditions, such as when conditional densities exist, conditioning can proceed using the classic Bayes' rule; however, it may not be possible to compute the density of a computable distribution (if the density exists at all). The probability and statistics literature contains many ad-hoc techniques for calculating conditional probabilities in special circumstances, and this state of affairs motivated much work on constructive definitions (such as those due to Tjur [Tju74], [Tju75], [Tju80], Pfanzagl [Pfa79], and Rao [Rao88], [Rao05]), but this work has often not been sensitive to issues of computability.

We recall the basics of the measure-theoretic approach to conditional distributions in Section 3, and in Section 4 we use notions from computable probability theory to consider the sense in which conditioning could be potentially computable.

1.4. Other Related Work. Conditional probabilities for computable distributions on finite, discrete sets are clearly computable, but may not be efficiently so. In this finite discrete setting, there are already interesting questions of computational complexity, which have been explored by a number of authors through extensions of Levin's theory of average-case complexity [Lev86]. For example, under cryptographic assumptions, it is difficult to sample from the conditional distribution of a uniformly-distributed binary string of length n given its image under a one-way function. This can be seen to follow from the work of Ben-David, Chor, Goldreich, and Luby [BCGL92] in their theory of polynomial-time samplable distributions, which has since been extended by Yamakami [Yam99] and others. Extending these

complexity results to the more general setting considered here could bear on the practice of statistical AI and machine learning.

Osherson, Stob, and Weinstein [OSW88] study learning theory in the setting of *identifiability in the limit* (see [Gol67] and [Put65] for more details on this setting) and prove that a certain type of “computable Bayesian” learner fails to identify the index of a (computably enumerable) set that is “computably identifiable” in the limit. More specifically, a so-called “Bayesian learner” is required to return an index for a set with the highest conditional probability given a finite prefix of an infinite sequence of random draws from the unknown set. An analysis by Roy [Roy11] of their construction reveals that the conditional distribution of the index given the infinite sequence is an everywhere discontinuous function (on every measure one set), hence noncomputable for much the same reason as our elementary construction involving a mixture of measures concentrated on the rationals and on the irrationals (see Section 5). As we argue, it is more appropriate to study the operator when it is restricted to those random variables whose conditional distributions admit versions that are continuous everywhere, or at least on a measure one set.

Our work is distinct from the study of conditional distributions with respect to priors that are universal for partial computable functions (as defined using Kolmogorov complexity) by Solomonoff [Sol64], Zvonkin and Levin [ZL70], and Hutter [Hut07]. The computability of conditional distributions also has a rather different character in Takahashi’s work on the algorithmic randomness of points defined using universal Martin-Löf tests [Tak08]. The objects with respect to which one is conditioning in these settings are typically *computably enumerable*, but not computable. In the present paper, we are interested in the problem of computing conditional distributions of random variables that are *computable*, even though the conditional distribution may itself be noncomputable.

In the most abstract setting, conditional probabilities can be constructed as Radon-Nikodym derivatives. In work motivated by questions in algorithmic randomness, Hoyrup and Rojas [HR11] study notions of computability for absolute continuity and for Radon-Nikodym derivatives as elements in L^1 , i.e., the space of integrable functions. They demonstrate that there are computable measures whose Radon-Nikodym derivatives are not computable as points in L^1 , but these counterexamples do not correspond with conditional probabilities of computable random variables. Hoyrup, Rojas and Weihrauch [HRW11] show an equivalence between the problem of computing general Radon-Nikodym derivatives as elements in L^1 and computing the characteristic function of computably enumerable sets. However, conditional probabilities are a special case of Radon-Nikodym derivatives, and moreover, a computable element in L^1 is not well-defined at points, and so is not ideal for statistical purposes. Using their machinery, we demonstrate the non- L^1 -computability of our main construction. But the main goal of our paper is to provide a detailed analysis of the situation where it makes sense to ask for the conditional probability at points, which is the more relevant scenario for statistical inference.

1.5. Summary of Results. Following our presentation of computable probability theory and conditional probability in Sections 2 through 4, we provide our main positive and negative results about the computability of conditional probability, which we now summarize.

In Proposition 5.1, we construct a pair (X, C) of computable random variables such that every version of the conditional probability $\mathbf{P}[C = 1|X]$ is discontinuous everywhere, even when restricted to a \mathbf{P}_X -measure one subset. (We make these notions precise in Section 4.) The construction makes use of the elementary fact that the indicator function for the rationals in the unit interval—the so-called Dirichlet function—is itself nowhere continuous.

Because every function computable on a domain D is continuous on D , discontinuity is a fundamental barrier to computability, and so this construction rules out the possibility of a completely general algorithm for conditioning. A natural question is whether conditioning is a computable operation when we restrict the operator to random variables for which *some* version of the conditional distribution is continuous everywhere, or at least on a measure one set.

Our central result, Theorem 6.7, states that conditioning is not a computable operation on computable random variables, even in this restricted setting. We construct a pair (X, N) of computable random variables such that there is a version of the conditional distribution $\mathbf{P}[N|X]$ that is continuous on a measure one set, but no version of $\mathbf{P}[N|X]$ is computable. (Indeed, every individual conditional probability fails to be even lower semicomputable on any set of sufficient measure.) Moreover, the noncomputability of $\mathbf{P}[N|X]$ is at least as hard as the halting problem, in that if some oracle A computes $\mathbf{P}[N|X]$, then A computes the halting problem. The construction involves encoding the halting times of all Turing machines into the conditional distribution $\mathbf{P}[N|X]$, while ensuring that the joint distribution remains computable.

In Theorem 7.6 we strengthen our central result by constructing a pair of computable random variables whose conditional distribution is noncomputable but has an everywhere continuous version with infinitely differentiable conditional probabilities. This construction proceeds by smoothing out the distribution constructed in Theorem 6.7, but in such a way that one can still compute the halting problem relative to the conditional distribution.

Despite the noncomputability of conditioning in general, conditional distributions are often computable in practice. We provide some explanation of this phenomenon by characterizing several circumstances in which conditioning *is* a computable operation. Under suitable computability hypotheses, conditioning is computable in the discrete setting (Lemma 8.1) and where there is a conditional density (Corollary 8.8).

We also characterize a situation in which conditioning is possible in the presence of noisy data, capturing many natural models in science and engineering. Let U, V and E be computable random variables, and suppose that \mathbf{P}_E is absolutely continuous with a bounded computable density p_E and E is independent of U and V . We can think of $U + E$ as the corruption of an idealized measurement U by independent

source of additive error E . In Corollary 8.9, we show that the conditional distribution $\mathbf{P}[(U, V) | U + E]$ is computable (even if $\mathbf{P}[(U, V) | U]$ is not). Finally, we discuss how symmetry can contribute to the computability of conditional distributions.

2. COMPUTABLE PROBABILITY THEORY

We now give some background on computable probability theory, which will enable us to formulate our results. The foundations of the theory include notions of computability for probability measures developed by Edalat [Eda96], Weihrauch [Wei99], Schroeder [Sch07b], and Gács [Gác05]. Computable probability theory itself builds off notions and results in computable analysis. For a general introduction to this approach to real computation, see Weihrauch [Wei00], Braverman [Bra05] or Braverman and Cook [BC06].

2.1. Computable and C.e. Reals. We first recall some elementary definitions from computability theory (see, e.g. Rogers [Rog87, Ch. 5]). We say that a set of natural numbers (potentially in some correspondence with, e.g., rationals, integers, or other finitely describable objects with an implicit enumeration) is *computably enumerable* (c.e.) when there is a computer program that outputs every element of the set eventually. We say that a sequence of sets $\{B_n\}$ is c.e. *uniformly in n* when there is a computer program that, on input n , outputs every element of B_n eventually. A set is co-c.e. when its complement is c.e. (and so the (uniformly) computable sets are precisely those that are both (uniformly) c.e. and co-c.e.).

We now recall basic notions of computability for real numbers (see, e.g., [Wei00, Ch. 4.2] or [Nie09, Ch. 1.8]). We say that a real r is a *c.e. real* when the set of rationals $\{q \in \mathbb{Q} : q < r\}$ is c.e. Similarly, a *co-c.e. real* is one for which $\{q \in \mathbb{Q} : q > r\}$ is c.e. (C.e. and co-c.e. reals are sometimes called *left-c.e.* and *right-c.e.* reals, respectively.) A real r is *computable* when it is both c.e. and co-c.e. Equivalently, a real is computable when there is a program that approximates it to any given accuracy (e.g., given an integer k as input, the program reports a rational that is within 2^{-k} of the real). A function $f : \mathbb{N} \rightarrow \mathbb{R}$ is lower (upper) semicomputable when $f(n)$ is a c.e. (co-c.e.) real, uniformly in n (or more precisely, when the rational lower (upper) bounds of $f(n)$ are c.e. uniformly in n). The function f is computable if and only if it is both lower and upper semicomputable.

2.2. Computable Metric Spaces. Computable metric spaces, as developed in computable analysis [Hem02], [Wei93] and effective domain theory [JB97], [EH98], provide a convenient framework for formulating results in computable probability theory. For consistency, we largely use definitions from [HR09a] and [GHR10]. Additional details about computable metric spaces can also be found in [Wei00, Ch. 8.1] and [Gác05, §B.3].

Definition 2.1 (Computable metric space [GHR10, Def. 2.3.1]). A **computable metric space** is a triple (S, δ, \mathcal{D}) for which δ is a metric on the set S satisfying

- (1) (S, δ) is a complete separable metric space;

- (2) $\mathcal{D} = \{s_i\}_{i \in \mathbb{N}}$ is an enumeration of a dense subset of S , called **ideal points**; and,
 (3) the real numbers $\delta(s_i, s_j)$ are computable, uniformly in i and j .

Let $B(s_i, q_j)$ denote the ball of radius q_j centered at s_i . We call

$$\mathcal{B}_S := \{B(s_i, q_j) : s_i \in \mathcal{D}, q_j \in \mathbb{Q}, q_j > 0\} \quad (1)$$

the **ideal balls of S** , and fix the canonical enumeration of them induced by that of \mathcal{D} and \mathbb{Q} .

For example, the set $\{0, 1\}$ is a computable metric space under the discrete metric, characterized by $\delta(0, 1) = 1$. Cantor space, the set $\{0, 1\}^\infty$ of infinite binary sequences, is a computable metric space under its usual metric and the dense set of eventually constant strings (under a standard enumeration of finite strings). The set \mathbb{R} of real numbers is a computable metric space under the Euclidean metric with the dense set \mathbb{Q} of rationals (under its standard enumeration).

We let \mathcal{B}_S denote the Borel σ -algebra on a metric space S , i.e., the σ -algebra generated by the open balls of S . In this paper, measurable functions will always be with respect to the Borel σ -algebra of a metric space.

Definition 2.2 (Computable point [GHR10, Def. 2.3.2]). Let (S, δ, \mathcal{D}) be a computable metric space. A point $x \in S$ is **computable** when there is a program that enumerates a sequence $\{x_i\}$ in \mathcal{D} where $\delta(x_i, x) < 2^{-i}$ for all i . We call such a sequence $\{x_i\}$ a **representation** of the point x .

Remark 2.3. A real $\alpha \in \mathbb{R}$ is computable (as in Section 2.1) if and only if α is a computable point of \mathbb{R} (as a computable metric space). Although most of the familiar reals are computable, there are only countably many computable reals, and so almost every real is not computable.

The notion of a c.e. open set (or Σ_1^0 class) is fundamental in classical computability theory, and admits a simple definition in an arbitrary computable metric space.

Definition 2.4 (C.e. open set [GHR10, Def. 2.3.3]). Let (S, δ, \mathcal{D}) be a computable metric space with the corresponding enumeration $\{B_i\}_{i \in \mathbb{N}}$ of the ideal open balls \mathcal{B}_S . We say that $U \subseteq S$ is a **c.e. open set** when there is some c.e. set $E \subseteq \mathbb{N}$ such that $U = \bigcup_{i \in E} B_i$.

Note that the class of c.e. open sets is closed under computable unions and finite intersections.

A computable function can be thought of as a continuous function whose local modulus of continuity is witnessed by a program. It is important to consider the computability of *partial* functions, because many natural and important random variables are continuous only on a measure one subset of their domain.

Definition 2.5 (Computable partial function [GHR10, Def. 2.3.6]). Let $(S, \delta_S, \mathcal{D}_S)$ and $(T, \delta_T, \mathcal{D}_T)$ be computable metric spaces, the latter with the corresponding enumeration $\{B_n\}_{n \in \mathbb{N}}$ of the ideal open balls \mathcal{B}_T . A function $f : S \rightarrow T$ is said to be **computable on $R \subseteq S$** when there is a computable sequence $\{U_n\}_{n \in \mathbb{N}}$ of c.e.

open sets $U_n \subseteq S$ such that $f^{-1}[B_n] \cap R = U_n \cap R$ for all $n \in \mathbb{N}$. We call such a sequence $\{U_n\}_{n \in \mathbb{N}}$ a **witness** to the computability of f .

In particular, if f is computable on R , then the inverse image of c.e. open sets are c.e. open (in R) sets, and so we can see computability as a natural restriction on continuity.

Remark 2.6. Let S and T be computable metric spaces. If $f : S \rightarrow T$ is computable on some subset $R \subseteq S$, then for every *computable* point $x \in R$, the point $f(x)$ is also computable. One can show that f is computable on R when there is a program that uniformly transforms representations of points in R to representations of points in S . (For more details, see [HR09a, Prop. 3.3.2].)

