On the Computability of Conditional Probability

NATHANAEL L. ACKERMAN, Harvard University

CAMERON E. FREER, Massachusetts Institute of Technology

DANIEL M. ROY, University of Toronto

As inductive inference and machine-learning methods in computer science see continued success, researchers are aiming to describe ever more complex probabilistic models and inference algorithms. It is natural to ask whether there is a universal computational procedure for probabilistic inference. We investigate the computability of conditional probability, a fundamental notion in probability theory, and a cornerstone of Bayesian statistics. We 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 bounded density.

CCS Concepts: • **Mathematics of computing** → **Probability and statistics**; • **Theory of computation** → *Probabilistic computation*;

Additional Key Words and Phrases: Computable probability theory, conditional probability, real computation

ACM Reference format:

Nathanael L. Ackerman, Cameron E. Freer, and Daniel M. Roy. 2019. On the Computability of Conditional Probability. *J. ACM* 66, 3, Article 23 (June 2019), 40 pages.

<https://doi.org/10.1145/3321699>

A preliminary version of this article appears as “Noncomputable conditional distributions” in *Proceedings of the 26th Annual IEEE Symposium on Logic in Computer Science*, 107–116 (Ackerman et al. 2011), which was based upon results first appearing in an arXiv preprint (Ackerman et al. 2010).

C. E. Freer was partially supported by NSF grants DMS-0901020 and DMS-0800198, DARPA Contract Award No. FA8750-14-C-0001, and a grant from the John Templeton Foundation.

D. M. Roy was partially supported by graduate fellowships from the National Science Foundation and MIT Lincoln Laboratory, a Newton International Fellowship, a Research Fellowship at Emmanuel College, and an Ontario Early Researcher Award.

Work on this publication was also made possible through the support of ARO grant W911NF-13-1-0212, ONR grant N00014-13-1-0333, DARPA Contract Award No. FA8750-14-2-0004, and Google’s “Rethinking AI” project. The opinions expressed in this publication are those of the authors and do not necessarily reflect the views of the John Templeton Foundation or the U.S. Government.

Authors’ addresses: N. L. Ackerman, Harvard University, Department of Mathematics, One Oxford St., Cambridge, MA, 02138-2901; email: nate@math.harvard.edu; C. E. Freer, Massachusetts Institute of Technology, Department of Brain and Cognitive Sciences, 77 Massachusetts Ave., Cambridge, MA, 02139-4301; email: freer@mit.edu; D. M. Roy, University of Toronto, Department of Statistical Sciences, 100 St. George St., Toronto, ON, M5S 3G3 and Vector Institute, MaRS Centre, West Tower, 661 University Ave., Suite 710, Toronto, ON M5G 1M1; email: droy@utstat.toronto.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

0004-5411/2019/06-ART23 \$15.00

<https://doi.org/10.1145/3321699>

1 INTRODUCTION

The use of probability to reason about uncertainty has wide-ranging applications in science and engineering, and some of the most important computational problems relate to *conditioning*, which is used to perform Bayesian inductive reasoning in probabilistic models. As researchers have faced more complex phenomena, their representations have also increased in complexity, which in turn has led to more complicated inference algorithms. It is natural to ask whether there is a universal inference algorithm—in other words, whether it is possible to automate probabilistic reasoning via a general procedure that can compute conditional probabilities for an *arbitrary* computable joint distribution.

We demonstrate that there are computable joint distributions with noncomputable conditional distributions. As a consequence, no general algorithm for computing conditional probabilities can exist. 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 vein, we describe broadly applicable conditions under which conditional probabilities are computable.

We begin by describing a setting, *probabilistic programming*, that motivates the search for these results. We proceed to describe the technical frameworks for our results, *computable probability theory* and the modern formulation of *conditional probability*. We then highlight related work and end the introduction with a summary of results of the article.