Remark 2.7. Suppose that $f : S \rightarrow T$ is computable on $R \subseteq S$ with $\{U_n\}_{n \in \mathbb{N}}$ a witness to the computability of f . One can show that there is an effective G_δ set $R' \supseteq R$ and a function $f' : S \rightarrow T$ such that f' is computable on R' , the restriction of f' to R and f are equal as functions, and $\{U_n\}_{n \in \mathbb{N}}$ is a witness to the computability of f' . Furthermore, a G_δ -code for R' can be chosen uniformly in the witness $\{U_n\}_{n \in \mathbb{N}}$. One could consider such an f' to be a canonical representative of the computable partial function f with witness $\{U_n\}_{n \in \mathbb{N}}$. Note, however, that the G_δ -set chosen depends not just on f , but also on the witness $\{U_n\}_{n \in \mathbb{N}}$. In particular, it is possible that two distinct witnesses to the computability of f could result in distinct G_δ -sets.

2.3. Computable Random Variables. Intuitively, a random variable maps an input source of randomness to an output, inducing a distribution on the output space. Here we will use a sequence of independent fair coin flips as our source of randomness. We formalize this via the probability space $(\{0, 1\}^\infty, \mathcal{F}, \mathbf{P})$, where $\{0, 1\}^\infty$ is the product space of infinite binary sequences, \mathcal{F} is its Borel σ -algebra (generated by the set of basic clopen cylinders extending each finite binary sequence), and \mathbf{P} is the product measure of the uniform distribution on $\{0, 1\}$. Henceforth we will take $(\{0, 1\}^\infty, \mathcal{F}, \mathbf{P})$ to be the basic probability space, unless otherwise stated.

For a measure space $(\Omega, \mathcal{G}, \mu)$, a set $E \in \mathcal{G}$ is a μ -**null set** when $\mu E = 0$. More generally, for $p \in [0, \infty]$, we say that E is a μ -**measure p set** when $\mu E = p$. A relation between functions on Ω is said to hold μ -**almost everywhere** (abbreviated μ -**a.e.**) if it holds for all $\omega \in \Omega$ outside of a μ -null set. When μ is a probability measure, then we may instead say that the relation holds for μ -**almost all** ω (abbreviated μ -**a.a.**). We say that an event $E \in \mathcal{G}$ occurs μ -**almost surely** (abbreviated μ -**a.s.**) when $\mu E = 1$. In each case, we may drop the prefix μ - when it is clear from context (in particular, when it holds of \mathbf{P}).

We will use a SANS SERIF font for random variables.

Definition 2.8 (Random variable and its distribution). Let S be a computable metric space. A **random variable in S** is a function $X : \{0, 1\}^\infty \rightarrow S$ that is measurable with respect to the Borel σ -algebras of $\{0, 1\}^\infty$ and S . For a measurable subset $A \subseteq S$, we let $\{X \in A\}$ denote the inverse image

$$X^{-1}[A] = \{\omega \in \{0, 1\}^\infty : X(\omega) \in A\}, \quad (2)$$

and for $x \in S$ we similarly define the event $\{X = x\}$. We will write \mathbf{P}_X for the **distribution of X** , which is the measure on S defined by $\mathbf{P}_X(\cdot) := \mathbf{P}\{X \in \cdot\}$.

Definition 2.9 (Computable random variable). Let S be a computable metric space. Then a random variable X in S is a **computable random variable** when X is computable on some \mathbf{P} -measure one subset of $\{0, 1\}^\infty$.

More generally, for a probability measure μ and function f between computable metric spaces, we say that f is **μ -almost computable** when it is computable on a μ -measure one set. (See [HR09a] for further development of the theory of almost computable functions.)

Intuitively, X is a computable random variable when there is a program that, given access to an oracle bit tape $\omega \in \{0, 1\}^\infty$, outputs a representation of the point $X(\omega)$ (i.e., enumerates a sequence $\{x_i\}$ in \mathcal{D} where $\delta(x_i, X(\omega)) < 2^{-i}$ for all i), for all but a measure zero subset of bit tapes $\omega \in \{0, 1\}^\infty$.

Even though the source of randomness is a sequence of discrete bits, there are computable random variables with *continuous* distributions, such as a uniform random variable (gotten by subdividing the interval according to the random bittape) or an i.i.d.-uniform sequence (by splitting up the given element of $\{0, 1\}^\infty$ into countably many disjoint subsequences and dovetailing the constructions). (For details, see [FR10, Ex. 3, 4].) All of the standard distributions (standard normal, uniform, geometric, exponential, etc.) found in probability textbooks, as well the transformations of these distributions by computable (or almost computable) functions, are easily shown to be computable distributions.

It is crucial that we consider random variables that are computable only on a \mathbf{P} -measure one subset of $\{0, 1\}^\infty$. Consider the following example: For a real $\alpha \in [0, 1]$, we say that a binary random variable $X : \{0, 1\}^\infty \rightarrow \{0, 1\}$ is a **Bernoulli(α)** random variable when $\mathbf{P}_X\{1\} = \alpha$. There is a Bernoulli($\frac{1}{2}$) random variable that is computable on all of $\{0, 1\}^\infty$, given by the program that simply outputs the first bit of the input sequence. Likewise, when α is **dyadic** (i.e., a rational with denominator a power of 2), there is a Bernoulli(α) random variable that is computable on all of $\{0, 1\}^\infty$. However, this is not possible for any other choices of α (e.g., $\frac{1}{3}$).

Proposition 2.10. *Let $\alpha \in [0, 1]$ be a nondyadic real. Every Bernoulli(α) random variable $X : \{0, 1\}^\infty \rightarrow \{0, 1\}$ is discontinuous, hence not computable on all of $\{0, 1\}^\infty$.*

Proof. Assume X is continuous. Let $Z_0 := X^{-1}(0)$ and $Z_1 := X^{-1}(1)$. Then $\{0, 1\}^\infty = Z_0 \cup Z_1$, and so both are closed (as well as open). The compactness of $\{0, 1\}^\infty$ implies that these closed subspaces are also compact, and so Z_0 and Z_1 can each be written as the finite disjoint union of clopen basis elements. But each of these elements has dyadic measure, hence their sum cannot be either α or $1 - \alpha$, contradicting the fact that $\mathbf{P}(Z_1) = 1 - \mathbf{P}(Z_0) = \alpha$. \square

On the other hand, for an arbitrary computable $\alpha \in [0, 1]$, consider the random variable X_α given by $X_\alpha(x) = 1$ if $\sum_{i=0}^\infty x_i 2^{-i-1} < \alpha$ and 0 otherwise. This construction, due to [Man73], is a Bernoulli(α) random variable and is computable on every

point of $\{0, 1\}^\infty$ other than a binary expansion of α . Not only are these random variables computable, but they can be shown to be optimal in their use of input bits, via the classic analysis of rational-weight coins by Knuth and Yao [KY76]. Hence it is natural to admit as computable random variables those measurable functions that are computable only on a \mathbf{P} -measure one subset of $\{0, 1\}^\infty$, as we have done.

2.4. Computable Probability Measures. In this section, we introduce the class of computable probability measures and describe their connection with computable random variables.

Let $(S, \delta_S, \mathcal{D}_S)$ be a computable metric space, and let $\mathcal{B}(S)$ denote its Borel sets. We will denote by $\mathcal{M}(S)$ the set of (Borel) measures on S and by $\mathcal{M}_1(S)$ the subset which are probability measures. Consider the subset $\mathcal{D}_P \subseteq \mathcal{M}_1(S)$ comprised of those probability measures that are concentrated on a finite subset of \mathcal{D}_S and where the measure of each atom is rational, i.e., $\nu \in \mathcal{D}_P$ if and only if $\nu = \sum_{i=1}^k q_i \delta_{t_i}$ for some rationals $q_i \geq 0$ such that $\sum_{i=1}^k q_i = 1$ and some points $t_i \in \mathcal{D}_S$, where δ_{t_i} denotes the Dirac delta on t_i . Gács [Gács05, §B.6.2] shows that \mathcal{D}_P is dense in the Prokhorov metric δ_P given by

$$\delta_P(\mu, \nu) := \inf \{ \varepsilon > 0 : \forall A \in \mathcal{B}(S), \mu(A) \leq \nu(A^\varepsilon) + \varepsilon \}, \quad (3)$$

where

$$A^\varepsilon := \{ p \in S : \exists q \in A, \delta_S(p, q) < \varepsilon \} = \bigcup_{p \in A} B_\varepsilon(p) \quad (4)$$

is the ε -neighborhood of A and $B_\varepsilon(p)$ is the open ball of radius ε about p . Moreover, $(\mathcal{M}_1(S), \delta_P, \mathcal{D}_P)$ is a computable metric space. (See also [HR09a, Prop. 4.1.1].) We say that $\mu \in \mathcal{M}_1(S)$ is a computable (Borel) probability measure when μ is a computable point in $\mathcal{M}_1(S)$ as a computable metric space. Note that the measure \mathbf{P} on $\{0, 1\}^\infty$ is a computable probability measure.

We can characterize the class of computable probability measures in terms of the computability of the measure of open sets.

Theorem 2.11 ([HR09a, Thm. 4.2.1]). *Let $(S, \delta_S, \mathcal{D}_S)$ be a computable metric space. A probability measure $\mu \in \mathcal{M}_1(S)$ is computable if and only if the measure $\mu(A)$ of a c.e. open set $A \subseteq S$ is a c.e. real, uniformly in A .*

Definition 2.12 (Computable probability space [GHR10, Def. 2.4.1]). A **computable probability space** is a pair (S, μ) where S is a computable metric space and μ is a computable probability measure on S .

Let (S, μ) be a computable probability space. We know that the measure of a c.e. open set A is a c.e. real, but is not in general a computable real. On the other hand, if A is a decidable subset (i.e., $S \setminus A$ is c.e. open) then $\mu(S \setminus A)$ a c.e. real, and therefore, by the identity $\mu(A) + \mu(S \setminus A) = 1$, we have that $\mu(A)$ is a computable real. In connected spaces, the only decidable subsets are the empty set and the whole space. However, there exists a useful surrogate when dealing with measure spaces.

Definition 2.13 (Almost decidable set [GHR10, Def. 3.1.3]). Let (S, μ) be a computable probability space. A (Borel) measurable subset $A \subseteq S$ is said to be **μ -almost decidable** when there are two c.e. open sets U and V such that $U \subseteq A$ and $V \subseteq S \setminus A$ and $\mu(U) + \mu(V) = 1$.

The following is immediate.

Lemma 2.14 ([GHR10, Prop. 3.1.1]). *Let (S, μ) be a computable probability space, and let A be μ -almost decidable. Then $\mu(A)$ is a computable real.*

While we may not be able to compute the probability measure of ideal balls, we can compute a new basis of ideal balls for which we can. (See also Bosserhoff [Bos08, Lem. 2.15].)

Lemma 2.15 ([GHR10, Thm. 3.1.2]). *Let (S, μ) be a computable probability space, and let \mathcal{D}_S be the ideal points of S with standard enumeration $\{d_i\}_{i \in \mathbb{N}}$. There is a computable sequence $\{r_j\}_{j \in \mathbb{N}}$ of reals, dense in the positive reals, such that the balls $\{B(d_i, r_j)\}_{i, j \in \mathbb{N}}$ form a basis of μ -almost decidable sets.*

We now show that every c.e. open set is the union of a computable sequence of almost decidable subsets.

Lemma 2.16 (Almost decidable subsets). *Let (S, μ) be a computable probability space and let V be a c.e. open set. Then, uniformly in V , we can compute a sequence of μ -almost decidable sets $\{V_k\}_{k \in \mathbb{N}}$ such that, for each k , $V_k \subseteq V_{k+1}$, and $\bigcup_{k \in \mathbb{N}} V_k = V$.*

Proof. Let $\{B_k\}_{k \in \mathbb{N}}$ be a standard enumeration of the ideal balls of S where $B_k = B(d_{m_k}, q_{l_k})$, and let $E \subseteq \mathbb{N}$ be a c.e. set such that $V = \bigcup_{k \in E} B_k$. Let $\{B(d_i, r_j)\}_{i, j \in \mathbb{N}}$ form a basis of μ -almost decidable sets, as shown to be computable by Lemma 2.15. Consider the c.e. set

$$F_k := \{(i, j) : \delta_S(d_i, d_{m_k}) + r_j < q_{l_k}\}. \quad (5)$$

Because $\{d_i\}_{i \in \mathbb{N}}$ is dense in S and $\{r_j\}_{j \in \mathbb{N}}$ is dense in the positive reals we have for each $k \in \mathbb{N}$ that $B_k = \bigcup_{(i, j) \in F_k} B(d_i, r_j)$. In particular this implies that the set $F := \bigcup_{k \in E} F_k$ is a c.e. set with $V = \bigcup_{(i, j) \in F} B(d_i, r_j)$. Let $\{(i_n, j_n)\}_{n \in \mathbb{N}}$ be a c.e. enumeration of F and let $V_k := \bigcup_{n \leq k} B(d_{i_n}, r_{j_n})$, which is almost decidable. By construction, for each k , $V_k \subseteq V_{k+1}$, and $\bigcup_{k \in \mathbb{N}} V_k = V$. \square

Using the notion of an almost decidable set, we have the following characterization of computable measures.

Corollary 2.17. *Let S be a computable metric space and let $\mu \in \mathcal{M}_1(S)$ be a probability measure on S . Then μ is computable if the measure $\mu(A)$ of every μ -almost decidable set A is a computable real, uniformly in A .*

Proof. Let V be a c.e. open set of S . By Theorem 2.11, it suffices to show that $\mu(V)$ is a c.e. real, uniformly in V . By Lemma 2.16, we can compute a nested sequence $\{V_k\}_{k \in \mathbb{N}}$ of μ -almost decidable sets whose union is V . Because V is open, $\mu(V) = \sup_{k \in \mathbb{N}} \mu(V_k)$. By hypothesis, $\mu(V_k)$ is a computable real for each k , and so the supremum is a c.e. real, as desired. \square

Computable random variables have computable distributions.