1.1 Probabilistic Programming

Within probabilistic artificial intelligence and machine learning, *probabilistic programming* provides formal languages and algorithms for describing and computing answers from probabilistic models. Probabilistic programming languages themselves build on modern programming languages and their facilities for recursion, abstraction, modularity, and so on, 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. Probabilistic programming languages have been the focus of a long tradition of research within programming languages, model checking, and formal methods. For some of the early approaches within the AI and machine-learning community, see, e.g., the languages PHA (Poole 1991), IBAL (Pfeffer 2001), Markov Logic (Richardson and Domingos 2006), λ_o (Park et al. 2008), Church (Goodman et al. 2008), HANSEI (Kiselyov and Shan 2009), and Infer.NET (Minka et al. 2010).

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 incomplete. Our results explain why this is necessarily the case.

1.2 Computable Probability Theory

To study computable probability theory and the computability of conditioning, we work within the framework of Type-2 Theory of Effectivity (TTE) and use appropriate representations for topological and measurable objects such as distributions, random variables, and maps between them. This framework builds upon and contains as a special case ordinary Turing computation on discrete spaces and gives us a basis for precisely describing the operations that probabilistic programming languages are capable of performing.

In particular, we study the computability of distributions on computable Polish spaces including, e.g., certain spaces of 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 (1933) gave an axiomatic characterization of conditional probabilities and an abstract construction of them using Radon–Nikodym derivatives, but this definition and construction do not yield a general recipe for their calculation. There is a further problem: 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 article, we are therefore especially interested in situations where it makes sense to ask for the conditional distribution given a particular point. One of our main results is in the setting where there is a unique continuous conditional distribution. In this case, conditioning yields a canonical answer, which is a natural desideratum for statistical applications.

A large body of work in probability and statistics is concerned with the derivation of conditional probabilities and distributions in special circumstances, each situation often requiring some special insight into the structure of the answer, especially when it was desirable for conditional probabilities and distributions to be defined at points, as in Bayesian statistical applications. This state of affairs motivated work on constructive definitions of conditioning (such as those due to Tjur (1974, 1975, 1980), Pfanzagl (1979), and Rao (1988, 2005)), although this work has not been sensitive to issues of computability.

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), as we describe in Section 9.2.

We recall the basics of the measure-theoretic approach to conditional probability 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

We now describe several other connections between conditional probability and computation.

1.4.1 Complexity Theory of Finite Discrete Distributions. 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 (Levin 1986). 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 et al. (1992) in their theory of polynomial-time samplable distributions, which has since been extended by Yamakami (1999) and others. Other positive and negative complexity results have been obtained in the particular case of Bayesian networks by Cooper (1990) and Dagum and Luby (1993, 1997). Extending these complexity results to the more general setting considered here could bear on the practice of statistical AI and machine learning.

1.4.2 Computable Bayesian Learners. Osherson et al. (1988) study learning theory in the setting of *identifiability in the limit* (see Gold (1967) and Putnam (1965) 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 “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 (2011) 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, in the context of statistical analysis, it is more appropriate to study the conditioning 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.

1.4.3 Induction with Respect to Universal Priors. 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 (1964), Zvonkin and Levin (1970), and Hutter (2007). 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 (Takahashi 2008). The objects with respect to which one is conditioning in these settings are typically not computable (e.g., the universal semimeasure is merely lower semicomputable). In the present article, 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.

1.4.4 Radon–Nikodym Derivatives. In the abstract setting, conditional probabilities are (suitably measurable) Radon–Nikodym derivatives. In work motivated by questions in algorithmic randomness, Hoyrup and Rojas (2011) study notions of computability for absolute continuity and for Radon–Nikodym derivatives as elements in L^1 , i.e., the space of integrable functions. Hoyrup et al. (2011) then show an equivalence between the problem of computing Radon–Nikodym derivatives as elements in L^1 and computing the characteristic function of computably enumerable sets. The noncomputability of the Radon–Nikodym derivative operator is demonstrated by a pair μ, ν of computable measures whose Radon–Nikodym derivative $d\mu/d\nu$ is not computable as an element in $L^1(\nu)$. However, the Radon–Nikodym derivatives they study do not correspond to conditional probabilities, and so the computability of the operator restricted to those maps arising in