Proposition 2.18 ([GHR10, Prop. 2.4.2]). *Let X be a computable random variable in a computable metric space S . Then its distribution is a computable point in the computable metric space $\mathcal{M}_1(S)$.*

On the other hand, one can show that given a computable point μ in $\mathcal{M}_1(S)$, one can construct an i.i.d.- μ sequence of computable random variables in S .

3. CONDITIONAL PROBABILITIES AND DISTRIBUTIONS

Informally, the *conditional probability* of an event B given an event A is the likelihood that the event B occurs, given the knowledge that the event A has occurred.

Definition 3.1 (Conditioning with respect to a single event). Let S be a measurable space and let $\mu \in \mathcal{M}_1(S)$ be a probability measure on S . Let $A, B \subseteq S$ be measurable sets, and suppose that $\mu(A) > 0$. Then the **conditional probability of B given A** , written $\mu(B|A)$, is defined by

$$\mu(B|A) = \frac{\mu(B \cap A)}{\mu(A)}. \quad (6)$$

Note that for any fixed measurable set $A \subseteq S$ with $\mu(A) > 0$, the function $\mu(\cdot|A)$ is a probability measure. This notion of conditioning is well-defined precisely when $\mu(A) > 0$, and so is insufficient for defining the conditional probability given the event that a *continuous* random variable takes a particular value, as such an event has measure zero.

We will often be interested in the case where B and A are measurable sets of the form $\{Y \in D\}$ and $\{X \in C\}$. In this case, we define the abbreviation

$$\mathbf{P}\{Y \in D \mid X \in C\} := \mathbf{P}(\{Y \in D\} \mid \{X \in C\}). \quad (7)$$

Again, this is well-defined when $\mathbf{P}\{X \in C\} > 0$. As a special case, when $C = \{x\}$ is an atom, we obtain the notation

$$\mathbf{P}\{Y \in D \mid X = x\}. \quad (8)$$

The modern formulation of conditional probability is due to Kolmogorov [Kol33], and gives a consistent solution to the problem of conditioning on the value of general (and in particular, continuous) random variables. (See Kallenberg [Kal02, Chp. 6] for a rigorous treatment.)

Definition 3.2 (Conditioning with respect to a random variable). Let X and Y be random variables in measurable spaces S and T , respectively, and let $B \subseteq T$ be a measurable set. Then a measurable function $\mathbf{P}[Y \in B|X]$ from S to $[0, 1]$ is (a version of) the **conditional probability of $Y \in B$ given X** when

$$\mathbf{P}\{X \in A, Y \in B\} = \int_A \mathbf{P}[Y \in B|X] d\mathbf{P}_X \quad (9)$$

for all measurable sets $A \subseteq S$. Moreover, $\mathbf{P}[Y \in B|X]$ is uniquely defined up to a \mathbf{P}_X -null set (i.e., almost surely unique) and so we may sensibly refer to *the* conditional probability when we mean a generic element of the equivalence class of versions.

Remark 3.3. It is also common to define a conditional probability given \mathbf{X} as a $\sigma(\mathbf{X})$ -measurable random variable \mathbf{P} that is the Radon-Nikodym derivative of the measure $\mathbf{P}(\cdot \cap \{\mathbf{Y} \in B\})$ with respect to \mathbf{P} , both viewed as measures on the σ -algebra $\sigma(\mathbf{X})$. (See Kallenberg [Kal02, Chp. 6] for such a construction.) However there is a close relationship between these two definitions: In particular, there exists a measure function h from S to $[0, 1]$ such that $\mathbf{P} = h(\mathbf{X})$ a.s. We have taken h to be our definition of the conditional probability as it is more natural to have a function from S in the statistical setting. (We take the same tact when defining conditional distributions.) In Propositions 6.5 and 6.10, we demonstrate that both definitions of conditional probability admit noncomputable versions.

It is natural to want to consider not just individual conditional probabilities but the entire conditional distribution $\mathbf{P}[\mathbf{Y} \in \cdot | \mathbf{X}]$ of \mathbf{Y} given \mathbf{X} . In order to define conditional distributions, we first recall the notion of a probability kernel. (For more details, see, e.g., [Kal02, Ch. 3, 6].)

Definition 3.4 (Probability kernel). Let S and T be measurable spaces. A function $\kappa : S \times \mathcal{B}_T \rightarrow [0, 1]$ is called a **probability kernel (from S to T)** when

- (1) for every $s \in S$, the function $\kappa(s, \cdot)$ is a probability measure on T ; and
- (2) for every $B \in \mathcal{B}_T$, the function $\kappa(\cdot, B)$ is measurable.

It can be shown that κ is a probability kernel from S to T if and only if $s \mapsto \kappa(s, \cdot)$ is a measurable map from S to $\mathcal{M}_1(T)$ [Kal02, Lem. 1.40].

Definition 3.5 (Conditional distribution). Let \mathbf{X} and \mathbf{Y} be random variables in measurable spaces S and T , respectively. A probability kernel κ is called a **(regular) version of the conditional distribution $\mathbf{P}[\mathbf{Y} \in \cdot | \mathbf{X}]$** when it satisfies

$$\mathbf{P}\{\mathbf{X} \in A, \mathbf{Y} \in B\} = \int_A \kappa(x, B) \mathbf{P}_{\mathbf{X}}(dx), \quad (10)$$

for all measurable sets $A \subseteq S$ and $B \subseteq T$.

We will simply write $\mathbf{P}[\mathbf{Y} | \mathbf{X}]$ in place of $\mathbf{P}[\mathbf{Y} \in \cdot | \mathbf{X}]$.

Definition 3.6. Let μ be a measure on a topological space S with open sets \mathcal{S} . Then the **support of μ** , written $\text{supp}(\mu)$, is defined to be the set of points $x \in S$ such that all open neighborhoods of x have positive measure, i.e.,

$$\text{supp}(\mu) := \{x \in S : \forall B \in \mathcal{S} (x \in B \implies \mu(B) > 0)\}. \quad (11)$$

Given any two versions of a conditional distribution, they need only agree almost everywhere. However, they will agree at points of continuity in the support:

Lemma 3.7. *Let \mathbf{X} and \mathbf{Y} be random variables in topological spaces S and T , respectively, and suppose that κ_1, κ_2 are versions of the conditional distribution $\mathbf{P}[\mathbf{Y} | \mathbf{X}]$. Let $x \in S$ be a point of continuity of both of the maps $x \mapsto \kappa_i(x, \cdot)$ for $i = 1, 2$. If $x \in \text{supp}(\mathbf{P}_{\mathbf{X}})$, then $\kappa_1(x, \cdot) = \kappa_2(x, \cdot)$.*

Proof. Fix a measurable set $A \subseteq Y$ and define $g(\cdot) := \kappa_1(\cdot, A) - \kappa_2(\cdot, A)$. We know that $g = 0$ $\mathbf{P}_{\mathbf{X}}$ -a.e., and also that g is continuous at x . Assume, for the purpose of

contradiction, that $g(x) = \varepsilon > 0$. By continuity, there is an open neighborhood B of x , such that $g(B) \in (\frac{\varepsilon}{2}, \frac{3\varepsilon}{2})$. But $x \in \text{supp}(\mathbf{P}_X)$, hence $\mathbf{P}_X(B) > 0$, contradicting $g = 0$ \mathbf{P}_X -a.e. \square

When conditioning on a discrete random variable, i.e. one whose image is a discrete set, it is well known that a version of the conditional distribution can be built by elementary conditioning with respect to single events.

Lemma 3.8. *Let X and Y be random variables on measurable spaces S and T , respectively. Suppose that X is a discrete random variable with support $R \subseteq S$, and let ν be an arbitrary probability measure on T . Define the function $\kappa : S \times \mathcal{B}_T \rightarrow [0, 1]$ by*

$$\kappa(x, B) := \mathbf{P}\{Y \in B \mid X = x\} \quad (12)$$

for all $x \in R$ and $\kappa(x, \cdot) = \nu(\cdot)$ for $x \notin R$. Then κ is a version of the conditional distribution $\mathbf{P}\{Y \mid X\}$.

Proof. The function κ , given by

$$\kappa(x, B) := \mathbf{P}\{Y \in B \mid X = x\} \quad (13)$$

for all $x \in R$ and $\kappa(x, \cdot) = \nu(\cdot)$ for $x \notin R$, is well-defined because $\mathbf{P}\{X = x\} > 0$ for all $x \in R$, and so the right hand side of Equation (13) is well-defined. Furthermore, $\mathbf{P}\{X \in R\} = 1$ and so κ is characterized by Equation (13) for \mathbf{P}_X -almost all x . Finally, for all measurable sets $A \subseteq S$ and $B \subseteq T$, we have

$$\int_A \kappa(x, B) \mathbf{P}_X(dx) = \sum_{x \in R \cap A} \mathbf{P}\{Y \in B \mid X = x\} \mathbf{P}\{X = x\} \quad (14)$$

$$= \sum_{x \in R \cap A} \mathbf{P}\{Y \in B, X = x\} \quad (15)$$

$$= \mathbf{P}\{Y \in B, X \in A\}, \quad (16)$$

and so κ is a version of the conditional distribution $\mathbf{P}\{Y \mid X\}$. \square

4. COMPUTABLE CONDITIONAL PROBABILITIES AND DISTRIBUTIONS

We begin by demonstrating the computability of elementary conditional probability given positive-measure events that are almost decidable. We then return to the abstract setting and lay the foundations for the remainder of the paper.

Lemma 4.1 ([GHR10, Prop. 3.1.2]). *Let (S, μ) be a computable probability space and let A be an almost decidable subset of S satisfying $\mu(A) > 0$. Then $\mu(\cdot \mid A)$ is a computable probability measure.*

Proof. By Corollary 2.17, it suffices to show that $\frac{\mu(B \cap A)}{\mu(A)}$ is computable for an almost decidable set B . But then $B \cap A$ is almost decidable and so its measure, the numerator, is a computable real. The denominator is likewise the measure of an almost decidable set, hence a computable real. Finally, the ratio of two computable reals is computable. \square

Conditioning in the abstract setting is more involved. Having defined the abstract notion of conditional probabilities and conditional distributions in Section 3, we now define notions of computability for these objects, starting with conditional distributions.

Definition 4.2 (Computable probability kernel). Let S and T be computable metric spaces and let $\kappa : S \times \mathcal{B}_T \rightarrow [0, 1]$ be a probability kernel from S to T . Then we say that κ is a **computable (probability) kernel** when the map $\phi_\kappa : S \rightarrow \mathcal{M}_1(T)$ given by $\phi_\kappa(s) := \kappa(s, \cdot)$ is a computable function. Similarly, we say that κ is computable on a subset $D \subseteq S$ when ϕ_κ is computable on D .

The following correspondence will allow us to derive other characterizations of computable probability kernels. The proof, however, will use that fact that given a sequence of ideal balls $B(d_1, q_1), \dots, B(d_n, q_n)$ and an ideal ball $B(d^*, q^*)$ we can semi-decide when $B(d^*, q^*) \subseteq \bigcup_{i \leq n} B(d_i, q_i)$ (uniformly in the indexes of the ideal balls).

Henceforth, we will make the assumption that our computable metric spaces have the property. This holds for all the specific spaces $(\mathbb{R}^k, \{0, 1\}, \{0, 1\}^\infty, \mathbb{N}, \text{etc.})$ that we consider.

Proposition 4.3. *Let (T, δ, \mathcal{D}) be a computable metric space. Let \mathcal{T} be the collection of sets of the form*

$$P_{A,q} = \{\mu \in \mathcal{M}_1(T) : \mu(A) > q\} \quad (17)$$

where A is a c.e. open subset of T and q is a rational. Then the elements of \mathcal{T} are c.e. open subsets of $\mathcal{M}_1(T)$ uniformly in A and q .

Proof. Note that \mathcal{T} is a subbasis for the weak topology induced by the Prokhorov metric (see [Sch07a, Lem. 3.2]).

Let $P = P_{A,q}$ for a rational q and c.e. open subset A .

We can write $A = \bigcup_{n \in \mathbb{N}} B(d_n, r_n)$ for a sequence (computable uniformly in A) of ideal balls in T with centers $d_n \in \mathcal{D}_T$ and rational radii r_n . Define $A_m := \bigcup_{n \leq m} B(d_n, r_n)$. Then $A_m \subseteq A_{m+1}$ and $A = \bigcup_m A_m$. Writing

$$P_m := \{\mu \in \mathcal{M}_1(T) : \mu(A_m) > q\}, \quad (18)$$

we have $P = \bigcup_m P_m$. In order to show that P is c.e. open uniformly in q and A , it suffices to show that P_m is c.e. open, uniformly in q , m and A .

Let \mathcal{D}_P be the ideal points of $\mathcal{M}_1(T)$ (see Section 2.4), let $\nu \in \mathcal{D}_P$, and let R be the finite set on which it concentrates. Gács [Gács05, Prop. B.17] characterizes the ideal ball E centered at ν

$$\mu(C^\varepsilon) > \nu(C) - \varepsilon \quad (19)$$

for all subsets $C \subseteq R$, where $C^\varepsilon = \bigcup_{t \in C} B(t, \varepsilon)$.

It is straightforward to show that $E \subseteq P_m$ if and only if $\nu(C_m) \geq q + \varepsilon$, where

$$C_m := \{t \in R : B(t, \varepsilon) \subseteq A_m\}. \quad (20)$$

Note that C_m is a decidable subset of R (uniformly in m , A , and E) and that $\nu(C_m)$ is a rational and so we can decide whether $E \subseteq P_m$, showing that P_m is c.e. open, uniformly in q , m , and A . \square

Recall that a lower semicomputable function from a computable metric space to $[0, 1]$ is one for which the preimage of $(q, 1]$ is c.e. open, uniformly in rationals q . Furthermore, we say that a function f from a computable metric space S to $[0, 1]$ is *lower semicomputable on $D \subseteq S$* when there is a uniformly computable sequence $\{U_q\}_{q \in \mathbb{Q}}$ of c.e. open sets such that

$$f^{-1}[(q, 1]] \cap D = U_q \cap D. \quad (21)$$

We can also interpret a computable probability kernel κ as a computable map sending each c.e. open set $A \subseteq T$ to a lower semicomputable function $\kappa(\cdot, A)$.

Lemma 4.4. *Let S and T be computable metric spaces, let κ be a probability kernel from S to T , and let $D \subseteq S$. If ϕ_κ is computable on D then $\kappa(\cdot, A)$ is lower semicomputable on D uniformly in the c.e. open set A .*

Proof. Let $q \in (0, 1)$ be a rational, and let A be a c.e. open set. Define $I := (q, 1]$. Then $\kappa(\cdot, A)^{-1}[I] = \phi_\kappa^{-1}[P]$, where

$$P := \{\mu \in \mathcal{M}_1(T) : \mu(A) > q\}. \quad (22)$$

This is an open set in the weak topology induced by the Prokhorov metric (see [Sch07a, Lem. 3.2]), and by Lemma 4.3, P is c.e. open.

By the computability of ϕ_κ , there is a c.e. open set V , uniformly computable in q and A such that

$$\kappa(\cdot, A)^{-1}[I] \cap D = \phi_\kappa^{-1}[P] \cap D = V \cap D, \quad (23)$$

and so $\kappa(\cdot, A)$ is lower semicomputable on D , uniformly in A . \square

In fact, when $A \subseteq T$ is a decidable set (i.e., when A and $T \setminus A$ are both c.e. open), $\kappa(\cdot, A)$ is a computable function.

Corollary 4.5. *Let S and T be computable metric spaces, let κ be a probability kernel from S to T computable on a subset $D \subseteq S$, and let $A \subseteq T$ be a decidable set. Then $\kappa(\cdot, A) : S \rightarrow [0, 1]$ is computable on D .*

Proof. By Lemma 4.4, $\kappa(\cdot, A)$ and $\kappa(\cdot, T \setminus A)$ are lower semicomputable on D . But $\kappa(x, A) = 1 - \kappa(x, T \setminus A)$ for all $x \in D$, and so $\kappa(\cdot, A)$ is upper semicomputable, and therefore computable, on D . \square

Although a conditional distribution may have many different versions, their computability as probability kernels does not differ (up to a change in domain by a null set).

Lemma 4.6. *Let X and Y be computable random variables on computable metric spaces S and T , respectively, and let κ be a version of a conditional distribution $\mathbf{P}[Y|X]$ that is computable on some \mathbf{P}_X -measure one set. Then any version of $\mathbf{P}[Y|X]$ is also computable on some \mathbf{P}_X -measure one set.*

Proof. Let κ be a version that is computable on a \mathbf{P}_X -measure one set D , and let κ' be any other version. Then $Z := \{s \in S : \kappa(s, \cdot) \neq \kappa'(s, \cdot)\}$ is a \mathbf{P}_X -null set, and $\kappa = \kappa'$ on $D \setminus Z$. Hence κ' is computable on the \mathbf{P}_X -measure one set $D \setminus Z$. \square

This observation motivates the following definition of computability for conditional distributions.

Definition 4.7 (Computable conditional distributions). Let X and Y be computable random variables on computable metric spaces S and T , respectively, and let κ be a version of the conditional distribution $\mathbf{P}[Y|X]$. We say that $\mathbf{P}[Y|X]$ is computable when κ is computable on a \mathbf{P}_X -measure one subset of S .

Note that this definition is analogous to our Definition 2.9 of a *computable random variable*. In fact, if κ is a version of a computable conditional distribution $\mathbf{P}[Y|X]$, then $\kappa(X, \cdot)$ is a ($\sigma(X)$ -measurable) computable random probability measure (i.e., a probability-measure-valued random variable).

Intuitively, a conditional distribution is computable when for some (and hence for any) version κ there is a program that, given as input a representation of a point $s \in S$, outputs a representation of the measure $\phi_\kappa(s) = \kappa(s, \cdot)$ for \mathbf{P}_X -almost all inputs s .

Suppose that $\mathbf{P}[Y|X]$ is computable, i.e., there is a version κ for which the map ϕ_κ is computable on some \mathbf{P}_X -measure one set $S' \subseteq S$. As noted in Definition 4.2, we will often abuse notation and say that κ is computable on S' . The restriction of ϕ_κ to S' is necessarily continuous (under the subspace topology on S'). We say that κ is **\mathbf{P}_X -almost continuous** when the restriction of ϕ_κ to some \mathbf{P}_X -measure one set is continuous. Thus when $\mathbf{P}[Y|X]$ is computable, there is some \mathbf{P}_X -almost continuous version.

We will need the following lemma in the proof of Lemma 6.3, but we postpone the proof until Section 8.2.

Lemma 4.8. *Let X and Y be random variables on metric spaces S and T , respectively, and let $R \subseteq S$. If a conditional density $p_{X|Y}(x|y)$ of X given Y is continuous on $R \times T$, positive, and bounded, then there is a version of the conditional distribution $\mathbf{P}[Y|X]$ that is continuous on R . In particular, if R is a \mathbf{P}_X -measure one subset, then there is a \mathbf{P}_X -almost continuous version.*

We now define computability for conditional probability.

Definition 4.9 (Computable conditional probability). Let X and Y be computable random variables in computable metric spaces S and T , respectively, and let $B \subseteq T$ be a \mathbf{P}_Y -almost decidable set. We say that the conditional probability $\mathbf{P}[Y \in B|X]$ is computable when it is computable (when viewed as a function from S to $[0, 1]$) on a \mathbf{P}_X -measure one set.

In Section 5 we describe a pair of computable random variables X, Y for which $\mathbf{P}[Y|X]$ is not computable, by virtue of no version being \mathbf{P}_X -almost continuous. In Section 6 we describe a pair of computable random variables X, Y for which there is a \mathbf{P}_X -almost continuous version of $\mathbf{P}[Y|X]$, but still no version that is computable on a \mathbf{P}_X -measure one set.

5. DISCONTINUOUS CONDITIONAL DISTRIBUTIONS

Our study of the computability of conditional distributions begins with a description of the following roadblock: a conditional distribution need not have *any* version that is continuous or even almost continuous (in the sense described in Section 4). This rules out computability.

We will work with the standard effective presentations of the spaces \mathbb{R} , \mathbb{N} , $\{0, 1\}$, as well as product spaces thereof, as computable metric spaces. For example, we will use \mathbb{R} under the Euclidean metric, along with the ideal points \mathbb{Q} under their standard enumeration.

Recall that a random variable C is a **Bernoulli**(p) random variable, or equivalently, a p -**coin**, when $\mathbf{P}\{C = 1\} = 1 - \mathbf{P}\{C = 0\} = p$. We call a $\frac{1}{2}$ -coin a **fair coin**. A random variable N is **geometric** when it takes values in $\mathbb{N} = \{0, 1, 2, \dots\}$ and satisfies

$$\mathbf{P}\{N = n\} = 2^{-(n+1)}, \quad \text{for } n \in \mathbb{N}. \quad (24)$$

A random variable that takes values in a discrete set is a **uniform** random variable when it assigns equal probability to each element. A continuous random variable U on the unit interval is **uniform** when the probability that it falls in the subinterval $[\ell, r]$ is $r - \ell$. It is easy to show that the distributions of these random variables are computable.

Let C , U , and N be independent computable random variables, where C is a fair coin, U is a uniform random variable on $[0, 1]$, and N is a geometric random variable. Fix a computable enumeration $\{r_i\}_{i \in \mathbb{N}}$ of the rational numbers (without repetition) in $(0, 1)$, and consider the random variable

$$X := \begin{cases} U, & \text{if } C = 1; \\ r_N, & \text{otherwise.} \end{cases} \quad (25)$$

It is easy to verify that X is a computable random variable.

Proposition 5.1. *Every version of the conditional distribution $\mathbf{P}[C|X]$ is discontinuous everywhere on every \mathbf{P}_X -measure one set. In particular, every version of the conditional probability $\mathbf{P}\{C = 0|X\}$ is discontinuous everywhere on every \mathbf{P}_X -measure one set.*

Proof. Note that $\mathbf{P}\{X \text{ rational}\} = \frac{1}{2}$ and, furthermore, $\mathbf{P}\{X = r_k\} = \frac{1}{2^{k+2}} > 0$. Therefore, any two versions of the conditional distribution $\mathbf{P}[C|X]$ must agree on *all* rationals in $[0, 1]$. In addition, because $\mathbf{P}_U \ll \mathbf{P}_X$, i.e.,

$$\mathbf{P}\{U \in A\} > 0 \implies \mathbf{P}\{X \in A\} > 0 \quad (26)$$

for all measurable sets $A \subseteq [0, 1]$, any two versions must agree on a Lebesgue-measure one set of the irrationals in $[0, 1]$. An elementary calculation shows that

$$\mathbf{P}\{C = 0 \mid X \text{ rational}\} = 1, \quad (27)$$

while

$$\mathbf{P}\{C = 0 \mid X \text{ irrational}\} = 0. \quad (28)$$

Therefore, all versions κ of $\mathbf{P}[\mathbf{C}|\mathbf{X}]$ satisfy, for $\mathbf{P}_\mathbf{X}$ -almost all x ,

$$\kappa(x, \{0\}) = \begin{cases} 1, & x \text{ rational;} \\ 0, & x \text{ irrational,} \end{cases} \quad (29)$$

where the right hand side, considered as a function of x , is the so-called Dirichlet function, a *nowhere continuous* function.

Suppose some version of the conditional probability $x \mapsto \kappa(x, \{0\})$ were continuous at a point y on a $\mathbf{P}_\mathbf{X}$ -measure one set R . Then there would exist an open interval I containing y such that the image of $I \cap R$ contains 0 or 1, but not both. However, R must contain all rationals in I and almost all irrationals in I . Furthermore, the image of every rational in $I \cap R$ is 1, and the image of almost every irrational in $I \cap R$ is 0, a contradiction. \square

Discontinuity is a fundamental obstacle to computability, but many conditional probabilities do admit continuous versions, and we can focus our attention on such settings, to rule out this objection. In particular, we might hope to be able to compute a conditional distribution of a pair of computable random variables when there is *some* version that is almost continuous or even continuous. However we will show that even this is not possible in general.

6. NONCOMPUTABLE ALMOST CONTINUOUS CONDITIONAL DISTRIBUTIONS

In this section, we construct a pair of random variables (\mathbf{X}, \mathbf{N}) that is computable, yet whose conditional distribution $\mathbf{P}[\mathbf{N}|\mathbf{X}]$ is not computable, despite the existence of a $\mathbf{P}_\mathbf{X}$ -almost continuous version.

Let M_n denote the n th Turing machine, under a standard enumeration, and let $h : \mathbb{N} \rightarrow \mathbb{N} \cup \{\infty\}$ be the map given by $h(n) = \infty$ if M_n does not halt (on input 0) and $h(n) = k$ if M_n halts (on input 0) at the k th step. Let $H = \{n \in \mathbb{N} : h(n) < \infty\}$ be the indices of those Turing machines that halt (on input 0). We now use h to define a pair (\mathbf{N}, \mathbf{X}) of computable random variables such that H is computable from the conditional distribution of \mathbf{N} given \mathbf{X} .

Let \mathbf{N} be a computable geometric random variable, \mathbf{C} a computable $\frac{1}{3}$ -coin, and \mathbf{U} and \mathbf{V} both computable uniform random variables on $[0, 1]$, all mutually independent. Let $\lfloor x \rfloor$ denote the greatest integer $y \leq x$. Note that $\lfloor 2^k \mathbf{V} \rfloor$ is uniformly distributed on $\{0, 1, 2, \dots, 2^k - 1\}$. Consider the derived random variables

$$\mathbf{X}_k := \frac{2\lfloor 2^k \mathbf{V} \rfloor + \mathbf{C} + \mathbf{U}}{2^{k+1}} \quad (30)$$

for $k \in \mathbb{N}$. Almost surely, the limit $\mathbf{X}_\infty := \lim_{k \rightarrow \infty} \mathbf{X}_k$ exists and satisfies $\mathbf{X}_\infty = \mathbf{V}$ a.s. Finally, we define $\mathbf{X} := \mathbf{X}_{h(\mathbf{N})}$.

Proposition 6.1. *The random variable \mathbf{X} is computable.*

Proof. Because \mathbf{U} and \mathbf{V} are computable and a.s. nondyadic, their (a.s. unique) binary expansions $\{\mathbf{U}_n : n \in \mathbb{N}\}$ and $\{\mathbf{V}_n : n \in \mathbb{N}\}$ are themselves computable random variables in $\{0, 1\}$, uniformly in n .

For each $k \geq 0$, define the random variable

$$D_k = \begin{cases} V_k, & h(\mathbf{N}) > k; \\ C, & h(\mathbf{N}) = k; \\ U_{k-h(\mathbf{N})-1}, & h(\mathbf{N}) < k. \end{cases} \quad (31)$$

By simulating $M_{\mathbf{N}}$ for k steps, we can decide whether $h(\mathbf{N})$ is less than, equal to, or greater than k . Therefore the random variables $\{D_k\}_{k \geq 0}$ are computable, uniformly in k . We now show that, with probability one, $\{D_k\}_{k \geq 0}$ is the binary expansion of X , thus showing that X is itself a computable random variable.

There are two cases to consider:

First, conditioned on $h(\mathbf{N}) = \infty$, we have that $D_k = V_k$ for all $k \geq 0$. In fact, $X = V$ when $h(\mathbf{N}) = \infty$, and so the binary expansions match.