the construction of conditional probabilities is not addressed by this work. The underlying notion of computability is another important difference. An element in $L^1(\nu)$ is an equivalence class of functions, every pair agreeing on a set of ν -measure one. Thus, one cannot, in general, evaluate these derivatives at points in a well-defined manner. Most Bayesian statisticians would be unfamiliar and perhaps unsatisfied with this notion of computability, especially in settings where their statistical models admit continuous versions of Radon–Nikodym derivatives that are unique, and thus well defined pointwise. Regardless, we will show that even in such settings, computing conditional probabilities is not possible in general, even in the weaker L^1 sense. However, our positive results do yield computable probabilities/distributions defined pointwise.

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. Recall that measurable functions are often defined only up to a measure-zero set; any two functions that agree almost everywhere are called *versions* of each other.

In Proposition 5.1, we construct random variables X and C that are computable on a $\mathbb{P}$ -measure one set, such that every version of the conditional distribution map $\mathbb{P}[C|X = \cdot]$ (i.e., a probability-measure-valued function f such that $f(X)$ is a regular version of the conditional distribution $\mathbb{P}[C|X]$) is discontinuous everywhere, even when restricted to a $\mathbb{P}_X$ -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.

In fact, even under this restriction, conditioning is not even continuous, let alone computable, as we show in Section 6. We further demonstrate that if some computer program purports to state true facts about the conditional distribution of a computable joint distribution provided as input, then we can uniformly find some other representation of the input distribution such that the program does not output any nontrivial fact about the conditional distribution.

Our central result, Theorem 7.6, provides a pair of random variables that are computable on a measure one set, but such that the conditional distribution of one variable given the other is not computable on any measure one set (though some version is continuous on a measure one set). The construction involves encoding the halting times of all Turing machines into the conditional distribution map while ensuring that the joint distribution remains computable. This result yields another proof of the noncomputability of the conditioning operation restricted to measures having conditional distributions that are continuous on a measure one set.

In Section 8, we extend our central result by constructing a pair of random variables, again computable on a measure one set, whose conditional distribution map is noncomputable but has an everywhere continuous version with infinitely differentiable conditional probability maps. This construction proceeds by smoothing out the distribution constructed in Section 7, but in such a way that one can still compute the halting problem relative to the conditional distribution. This result implies that conditioning is not a computable operation, even when we further restrict to the case where the conditional distribution has an everywhere continuous version.

Despite the noncomputability of conditioning in general, conditional distribution maps 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 (Proposition 9.2) and where there is a conditional density (Corollary 9.6).

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 where U and E are real-valued, and suppose that $\mathbb{P}_E$ 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 9.7, we show that the conditional distribution map $\mathbb{P}[(U, V)|U + E = \cdot]$ is computable (even if $\mathbb{P}[(U, V)|U = \cdot]$ is not). Finally, we discuss how symmetry, in the form of *exchangeability*, 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 (1996), Weihrauch (1999), Schröder (2007), and Gács (2005). Computable probability theory itself builds off notions and results in computable analysis, specifically the Type-2 Theory of Effectivity. For a general introduction to this approach to real computation, see Weihrauch (2000), Braverman (2005), or Braverman and Cook (2006).

2.1 Computable and Computable Enumerable Reals

We first recall some elementary definitions from computability theory (see, e.g., Rogers (1987, Ch. 5)). A set of natural numbers (potentially in some correspondence with, e.g., rationals, integers, or other finitely describable objects with an implicit enumeration) is *computable* when there is a computer program that, given k , outputs whether or not k is in the set. A set is *computably enumerable* (c.e.) when there is a computer program that outputs every element of the set eventually. Note that a set is computable when both it and its complement are c.e. We say that a sequence of sets $\{B_n\}$ is computable *uniformly in n* when there is a single computer program that, given n and k , outputs whether or not k is in B_n . We say that the sequence is c.e. *uniformly in n* when there is a computer program that, on input n , outputs every element of B_n eventually.