Condition on $h(\mathbf{N}) = m$ and let D denote the computable random real whose binary expansion is $\{D_k\}_{k \geq 0}$. We must then show that $D = X_m$ a.s. Note that

$$\lfloor 2^m X_m \rfloor = \lfloor 2^m V \rfloor = \sum_{k=0}^{m-1} 2^{m-1-k} V_k = \lfloor 2^m D \rfloor, \quad (32)$$

and thus the binary expansions agree for the first m digits. Finally, notice that $2^{m+1} X_m - 2 \lfloor 2^m X_m \rfloor = C + U$, and so the next binary digit of X_m is C , followed by the binary expansion of U , thus agreeing with D for all $k \geq 0$. \square

We now show that $\mathbf{P}[\mathbf{N}|X]$ is not computable, despite the existence of a \mathbf{P}_X -almost continuous version of $\mathbf{P}[\mathbf{N}|X]$. We begin by characterizing the conditional density of X given \mathbf{N} .

Lemma 6.2. *For each $k \in \mathbb{N} \cup \{\infty\}$, the distribution of X_k admits a density p_{X_k} with respect to Lebesgue measure on $[0, 1]$, where, for all $k < \infty$ and \mathbf{P}_X -almost all x ,*

$$p_{X_k}(x) = \begin{cases} \frac{4}{3}, & \lfloor 2^{k+1} x \rfloor \text{ even;} \\ \frac{2}{3}, & \lfloor 2^{k+1} x \rfloor \text{ odd,} \end{cases} \quad (33)$$

and $p_{X_\infty}(x) = 1$.

Proof. We have $X_\infty = V$ a.s. and so the constant function taking the value 1 is the density of X_∞ with respect to Lebesgue measure on $[0, 1]$.

Let $k \in \mathbb{N}$. With probability one, the integer part of $2^{k+1} X_k$ is $2 \lfloor 2^k V \rfloor + C$ while the fractional part is U . Therefore, the distribution of $2^{k+1} X_k$ (and hence X_k) admits a piecewise constant density with respect to Lebesgue measure.

In particular, $\lfloor 2^{k+1} X_k \rfloor \equiv C \pmod{2}$ almost surely and $2 \lfloor 2^k V \rfloor$ is independent of C and uniformly distributed on $\{0, 2, \dots, 2^{k+1} - 2\}$. Therefore,

$$\mathbf{P}\{\lfloor 2^{k+1} X_k \rfloor = \ell\} = 2^{-k} \cdot \begin{cases} \frac{2}{3}, & \ell \text{ even;} \\ \frac{1}{3}, & \ell \text{ odd,} \end{cases} \quad (34)$$




FIGURE 1. A visualization of (X, Y) , where Y is uniformly distributed and $N = \lfloor -\log_2 Y \rfloor$. Regions that appear (at low resolution) to be uniform can suddenly be revealed (at higher resolutions) to be patterned. Deciding whether the pattern is in fact uniform (or below the resolution of this printer/display) is tantamount to solving the halting problem, but it is possible to sample from this distribution nonetheless. Note that this is not a plot of the density, but instead a plot where the darkness of each pixel is proportional to its measure.

for every $\ell \in \{0, 1, \dots, 2^{k+1} - 1\}$. It follows immediately that the density p of $2^{k+1}X_k$ with respect to Lebesgue measure on $[0, 2^{k+1}]$ is given by

$$p(x) = 2^{-k} \cdot \begin{cases} \frac{2}{3}, & [x] \text{ even;} \\ \frac{1}{3}, & [x] \text{ odd.} \end{cases} \quad (35)$$

and so the density of X_k is obtained by rescaling: $p_{X_k}(x) = 2^{k+1} \cdot p(2^{k+1}x)$. \square

As X_k admits a density with respect to Lebesgue measure on $[0, 1]$ for all $k \in \mathbb{N} \cup \{\infty\}$, it follows that the conditional distribution of X given N admits a conditional density $p_{X|N}$ (with respect to Lebesgue measure on $[0, 1]$) given by

$$p_{X|N}(x|n) = p_{X_{h(n)}}(x). \quad (36)$$

Each of these densities is positive, bounded, and continuous on the nondyadic reals, and so they can be combined to form a \mathbf{P}_X -almost continuous version of the conditional distribution.

Lemma 6.3. *There is a \mathbf{P}_X -almost continuous version of $\mathbf{P}[N|X]$.*

Proof. By Bayes' rule (Lemma 8.6), the probability kernel κ given by

$$\kappa(x, B) := \frac{\sum_{n \in B} p_{\mathbf{X}|\mathbf{N}}(x|n) \cdot \mathbf{P}\{\mathbf{N} = n\}}{\sum_{n \in \mathbb{N}} p_{\mathbf{X}|\mathbf{N}}(x|n) \cdot \mathbf{P}\{\mathbf{N} = n\}} \quad \text{for } B \subseteq \mathbb{N} \quad (37)$$

is a version of the conditional distribution $\mathbf{P}[\mathbf{N}|\mathbf{X}]$. Moreover, $p_{\mathbf{X}|\mathbf{N}}$ is positive, bounded and continuous on $R \times \mathbb{N}$, where R is the $\mathbf{P}_{\mathbf{X}}$ -measure one set of nondyadic reals in the unit interval. It follows from Lemma 4.8 that κ is $\mathbf{P}_{\mathbf{X}}$ -almost continuous. \square

Lemma 6.4. *Let κ be a version of $\mathbf{P}[\mathbf{N}|\mathbf{X}]$. For all $m, n \in \mathbb{N}$ and $\mathbf{P}_{\mathbf{X}}$ -almost all x ,*

$$\tau_{m,n}(x) := 2^{m-n} \cdot \frac{\kappa(x, \{m\})}{\kappa(x, \{n\})} \in \begin{cases} \{\frac{1}{2}, 1, 2\}, & h(n), h(m) < \infty; \\ \{1\}, & h(n) = h(m) = \infty; \\ \{\frac{2}{3}, \frac{3}{4}, \frac{4}{3}, \frac{3}{2}\}, & \text{otherwise.} \end{cases}$$

Proof. Let $m, n \in \mathbb{N}$. By (37) and then (36), for $\mathbf{P}_{\mathbf{X}}$ -almost all x ,

$$\begin{aligned} \tau_{m,n}(x) &= 2^{m-n} \cdot \frac{p_{\mathbf{X}|\mathbf{N}}(x|m) \cdot \mathbf{P}\{\mathbf{N} = m\}}{p_{\mathbf{X}|\mathbf{N}}(x|n) \cdot \mathbf{P}\{\mathbf{N} = n\}} \\ &= \frac{p_{\mathbf{X}_{h(m)}}(x)}{p_{\mathbf{X}_{h(n)}}(x)}. \end{aligned}$$

By Lemma 6.3, we have $p_{\mathbf{X}_{\infty}}(x) = 1$ and $p_{\mathbf{X}_k}(x) \in \{\frac{2}{3}, \frac{4}{3}\}$ for all $k < \infty$ and $\mathbf{P}_{\mathbf{X}}$ -almost all x . The result follows by considering all possible combinations of values for each regime of values for $h(n)$ and $h(m)$. \square

Despite the computability of \mathbf{N} and \mathbf{X} , conditioning on \mathbf{X} leads to noncomputability.

Proposition 6.5. *For all k , the conditional probability $\mathbf{P}[\mathbf{N} = k|\mathbf{X}]$ is not lower semicomputable (hence not computable) on any measurable subset $R \subseteq [0, 1]$ where $\mathbf{P}_{\mathbf{X}}(R) > \frac{2}{3}$.*

Proof. Choose n to be the index of some Turing machine M_n that halts (on input 0), i.e., for which $h(n) < \infty$, let $m \in \mathbb{N}$, let κ be a version of the conditional distribution, let $\tau = \tau_{m,n}$ be as in Lemma 6.4, and let $V_{<\infty}$ and V_{∞} be disjoint c.e. open sets that contain $\{\frac{1}{2}, 1, 2\}$ and $\{\frac{2}{3}, \frac{3}{4}, \frac{4}{3}, \frac{3}{2}\}$, respectively. By Lemma 6.4, when $h(m) < \infty$, we have $\mathbf{P}_{\mathbf{X}}(\tau^{-1}[V_{<\infty}]) = 1$. Likewise, when $h(m) = \infty$, we have $\mathbf{P}_{\mathbf{X}}(\tau^{-1}[V_{\infty}]) = 1$.

Suppose that for some k that $\mathbf{P}[\mathbf{N} = k|\mathbf{X}]$ were lower semicomputable on some set R satisfying $\mathbf{P}_{\mathbf{X}}(R) > \frac{2}{3}$. Note that the numerator of Equation (37) is computable on the nondyadics, uniformly in $B = \{k\}$, and the denominator is independent of B . Therefore, because $\mathbf{P}[\mathbf{N} = k|\mathbf{X}] = \kappa(\cdot, \{k\})$ $\mathbf{P}_{\mathbf{X}}$ -a.e., it follows that $\mathbf{P}[\mathbf{N} = k|\mathbf{X}]$ is lower semicomputable on R , uniformly in k . But then, uniformly in k , we have that $\mathbf{P}[\mathbf{N} = k|\mathbf{X}]$ is computable on R because $\{k\}$ is a decidable subset of \mathbb{N} and

$$\mathbf{P}[\mathbf{N} = k|\mathbf{X}] = 1 - \sum_{j \neq k} \mathbf{P}[\mathbf{N} = j|\mathbf{X}] \quad \text{a.s.} \quad (38)$$

Therefore, τ is computable on R , uniformly in m .

It follows that there exist c.e. open sets $U_{<\infty}$ and U_∞ , uniformly in m , such that $\tau^{-1}[V_{<\infty}] \cap R = U_{<\infty} \cap R$ and thus $\mathbf{P}_X(U_{<\infty} \cap R) = \mathbf{P}_X(\tau^{-1}[V_{<\infty}] \cap R)$, and similarly for V_∞ and U_∞ .

If $h(m) < \infty$, then $\mathbf{P}_X(\tau^{-1}[V_{<\infty}]) = 1$ and thus

$$\mathbf{P}_X(U_{<\infty}) \geq \mathbf{P}_X(U_{<\infty} \cap R) = \mathbf{P}_X(\tau^{-1}[V_{<\infty}] \cap R) = \mathbf{P}_X(R) > \frac{2}{3}. \quad (39)$$

Similarly, $h(m) = \infty$ implies that $\mathbf{P}_X(U_\infty) > \frac{2}{3}$. Therefore, at least one of

$$\mathbf{P}_X(U_{<\infty}) > \frac{2}{3} \quad \text{or} \quad \mathbf{P}_X(U_\infty) > \frac{2}{3} \quad (40)$$

holds.

As $V_{<\infty} \cap V_\infty = \emptyset$, we have $U_{<\infty} \cap U_\infty \cap R = \emptyset$ and so $\mathbf{P}_X(U_{<\infty} \cap U_\infty) < 1 - \frac{2}{3}$. Thus, $\mathbf{P}_X(U_{<\infty}) + \mathbf{P}_X(U_\infty) \leq 1 + \mathbf{P}_X(U_{<\infty} \cap U_\infty) < 1 + (1 - \frac{2}{3}) = \frac{4}{3}$ and so we have that at most one (and hence exactly one) of (i) $\mathbf{P}_X(U_{<\infty}) > \frac{2}{3}$ and $h(m) < \infty$ or (ii) $\mathbf{P}_X(U_\infty) > \frac{2}{3}$ and $h(m) = \infty$.

But $\mathbf{P}_X(U_{<\infty})$ and $\mathbf{P}_X(U_\infty)$ are c.e. reals, uniformly in m , and so we can computably distinguish between cases (i) and (ii), and thus decide whether or not $h(m) < \infty$, or equivalently, whether or not $m \in H$, uniformly in m . But H is not computable and so we have a contradiction. \square

Corollary 6.6. *The conditional distribution $\mathbf{P}[\mathbb{N}|X]$ is not computable.*

Proof. As $\{n\}$ is a decidable subset of \mathbb{N} , uniformly in n , it follows from Corollary 4.5 that the conditional probability $\mathbf{P}[\mathbb{N} = n|X]$ is computable on a measure one set when the conditional distribution is computable. By Proposition 6.5, the former is not computable, and so the latter is not computable. \square

Note that these proofs show that not only is the conditional distribution $\mathbf{P}[\mathbb{N}|X]$ noncomputable, but in fact it computes the halting set H in the following sense. Although we have not defined the notion of an oracle that encodes $\mathbf{P}[\mathbb{N}|X]$, one could make this concept precise using, e.g., infinite strings in the Type-2 Theory of Effectivity (TTE). However, despite not having a definition of computability from this distribution, we can easily relativize the notion of computability for the distribution. In particular, an analysis of the above proof shows that if $\mathbf{P}[\mathbb{N}|X]$ is A -computable for an oracle $A \subseteq \mathbb{N}$, then A computes the halting set, i.e., $A \geq_T \emptyset'$.

Computable operations map computable points to computable points, and so we obtain the following consequence.

Theorem 6.7. *The operation $(X, Y) \mapsto \mathbf{P}[Y|X]$ of conditioning a pair of real-valued random variables is not computable, even when restricted to pairs for which there exists a \mathbf{P}_X -almost continuous version of the conditional distribution.*

Conditional probabilities are often thought about as generic elements of $L^1(\mathbf{P}_X)$ equivalence classes (i.e., functions that are equivalent up to \mathbf{P}_X -null sets.) However, Proposition 6.5 also rules out the computability of conditional probabilities in the weaker sense of so-called $L^1(\mathbf{P}_X)$ -computability.

Definition 6.8 ($L^1(\mu)$ computability [PER84]). Let μ be a computable measure on $[0, 1]$. Then a function $f \in L^1(\mu)$ is said to be $L^1(\mu)$ -**computable** if there is a computable sequence of rational polynomials f_n such that

$$\|f - f_n\|_1 = \int |f - f_n| d\mu \leq 2^{-n}. \quad (41)$$

If we let $\delta_1(f, g) = \int |f - g| d\mu$ then we can turn $L^1(\mu)$ into a computable metric space by quotienting out by the equivalence relation $\delta_1(f, g) = 0$ and letting the dense \mathcal{D}_1 consist of those equivalence classes containing a polynomial with rational coefficients. A function f is then $L^1(\mu)$ -computable if and only if its equivalence class is a computable point in $L^1(\mu)$ considered as a computable metric space.

The following result is an easy consequence of the fact that $L^1(\mu)$ -computable functions are computable “in probability”.

Proposition 6.9. *For all $n \in \mathbb{N}$, the conditional probability $\mathbf{P}[\mathbf{N} = n|\mathbf{X}]$ is not $L^1(\mathbf{P}_X)$ -computable.*

Proof. Let $n \in \mathbb{N}$. By Proposition 6.5, the conditional probability $\mathbf{P}[\mathbf{N} = n|\mathbf{X}]$ is not computable on any Borel set R such that $\mathbf{P}_X(R) > \frac{2}{3}$. On the other hand, a result by Hoyrup and Rojas [HR09b, Thm. 3.1.1] implies that a function f is $L^1(\mathbf{P}_X)$ -computable only if f is computable on some \mathbf{P}_X -measure $1 - 2^{-r}$ set, uniformly in r . Therefore, $\mathbf{P}[\mathbf{N} = n|\mathbf{X}]$ is not $L^1(\mathbf{P}_X)$ -computable. \square

In fact, this result can be strengthened to a statement about the noncomputability of each conditional probability $\mathbf{P}[\mathbf{N} = n|\mathbf{X}]$ when viewed as $\sigma(\mathbf{X})$ -measurable random variables (Remark 3.3), or, equivalently, $\sigma(\mathbf{X})$ -measurable Radon-Nikodym derivatives. (See Kallenberg [Kal02, Chp. 6] for a discussion of the relationship between conditional expectation, Radon-Nikodym derivatives and conditional probability.) In this form, we tighten a result of Hoyrup and Rojas [HR11, Prop. 3].

Proposition 6.10. *Let $n \in \mathbb{N}$, and then consider the conditional probability $\mathbf{P}[\mathbf{N} = n|\mathbf{X}]$ as a $\sigma(\mathbf{X})$ -measurable Radon-Nikodym derivative $\eta_n : \{0, 1\}^\infty \rightarrow [0, 1]$ of $\mathbf{P}(\cdot \cap \{\mathbf{N} = n\})$ with respect to \mathbf{P} . Then η_n is not lower semicomputable (and hence not computable) on any measurable subset $S \subseteq \{0, 1\}^\infty$ where $\mathbf{P}(S) > \frac{2}{3}$. In particular, η_n is not $L^1(\mathbf{P})$ -computable.*

Proof. Let κ be a version of the conditional distribution $\mathbf{P}[\mathbf{N}|\mathbf{X}]$, and, for $n \in \mathbb{N}$, let η_n be defined as above. In particular, note that, for all $n \in \mathbb{N}$ and \mathbf{P} -almost all ω ,

$$\eta_n(\omega) = \kappa(\mathbf{X}(\omega), \{n\}). \quad (42)$$

In particular, this implies that, for all $n, m \in \mathbb{N}$ and \mathbf{P} -almost all ω ,

$$2^{m-n} \cdot \frac{\eta_m(\omega)}{\eta_n(\omega)} \in \begin{cases} \{\frac{1}{2}, 1, 2\}, & h(n), h(m) < \infty; \\ \{1\}, & h(n) = h(m) = \infty; \\ \{\frac{2}{3}, \frac{3}{4}, \frac{4}{3}, \frac{3}{2}\}, & \text{otherwise.} \end{cases} \quad (43)$$

Using Equations (37), (42), and (43), one can express η_n as a ratio whose numerator is computable on a \mathbf{P} -measure one set, uniformly in n , and whose denominator is

positive for \mathbf{P} -almost all ω and independent of n . One can then show, along the same lines as Proposition 6.5, that η_n is not lower semicomputable on any set S as defined above, for otherwise we could decide H . This then rules out the $L^1(\mathbf{P})$ -computability of η_n . \square

It is natural to ask whether this construction can be extended to produce a pair of computable random variables whose conditional distribution is noncomputable but has an *everywhere continuous* version. We provide such a strengthening in Section 7.

Despite these results, many important questions remain: How badly noncomputable is conditioning, even restricted to these continuous settings? In what restricted settings is conditioning *computable*? In Section 8, we begin to address the latter of these.

7. NONCOMPUTABLE EVERYWHERE CONTINUOUS CONDITIONAL DISTRIBUTIONS

As we saw in Section 5, discontinuity poses a fundamental obstacle to the computability of conditional probabilities. As such, it is natural to ask whether we can construct a pair of random variables (Z, N) that is computable and admits an *everywhere* continuous version of the conditional distribution $\mathbf{P}[N|Z]$, which is itself nonetheless not computable. In fact, this is possible using a construction similar to that of (X, N) in Section 6.

In particular, if we think of the construction of the k th bit of X as an iterative process, we see that there are two distinct stages. During the first stage, which occurs so long as $k < h(N)$, the bits of X simply mimic those of the uniform random variable V . Then during the second stage, once $k \geq h(N)$, the bits mimic that of $\frac{1}{2}(C + U)$.

Our construction of Z will differ in the second stage, where the bits of Z will instead mimic those of a random variable S specially designed to smooth out the rough edges caused by the biased coin C , while still allowing us to encode the halting set. In particular, S will be absolutely continuous and will have an infinitely differentiable density.

We now make the construction precise. Let N, U, V and C be as in the first construction. We begin by defining several random variables from which we will construct S , and then Z .

Lemma 7.1. *There is a random variable F in $[0, 1]$ with the following properties:*

- (1) F is computable.
- (2) \mathbf{P}_F admits a density p_F with respect to Lebesgue measure (on $[0, 1]$) that is infinitely differentiable everywhere.
- (3) $p_F(0) = \frac{2}{3}$ and $p_F(1) = \frac{4}{3}$.
- (4) $\frac{d_+^n}{dx^n} p_F(0) = \frac{d_-^n}{dx^n} p_F(1) = 0$, for all $n \geq 1$ (where $\frac{d_-^n}{dx^n}$ and $\frac{d_+^n}{dx^n}$ are the left and right derivatives respectively).

(See Figure 2 for one such random variable.) Let F be as in Lemma 7.1, and independent of all earlier random variables mentioned. Note that F is almost surely




FIGURE 2. (left) The graph of the function defined by $f(x) = \exp\{-\frac{1}{1-x^2}\}$, for $x \in (-1, 1)$, and 0 otherwise, a C^∞ bump function whose derivatives at ± 1 are all 0. (right) A density $p(y) = \frac{2}{3} \left(\frac{\Phi(2y-1)}{\Phi(1)} + 1 \right)$, for $y \in (0, 1)$, of a random variable satisfying Lemma 7.1, where $\Phi(y) = \int_{-1}^y \exp\{-\frac{1}{1-x^2}\} dx$ is the integral of the bump function.

nondyadic and so the r -th bit F_r of F is a computable random variable, uniformly in r .

Let D be a computable random variable, independent of all earlier random variables mentioned, and uniformly distributed on $\{0, 1, \dots, 7\}$. Consider

$$S = \frac{1}{8} \times \begin{cases} F, & \text{if } D = 0; \\ 4 + (1 - F), & \text{if } D = 4; \\ 4C + (D \bmod 4) + U, & \text{otherwise.} \end{cases} \quad (44)$$

It is clear that S is itself a computable random variable, and straightforward to show that

(i) \mathbf{P}_S admits an infinitely differentiable density p_S with respect to Lebesgue measure on $[0, 1]$; and

(ii) For all $n \geq 0$, we have $\frac{d^n}{dx^n} p_S(0) = \frac{d^n}{dx^n} p_S(1)$.

(For a visualization of the density p_S see Figure 3.)

We say a real $x \in [0, 1]$ is **valid for \mathbf{P}_S** if $x \in (\frac{1}{8}, \frac{1}{2}) \cup (\frac{5}{8}, 1)$. In particular, when $D \notin \{0, 4\}$, then S is valid for \mathbf{P}_S . The following are then straightforward consequences of the construction of S and the definition of valid points:

(iii) If x is valid for \mathbf{P}_S then $p_S(x) \in \{\frac{2}{3}, \frac{4}{3}\}$.

(iv) The \mathbf{P}_S -measure of points valid for \mathbf{P}_S is $\frac{3}{4}$.

Next we define, for every $k \in \mathbb{N}$, the random variables Z_k mimicking the construction of X_k . Specifically, for $k \in \mathbb{N}$, define

$$Z_k := \frac{\lfloor 2^k V \rfloor + S}{2^k}, \quad (45)$$




FIGURE 3. Graph of the density of S , when constructed from F as given in Figure 2.

and let $Z_\infty := \lim_{k \rightarrow \infty} Z_k = V$ a.s. Then the n th bit of Z_k is

$$(Z_k)_n = \begin{cases} V_n, & n < k; \\ S_{n-k}, & n \geq k \end{cases} \quad \text{a.s.} \quad (46)$$

For $k < \infty$, we say that $x \in [0, 1]$ is **valid for \mathbf{P}_{Z_k}** if the fractional part of $2^k x$ is valid for \mathbf{P}_S , and we say that x is **valid for \mathbf{P}_{Z_∞}** for all x . Let A_k be the collection of x valid for \mathbf{P}_{Z_k} , and let $A_\infty = [0, 1]$. It is straightforward to verify that $\mathbf{P}_Z(A_k) \geq \frac{9}{16}$ for all $k < \infty$, and 1 for $k = \infty$.

It is also straightforward to show from (i) and (ii) above that \mathbf{P}_{Z_k} admits an infinitely differentiable density p_{Z_k} with respect to Lebesgue measure on $[0, 1]$.

To complete the construction, we define $Z := Z_{h(\mathbb{N})}$. The following results are analogous to those in the almost continuous construction:

Lemma 7.2. *The random variable Z is computable.*

Lemma 7.3. *There is an everywhere continuous version of $\mathbf{P}[\mathbb{N}|Z]$.*

Proof. By construction, the conditional density of Z is everywhere continuous, bounded, and positive. The result follows from Lemma 4.8 for $R = [0, 1]$. \square

Lemma 7.4. *Let κ be a version of the conditional distribution $\mathbf{P}[\mathbb{N}|Z]$. For all $m, n \in \mathbb{N}$ and \mathbf{P}_Z -almost all x , if x is valid for $\mathbf{P}_{Z_{h(n)}}$ and $\mathbf{P}_{Z_{h(m)}}$ then*

$$\tau_{m,n}(x) := 2^{m-n} \cdot \frac{\kappa(x, \{m\})}{\kappa(x, \{n\})} \in \begin{cases} \{\frac{1}{2}, 1, 2\}, & h(n), h(m) < \infty; \\ \{1\}, & h(n) = h(m) = \infty; \\ \{\frac{2}{3}, \frac{3}{4}, \frac{4}{3}, \frac{3}{2}\}, & \text{otherwise.} \end{cases}$$

Proposition 7.5. *For all k , the conditional probability $\mathbf{P}[\mathbb{N} = k|Z]$ is not lower semicomputable (hence not computable) on any measurable subset $R \subseteq [0, 1]$ where $\mathbf{P}_Z(R) > \frac{23}{24}$.*

Proof. Choose n to be the index of a Turing machine M_n that does not halt (on input 0) so that $h(n) = \infty$, let $m \in \mathbb{N}$, let κ be a version of the conditional distribution, let $\tau = \tau_{m,n}$ be defined as in Lemma 7.4, and let $V_{<\infty}$ and V_∞ be disjoint c.e. open sets containing the points $\{\frac{2}{3}, \frac{3}{4}, \frac{4}{3}, \frac{3}{2}\}$ and $\{1\}$, respectively. Notice that all $x \in [0, 1]$ are valid for $\mathbf{P}_{Z_{h(n)}} = \mathbf{P}_{Z_\infty}$ and so $A_{h(n)} \cap A_{h(m)} = A_{h(m)}$. Therefore, by Lemma 7.4, when $h(m) < \infty$, we have $\mathbf{P}_Z(\tau^{-1}[V_{<\infty}]) \geq \mathbf{P}_Z(A_{h(m)}) \geq \frac{9}{16}$. On the other hand, when $h(m) = \infty$, we have $\mathbf{P}_Z(\tau^{-1}[V_\infty]) = \mathbf{P}_Z(A_\infty) = 1$. The remainder of the proof is similar to that of Proposition 6.5, replacing X with Z , replacing $\frac{2}{3}$ with $\frac{23}{24}$, and noting, e.g., that $\mathbf{P}_Z(\tau^{-1}[V_{<\infty}] \cap R) > \frac{9}{16} - \frac{1}{24}$ when $h(m) < \infty$. In particular, were $\mathbf{P}[N = k|Z]$ lower semicomputable on R , we could computably distinguish whether or not $m \in H$, uniformly in m , which is a contradiction. \square

As before, it follows immediately that $\mathbf{P}[N|Z]$ is not computable. It is possible to carry on the same development, showing the non-computability of the conditional probabilities as elements in $L^1(\mathbf{P}_X)$ and $L^1(\mathbf{P})$. For simplicity, we state the following strengthening of Theorem 6.7.