We now recall basic notions of computability for real numbers (see, e.g., Weihrauch (2000, Ch. 4.2) or Nies (2009, Ch. 1.8)). We say that a real r is a *c.e. real* (sometimes called a *left-c.e. real*) when the set of rationals $\{q \in \mathbb{Q} : q < r\}$ is c.e. A real r is *computable* when both it and its negative are 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 semicomputable when $f(n)$ is a c.e. real, uniformly in n (i.e., when the collection of rationals less than $f(n)$ is c.e. uniformly in n). Likewise, a function is upper semicomputable when its negative is lower semicomputable. The function f is computable if and only if it is both lower and upper semicomputable.

2.2 Computable Polish Spaces

Recall that a **Polish space** is a topological space that admits a metric under which it is a complete separable metric space. Computable Polish spaces, as developed in computable analysis (Hemmerling 2002; Weihrauch 1993) and effective domain theory (Blanck 1997; Edalat and Heckmann 1998), provide a convenient framework for formulating results in computable probability theory. For consistency, we largely use definitions from Hoyrup and Rojas (2009b) and Galatolo et al.

(2010). Additional details about computable Polish spaces, sometimes called computable metric spaces or effective Polish spaces, can also be found in Weihrauch (2000, Ch. 8.1), Gács (2005, §B.3), and Moschovakis (2009, Ch. 3I).

Definition 2.1 (Computable Polish Space (Galatolo et al. 2010, Def. 2.3.1)). A **computable Polish space** is a triple $(S, \delta, \mathcal{D})$ 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 .

In particular, note that condition (1) implies that the topological space determined by the metric space (S, δ) is a Polish space.

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

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

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

Let $\mathcal{B}_S$ denote the Borel σ -algebra on a Polish space S , i.e., the σ -algebra generated by the open balls of S . Let $\mathcal{M}_1(S)$ denote the set of Borel probability measures on S .

In this article we primarily work with computable Polish spaces. As such, unless otherwise noted, the σ -algebras will always be the Borel σ -algebras on such spaces—in particular, making them standard Borel spaces. Measurable functions between Polish spaces will always be measurable with respect to the Borel σ -algebras. We will sometimes refer to measurable subsets of a probability space as *events*.

Example 2.2. The set $\{0, 1\}$ is a computable Polish space under the discrete metric, where $\delta(0, 1) = 1$.

Cantor space, the set $\{0, 1\}^\infty$ of infinite binary sequences, is a computable Polish 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 Polish space under the Euclidean metric with the dense set $\mathbb{Q}$ of rationals (under its standard enumeration).

Suppose we are given a finite sequence $(T_0, \delta_0, \mathcal{D}_0), \dots, (T_{n-1}, \delta_{n-1}, \mathcal{D}_{n-1})$ of computable Polish spaces. Then the product metric space $\prod_{i=0}^{n-1} T_i$ (with one of any of the equivalent standard product metrics) is a computable Polish space where the ideal points consist of all finite products of ideal points. Furthermore, given a countably infinite such sequence $(T_0, \delta_0, \mathcal{D}_0), (T_1, \delta_1, \mathcal{D}_1), \dots$ that is uniformly computable and has a fixed bound on the diameter, the product metric spaces consists of the metric space whose underlying space is $\prod_{i \in \mathbb{N}} T_i$ and whose metric is given by $\delta(x, y) = \sum_{i \in \mathbb{N}} 2^{-i} \delta_i(x, y)$; this, too, can be made into a computable Polish space, by taking the ideal points to be those sequences $(x_0, x_1, \dots)$ with each $x_i \in \mathcal{D}_i$ such that for all but finitely many terms i , the point x_i is the first element in the enumeration of $\mathcal{D}_i$. Note that in both the finite and infinite case, all projection maps are computable.