Theorem 7.6. *Let X and Y be computable real-valued random variables. Then operation $(X, Y) \mapsto \mathbf{P}[X|Y]$ of conditioning a pair of real-valued random variables is not computable, even when restricted to pairs for which there exists an everywhere continuous version of the conditional distribution.*

8. POSITIVE RESULTS

We now consider situations in which we can compute conditional distributions. Probabilistic methods have been widely successful in many settings, and so it is important to understand situations in which conditional inference is possible. We begin with the setting of discrete random variables.

8.1. Discrete Random Variables. A very common situation is that in which we condition with respect to a random variable taking values in a discrete set. As we will see, conditioning is always possible in this setting, as it reduces to the elementary notion of conditional probability with respect to single events.

Lemma 8.1 (Conditioning on a discrete random variable). *Let X and Y be computable random variables in computable metric spaces S and T , respectively, where S is a finite or countable and discrete set. Then the conditional distribution $\mathbf{P}[Y|X]$ is computable, uniformly in X and Y .*

Proof. Let κ be a version of the conditional distribution $\mathbf{P}[Y|X]$, and define $S_+ := \text{supp}(\mathbf{P}_X) = \{x \in S : \mathbf{P}_X\{x\} > 0\}$. We have $\mathbf{P}_X(S_+) = 1$, and so, by Lemma 3.8, κ is characterized by the equations

$$\kappa(x, \cdot) = \mathbf{P}\{Y \in \cdot | X = x\} = \frac{\mathbf{P}\{Y \in \cdot, X = x\}}{\mathbf{P}\{X = x\}}, \quad x \in S_+. \quad (47)$$

Let $B \subseteq T$ be an \mathbf{P}_Y -almost decidable. Because $\{x\}$ is decidable, $\{x\} \times B$ is $\mathbf{P}_{(X,Y)}$ -almost decidable, and so the numerator is a computable real, uniformly in $x \in S_+$ and B . Taking $B = T$, we see that the denominator, and hence the ratio is also

computable, uniformly in $x \in S_+$ and B . Thus, by Corollary 2.17, $x \mapsto \kappa(x, \cdot)$ is computable on a \mathbf{P}_X -measure one set. \square

Note that one could provide an alternative proof using the “rejection sampling” method, which is used to define the semantics of conditioning on discrete random variables in several probabilistic programming languages. This could, for example, be formalized by using least fixed points or unbounded search.

A related situation is one where a computable random variable X concentrates on a finite or countable subset of a computable metric space. In this case, Equation (47) is well-defined and characterizes every version of the conditional distribution, but the event $\{X = x\}$, for $x \in \text{supp}(\mathbf{P}_X)$, is not decidable. However, if the countable subset is discrete, then each such event is almost decidable. In order to characterize the computability of conditioning in this setting, we first define a notion of computability for discrete subsets of computable metric spaces.

Definition 8.2 (Computably discrete set). Let S be a computable metric space. We say that a (finite or countably infinite) subset $D \subseteq S$ is a **computably discrete subset** when there exists a function $f : S \rightarrow \mathbb{N}$ that is computable and injective on D . We call f the **witness** to the computable discreteness of D .

The following result is an easy consequence of the definition of computably discrete subsets and the computability of conditional distributions given discrete observations.

Lemma 8.3. *Let X and Y be computable random variables in computable metric spaces S and T , respectively, let $D \subseteq S$ be a computable discrete subset with witness f , and assume that $X \in D$ a.s. Then the conditional distribution $\mathbf{P}[Y|X]$ is computable, uniformly in X, Y and (the witness for) f .*

Proof. Let κ be a version of the conditional distribution of $\mathbf{P}[Y|X]$ and let κ_f be a version of the conditional distribution of $\mathbf{P}[Y|f(X)]$. Then, for $x \in D$,

$$\kappa(x, B) = \kappa_f(f(x), B), \quad (48)$$

and these equations completely characterize κ as $\mathbf{P}_X(D) = \mathbf{P}_{f(X)}(f(D)) = 1$. Note that $f(X)$ is a computable random variable in \mathbb{N} , and so by Lemma 8.1, κ_f is computable on $f(D)$, and so κ is computable on D . \square

8.2. Continuous and Dominated Settings. The most common way to calculate conditional probabilities is to use Bayes’ rule, which requires the existence of a conditional density. (Within statistics, a probability model is said to be *dominated* when there is a conditional density.) We first recall some basic definitions.

Definition 8.4 (Density). Let (S, \mathcal{A}, ν) be a measure space and let $f : S \rightarrow \mathbb{R}^+$ be an ν -integrable function. Then the function μ on \mathcal{A} given by

$$\mu(A) = \int_A f d\nu \quad (49)$$

for $A \in \mathcal{A}$ is a (finite) measure on (S, \mathcal{A}) and f is called a **density of μ (with respect to ν)**. Note that g is a density of μ with respect to ν if and only if $f = g$ ν -a.e.

Definition 8.5 (Conditional density). Let X and Y be random variables in (complete separable) metric spaces, let $\kappa_{X|Y}$ be a version of $\mathbf{P}[X|Y]$, and assume that there is a measure $\nu \in \mathcal{M}(S)$ and measurable function $p_{X|Y}(x|y) : S \times T \rightarrow \mathbb{R}^+$ such that $p_{X|Y}(\cdot|y)$ is a density of $\kappa_{X|Y}(y, \cdot)$ with respect to ν for \mathbf{P}_Y -almost all y . That is,

$$\kappa_{X|Y}(y, A) = \int_A p_{X|Y}(x|y) \nu(dx) \quad (50)$$

for measurable sets $A \subseteq S$ and \mathbf{P}_Y -almost all y . Then $p_{X|Y}(x|y)$ is called a **conditional density of X given Y (with respect to ν)**.

Common parametric families of distributions (e.g., exponential families like Gaussian, Gamma, etc.) admit conditional densities, and in these cases, the well-known Bayes' rule gives a formula for expressing the conditional distribution. We give a proof of this classic result for completeness.

Lemma 8.6 (Bayes' rule [Sch95, Thm. 1.13]). *Let X and Y be random variables as in Definition 3.5, let $\kappa_{X|Y}$ be a version of the conditional distribution $\mathbf{P}[X|Y]$, and assume that there exists a conditional density $p_{X|Y}(x|y)$. Then a function $\kappa : S \times \mathcal{B}_T \rightarrow [0, 1]$ satisfying*

$$\kappa_{Y|X}(x, B) = \frac{\int_B p_{X|Y}(x|y) \mathbf{P}_Y(dy)}{\int p_{X|Y}(x|y) \mathbf{P}_Y(dy)} \quad (51)$$

for those points x for which the denominator is positive and finite, is a version of the conditional distribution $\mathbf{P}[Y|X]$.

Proof. By Definition 3.5 and Fubini's theorem, for Borel sets $A \subseteq S$ and $B \subseteq T$, we have that

$$\mathbf{P}\{X \in A, Y \in B\} = \int_B \kappa_{X|Y}(y, A) \mathbf{P}_Y(dy) \quad (52)$$

$$= \int_B \left(\int_A p_{X|Y}(x|y) \nu(dx) \right) \mathbf{P}_Y(dy) \quad (53)$$

$$= \int_A \left(\int_B p_{X|Y}(x|y) \mathbf{P}_Y(dy) \right) \nu(dx). \quad (54)$$

Taking $B = T$, we have

$$\mathbf{P}_X(A) = \int_A \left(\int p_{X|Y}(x|y) \mathbf{P}_Y(dy) \right) \nu(dx). \quad (55)$$

Because $\mathbf{P}_X(S) = 1$, this implies that the set of points x for which the denominator of (51) is infinite has ν -measure zero, and thus \mathbf{P}_X -measure zero. Taking A to be the set of points x for which the denominator is zero, we see that $\mathbf{P}_X(A) = 0$. It

follows that (51) characterizes κ up to a \mathbf{P}_X -null set. Also by (55), we see that the denominator is a density of \mathbf{P}_X with respect to ν , and so we have

$$\int_A \kappa_{Y|X}(x, B) \mathbf{P}_X(dx) = \int_A \kappa_{Y|X}(x, B) \left(\int p_{X|Y}(x|y) \mathbf{P}_Y(dy) \right) \nu(dx). \quad (56)$$

Finally, by the definition of $\kappa_{Y|X}$, Equation (54), and the fact that the denominator is positive and finite for \mathbf{P}_X -almost all x , we see that $\kappa_{Y|X}$ is a version of the conditional distribution $\mathbf{P}[Y|X]$. \square

Comparing Bayes' rule (51) to the definition of conditional density (50), we see that any conditional density of Y given X (with respect to \mathbf{P}_Y) satisfies

$$p_{Y|X}(y|x) = \frac{p_{X|Y}(x|y)}{\int p_{X|Y}(x|y) \mathbf{P}_Y(dy)}, \quad (57)$$

for $\mathbf{P}_{(X,Y)}$ -almost all (x, y) . We can now give the proof of Lemma 4.8.

of Lemma 4.8. Let $\kappa_{Y|X}$ be given by (51), and let $B \subseteq T$ be an open set. By hypothesis, the map $\phi : S \rightarrow \mathcal{C}(T, \mathbb{R}^+)$ given by $\phi(x) = p_{X|Y}(x|\cdot)$ is continuous on R , while the indicator function $\mathbf{1}_B$ is lower semicontinuous. Integration of a lower semicontinuous function with respect to a probability measure is a lower semicontinuous operation, and so the map $x \mapsto \int \mathbf{1}_B \phi(x) d\mathbf{P}_Y$ is lower semicontinuous on R .

Note that for every $x \in R$, the function $\phi(x)$ is positive and bounded by hypothesis. Integration of a bounded continuous function with respect to a probability measure is a continuous operation, and so the map $x \mapsto \int \phi(x) d\mathbf{P}_Y$ is positive and continuous on R . Therefore the ratio in (51) is a lower semicontinuous function of $x \in R$ for fixed B , completing the proof. \square

Using the following well-known result about integration of computable functions, we can study when the conditional distribution characterized by Bayes' rule is computable.

Proposition 8.7 ([HR09a, Cor. 4.3.2]). *Let S be a computable metric space, and μ a computable probability measure on S . Let $f : S \rightarrow \mathbb{R}^+$ be a bounded computable function. Then $\int f d\mu$ is a computable real, uniformly in f .*

Corollary 8.8 (Density and independence). *Let U , V , and Y be computable random variables (in computable metric spaces), where Y is independent of V given U . Assume that there exists a conditional density $p_{Y|U}(y|u)$ of Y given U (with respect to ν) that is bounded and computable. Then the conditional distribution $\mathbf{P}[(U, V)|Y]$ is computable.*

Proof. Let $X = (U, V)$. Then $p_{Y|X}(y|(u, v)) = p_{Y|U}(y|u)$ is the conditional density of Y given X (with respect to ν). Therefore, the computability of the integrand and the existence of a bound imply, by Proposition 8.7, that $\mathbf{P}[(U, V)|Y]$ is computable. \square

8.3. Conditioning on Noisy Observations. As an immediate consequence of Corollary 8.8, we obtain the computability of the following common situation in probabilistic modeling: where the observed random variable has been corrupted by independent absolutely continuous noise.

Corollary 8.9 (Independent noise). *Let U be a computable random variable in a computable metric space and let V and E be computable random variables in \mathbb{R} . Define $Y = U + E$. If \mathbf{P}_E is absolutely continuous with a bounded computable density p_E and E is independent of U and V then the conditional distribution $\mathbf{P}[(U, V) \mid Y]$ is computable.*

Proof. We have that

$$p_{Y|U}(y|u) = p_E(y - u) \tag{58}$$

is the conditional density of Y given U (with respect to Lebesgue measure). The result then follows from Corollary 8.8. \square

Pour-El and Richards [PER89, Ch. 1, Thm. 2] show that a twice continuously differentiable computable function has a computable derivative (despite the fact that Myhill [Myh71] exhibits a computable function $[0, 1] \rightarrow \mathbb{R}$ whose derivative is continuous, but not computable). Therefore, noise with a sufficiently smooth distribution has a computable density, and by Corollary 8.9, a computable random variable corrupted by such noise still admits a computable conditional distribution.

Furthermore, Corollary 8.9 implies that noiseless observations cannot always be computably approximated by noisy ones. For example, even though an observation corrupted with zero mean Gaussian noise with standard deviation σ may recover the original condition as $\sigma \rightarrow 0$, by our main noncomputability result (Theorem 6.7) one cannot, in general, compute how small σ must be in order to bound the error introduced by noise.

This result is also analogous to a classical theorem of information theory. Hartley [Har28] and Shannon [Sha49] show that the capacity of a continuous real-valued channel without noise is infinite, yet the addition of, e.g., Gaussian noise with $\varepsilon > 0$ variance causes the channel capacity to drop to a finite amount. The Gaussian noise prevents all but a finite amount of information from being encoded in the bits of the real number. Similarly, the amount of information in a continuous observation is too much in general for a computer to be able to update a probabilistic model. However, the addition of noise, as above, is sufficient for making conditioning possible on a computer.

The computability of conditioning with noise, coupled with the noncomputability of conditioning in general, has significant implications for our ability to recover a signal when noise is added, and suggests several interesting questions. For example, suppose we have a uniformly computable sequence of noise $\{E_n\}_{n \in \mathbb{N}}$ with absolutely continuous, uniformly computable densities such that the magnitude of the densities goes to 0 in some sufficiently nice way, and consider $Y_n := U + E_n$. Such a situation could arise, e.g., when we have a signal with noise but some way to reduce the noise over time.

When there is a continuous version of $\mathbf{P}[(U, V)|Y]$, we have

$$\lim_{n \rightarrow \infty} \mathbf{P}[(U, V)|Y_n] = \mathbf{P}[(U, V)|Y] \quad \text{a.s.} \quad (59)$$

However, we know that the right side is, in general, noncomputable, despite the fact that each term in the limit on the left side is computable.

This raises several questions, such as: What do bounds on how fast the sequence $\{\mathbf{P}[(U, V)|Y_n]\}_{n \in \mathbb{N}}$ converges to $\mathbf{P}[(U, V)|Y]$ tell us about the computability of $\mathbf{P}[(U, V)|Y]$? What conditions on the relationship between U and the sequence $\{E_n\}_{n \in \mathbb{N}}$ will allow us to recover information about $\mathbf{P}[(U, V)|Y]$ from individual distributions $\mathbf{P}[(U, V)|Y_n]$?

8.4. Other Settings. Freer and Roy [FR10] show how to compute conditional distributions in the setting of *exchangeable sequences*. A classic result by de Finetti shows that exchangeable sequences of random variables are in fact conditionally i.i.d. sequences, conditioned on a random measure, often called the *directing random measure*. Freer and Roy describe how to transform an algorithm for sampling an *exchangeable sequence* into a rule for computing the posterior distribution of the directing random measure given observations. The result is a corollary of a computable version of de Finetti’s theorem [FR09], and covers a wide range of common scenarios in nonparametric Bayesian statistics (often where no conditional density exists).

ACKNOWLEDGMENTS

A preliminary version of this article appeared as “Noncomputable conditional distributions” in *Proceedings of the 26th Annual IEEE Symposium on Logic in Computer Science*, 107–116 [AFR11].

CEF was partially supported by NSF grants DMS-0901020 and DMS-0800198. His work on this publication was made possible through the support of a grant from the John Templeton Foundation. The opinions expressed in this publication are those of the authors and do not necessarily reflect the views of the John Templeton Foundation.

DMR was partially supported by graduate fellowships from the National Science Foundation and MIT Lincoln Laboratory and is supported by a Newton International Fellowship and Emmanuel College.

The authors would like to thank Jeremy Avigad, Henry Cohn, Leslie Kaelbling, Vikash Mansinghka, Hartley Rogers, Michael Sipser, and Joshua Tenenbaum for helpful discussions, Bill Thurston for a useful comment regarding Lemma 7.1, Martín Escardó for a helpful discussion regarding computable random variables, and Quinn Culver, Noah Goodman, Leslie Kaelbling, Bjørn Kjos-Hanssen, Vikash Mansinghka, Timothy O’Donnell, and Geoff Patterson for comments on earlier drafts.

REFERENCES

- [AFR11] N. L. Ackerman, C. E. Freer, and D. M. Roy, *Noncomputable conditional distributions*, Proc. of the 26th Ann. IEEE Symp. on Logic in Comput. Sci. (LICS 2011), IEEE Computer Society, 2011, pp. 107–116.

- [BC06] M. Braverman and S. Cook, *Computing over the reals: foundations for scientific computing*, Notices Amer. Math. Soc. **53** (2006), no. 3, 318–329.
- [BCGL92] S. Ben-David, B. Chor, O. Goldreich, and M. Luby, *On the theory of average case complexity*, J. Comput. System Sci. **44** (1992), no. 2, 193–219.
- [Bos08] V. Bosserhoff, *Notions of probabilistic computability on represented spaces*, J.UCS **14** (2008), no. 6, 956–995.
- [Bra05] M. Braverman, *On the complexity of real functions*, FOCS'05—Proc. of the 46th Ann. IEEE Symp. on Foundations of Comput. Sci., IEEE Computer Society, 2005, pp. 155–164.
- [Eda96] A. Edalat, *The Scott topology induces the weak topology*, Proc. of the 11th Ann. IEEE Symp. on Logic in Comput. Sci. (New Brunswick, NJ, 1996), IEEE Comput. Soc. Press, Los Alamitos, CA, 1996, pp. 372–381.
- [EH98] A. Edalat and R. Heckmann, *A computational model for metric spaces*, Theoret. Comput. Sci. **193** (1998), no. 1-2, 53–73.
- [FR09] C. E. Freer and D. M. Roy, *Computable exchangeable sequences have computable de Finetti measures*, Mathematical Theory and Computational Practice (CiE 2009), Proc. of the 5th Conf. on Computability in Europe (K. Ambos-Spies, B. Löwe, and W. Merkle, eds.), Lecture Notes in Comput. Sci., vol. 5635, Springer, 2009, pp. 218–231.
- [FR10] ———, *Posterior distributions are computable from predictive distributions*, Proc. of the 13th Int. Conf. on Artificial Intelligence and Statistics (AISTATS 2010) (Y. W. Teh and M. Titterton, eds.), JMLR: W&CP 9, 2010, pp. 233–240.
- [Gác05] P. Gács, *Uniform test of algorithmic randomness over a general space*, Theoret. Comput. Sci. **341** (2005), no. 1-3, 91–137.
- [GHR10] S. Galatolo, M. Hoyrup, and C. Rojas, *Effective symbolic dynamics, random points, statistical behavior, complexity and entropy*, Inform. and Comput. **208** (2010), no. 1, 23–41.
- [GMR⁺08] N. D. Goodman, V. K. Mansinghka, D. M. Roy, K. Bonawitz, and J. B. Tenenbaum, *Church: a language for generative models*, Proc. of the 24th Conf. on Uncertainty in Artificial Intelligence, 2008.
- [Gol67] E. M. Gold, *Language identification in the limit*, Inform. and Control **10** (1967), no. 5, 447–474.
- [Har28] R. Hartley, *Transmission of information*, Bell System Technical J. (1928), 535–563.
- [Hem02] A. Hemmerling, *Effective metric spaces and representations of the reals*, Theoret. Comput. Sci. **284** (2002), no. 2, 347–372, Computability and complexity in analysis (Castle Dagstuhl, 1999).
- [HR09a] M. Hoyrup and C. Rojas, *Computability of probability measures and Martin-Löf randomness over metric spaces*, Inform. and Comput. **207** (2009), no. 7, 830–847.
- [HR09b] M. Hoyrup and C. Rojas, *Applications of effective probability theory to Martin-Löf randomness*, Automata, Languages and Programming (S. Albers, A. Marchetti-Spaccamela, Y. Matias, S. Nikolettseas, and W. Thomas, eds.), Lecture Notes in Comput. Sci., vol. 5555, Springer, 2009, pp. 549–561.
- [HR11] M. Hoyrup and C. Rojas, *Absolute continuity of measures and preservation of randomness*, Preprint, <http://www.loria.fr/~hoyrup/abscont.pdf>, 2011.
- [HRW11] M. Hoyrup, C. Rojas, and K. Weihrauch, *Computability of the Radon-Nikodym derivative*, Models of Computation in Context, Proc. of the 7th Conf. on Computability in Europe (CiE 2011) (B. Löwe, D. Normann, I. N. Soskov, and A. A. Soskova, eds.), Lecture Notes in Comput. Sci., vol. 6735, Springer, 2011, pp. 132–141.
- [Hut07] M. Hutter, *On universal prediction and Bayesian confirmation*, Theoret. Comput. Sci. **384** (2007), no. 1, 33–48.
- [JB97] Jens and Blanck, *Domain representability of metric spaces*, Annals of Pure and Applied Logic **83** (1997), no. 3, 225–247.

- [Kal02] O. Kallenberg, *Foundations of modern probability*, 2nd ed., Springer, New York, 2002.
- [Kol33] A. N. Kolmogorov, *Grundbegriffe der Wahrscheinlichkeitsrechnung*, Springer, 1933.
- [KS09] O. Kiselyov and C. Shan, *Embedded probabilistic programming*, Domain-Specific Languages (W. M. Taha, ed.), Lecture Notes in Computer Science, vol. 5658, Springer, 2009, pp. 360–384.
- [KY76] D. E. Knuth and A. C. Yao, *The complexity of nonuniform random number generation*, Algorithms and complexity (Proc. Sympos., Carnegie-Mellon Univ., Pittsburgh, Pa., 1976), Academic Press, New York, 1976, pp. 357–428.
- [Lev86] L. A. Levin, *Average case complete problems*, SIAM J. Comput. **15** (1986), no. 1, 285–286.
- [Man73] I. Mann, *Probabilistic recursive functions*, Trans. Amer. Math. Soc. **177** (1973), 447–467.
- [MWGK10] T. Minka, J. Winn, J. Guiver, and D. Knowles, *Infer.NET 2.4*, 2010, Microsoft Research Cambridge. <http://research.microsoft.com/infernet>.
- [Myh71] J. Myhill, *A recursive function, defined on a compact interval and having a continuous derivative that is not recursive*, Michigan Math. J. **18** (1971), 97–98.
- [Nie09] A. Nies, *Computability and randomness*, Oxford Logic Guides, vol. 51, Oxford University Press, Oxford, 2009.
- [OSW88] D. N. Osherson, M. Stob, and S. Weinstein, *Mechanical learners pay a price for Bayesianism*, J. Symbolic Logic **53** (1988), no. 4, 1245–1251.
- [PER84] M. B. Pour-El and I. Richards, *L^p -computability in recursive analysis*, Proc. Amer. Math. Soc. **92** (1984), no. 1, 93–97.
- [PER89] M. B. Pour-El and J. I. Richards, *Computability in analysis and physics*, Perspectives in Mathematical Logic, Springer-Verlag, Berlin, 1989.
- [Pfa79] J. Pfanzagl, *Conditional distributions as derivatives*, Ann. Probab. **7** (1979), no. 6, 1046–1050.
- [Pfe01] A. Pfeffer, *IBAL: A probabilistic rational programming language*, Proc. of the 17th Int. Joint Conf. on Artificial Intelligence, Morgan Kaufmann Publ., 2001, pp. 733–740.
- [Poo91] D. Poole, *Representing Bayesian networks within probabilistic Horn abduction*, Proc. of the 7th Conf. on Uncertainty in Artificial Intelligence, 1991, pp. 271–278.
- [PPT08] S. Park, F. Pfennig, and S. Thrun, *A probabilistic language based on sampling functions*, ACM Trans. Program. Lang. Syst. **31** (2008), no. 1, 1–46.
- [Put65] H. Putnam, *Trial and error predicates and the solution to a problem of Mostowski*, J. Symbolic Logic **30** (1965), 49–57.
- [Rao88] M. M. Rao, *Paradoxes in conditional probability*, J. Multivariate Anal. **27** (1988), no. 2, 434–446.
- [Rao05] ———, *Conditional measures and applications*, second ed., Pure and Applied Mathematics, vol. 271, Chapman & Hall/CRC, Boca Raton, FL, 2005.
- [RD06] M. Richardson and P. Domingos, *Markov logic networks*, Machine Learning **62** (2006), no. 1–2, 107–136.
- [Rog87] H. Rogers, Jr., *Theory of recursive functions and effective computability*, second ed., MIT Press, Cambridge, MA, 1987.
- [Roy11] D. M. Roy, *Computability, inference and modeling in probabilistic programming*, Ph.D. thesis, Massachusetts Institute of Technology, 2011.
- [Sch95] M. J. Schervish, *Theory of statistics*, Springer Series in Statistics, Springer-Verlag, New York, 1995.
- [Sch07a] M. Schröder, *Admissible representations for probability measures*, Math. Log. Q. **53** (2007), no. 4–5, 431–445.
- [Sch07b] ———, *Admissible representations of probability measures*, Proceedings of the Third International Conference on Computability and Complexity in Analysis (CCA 2006), Electron. Notes Theor. Comput. Sci., vol. 167, Elsevier Sci. B. V., Amsterdam, 2007, pp. 61–78 (electronic).

- [Sha49] C. E. Shannon, *Communication in the presence of noise*, Proc. I.R.E. **37** (1949), 10–21.
- [Sol64] R. J. Solomonoff, *A formal theory of inductive inference II*, Inform. and Control **7** (1964), 224–254.
- [Tak08] H. Takahashi, *On a definition of random sequences with respect to conditional probability*, Inform. and Comput. **206** (2008), no. 12, 1375–1382.
- [Tju74] T. Tjur, *Conditional probability distributions*, Lecture Notes, no. 2, Institute of Mathematical Statistics, University of Copenhagen, Copenhagen, 1974.
- [Tju75] ———, *A constructive definition of conditional distributions*, Preprint 13, Institute of Mathematical Statistics, University of Copenhagen, Copenhagen, 1975.
- [Tju80] ———, *Probability based on Radon measures*, Wiley Series in Probability and Mathematical Statistics, John Wiley & Sons Ltd., Chichester, 1980.
- [Wei93] K. Weihrauch, *Computability on computable metric spaces*, Theoret. Comput. Sci. **113** (1993), no. 2, 191–210.
- [Wei99] ———, *Computability on the probability measures on the Borel sets of the unit interval*, Theoret. Comput. Sci. **219** (1999), no. 1-2, 421–437.
- [Wei00] ———, *Computable analysis: an introduction*, Springer-Verlag, Berlin, 2000.
- [Yam99] T. Yamakami, *Polynomial time samplable distributions*, J. Complexity **15** (1999), no. 4, 557–574.
- [ZL70] A. K. Zvonkin and L. A. Levin, *The complexity of finite objects and the basing of the concepts of information and randomness on the theory of algorithms*, Uspehi Mat. Nauk **25** (1970), no. 6 (156), 85–127.

DEPARTMENT OF MATHEMATICS, HARVARD UNIVERSITY, ONE OXFORD STREET, CAMBRIDGE, MA 02138–2901, USA

E-mail address: `nate@math.harvard.edu`

COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY, MASSACHUSETTS INSTITUTE OF TECHNOLOGY, 32 VASSAR STREET, CAMBRIDGE, MA 02139–4309, USA

E-mail address: `freer@math.mit.edu`

DEPARTMENT OF ENGINEERING, UNIVERSITY OF CAMBRIDGE, TRUMPINGTON STREET, CAMBRIDGE CB2 1PZ, UK

E-mail address: `d.roy@eng.cam.ac.uk`