Definition 2.3 (Computable Point (Galatolo et al. 2010, Def. 2.3.2)). Let $(S, \delta, \mathcal{D})$ be a computable Polish space with $\mathcal{D} = \{s_j\}_{j \in \mathbb{N}}$ and $x \in S$. Given a sequence $\{i_k\}_{k \in \mathbb{N}}$ of natural numbers, we say that the sequence $\{s_{i_k}\}_{k \in \mathbb{N}}$ of elements of $\mathcal{D}$ is a **representation** of the point x if $\delta(s_{i_k}, x) < 2^{-k}$ for all k . When $\{i_k\}_{k \in \mathbb{N}}$ is a computable sequence such that $\{s_{i_k}\}_{k \in \mathbb{N}}$ is a representation of x , we say that $\{s_{i_k}\}_{k \in \mathbb{N}}$ is a **computable representation**, and that the point x is **computable**.

Remark 2.4. 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 Polish 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 Polish space.

Definition 2.5 (C.e. Open Set (Galatolo et al. 2010, Def. 2.3.3)). Let $(S, \delta, \mathcal{D})$ be a computable Polish 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 2.6 (Computable Partial Function (Galatolo et al. 2010, Def. 2.3.6)). Let $(S, \delta_S, \mathcal{D}_S)$ and $(T, \delta_T, \mathcal{D}_T)$ be computable Polish 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 **continuous on** $R \subseteq S$ when f restricted to R is continuous as a function from R , under the subspace topology to 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 .

Note that the notion of being computable on a set R can be relativized to an oracle $A \subseteq \mathbb{N}$ in the obvious way. A function is continuous on R if and only if it is A -computable on R for some oracle A .

Remark 2.7. Let S and T be computable Polish 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 an oracle Turing machine that, upon being fed a representations of points $x \in R$ on its oracle tape, computes representations of their images $f(x) \in S$. (For more details, see Hoyrup and Rojas (2009b, Prop. 3.3.2)).

2.3 Notions of Computability for Functions on Probability Spaces

The standard notion of computability of functions between computable Polish spaces is too restrictive in most cases when the inputs to these functions are points in a probability space. For example, the Heaviside function $f(x) = \mathbf{1}(x \geq 0)$ is not computable on any set containing a neighborhood of 0. However, we can reliably compute the image of a Gaussian random variable under f , because the Gaussian random variable is nonzero with probability one, and f is computable on $\mathbb{R} \setminus \{0\}$.

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 predicate P on Ω is said to hold **μ -almost everywhere** (abbreviated **μ -a.e.**) if the event $E_P = \{\omega \in \Omega : P(\omega) \text{ does not hold}\}$ is a μ -null set. When E_P is a μ -null set but μ is a probability measure, we will instead say the event P holds **μ -almost surely**, and we likewise 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 $\mathbb{P}$).

Definition 2.8. Let S and T be Polish spaces and μ a probability measure on S . A measurable function $f : S \rightarrow T$ is **μ -almost continuous** when it is continuous on a μ -measure one set. When

S and T are computable Polish spaces, the measurable function f is **μ -almost computable** when it is computable on a μ -measure one set.

(See Hoyrup and Rojas (2009b) for further development of the theory of almost computable functions.) The following result relates μ -almost continuity to μ -a.e. continuity, i.e., the set of continuity points being a μ -measure one set. The proofs of the following proposition and lemma are due to François Dorais et al. (2013).

PROPOSITION 2.9. *Let X and Y be Polish spaces, let $f : X \rightarrow Y$ be a μ -almost continuous function, and let μ be a probability measure on X . Then there is a μ -a.e. continuous $g : X \rightarrow Y$ that agrees with f μ -a.e.*

We will need the following technical lemma. Recall that a G_δ set is a countable intersection of open sets.

LEMMA 2.10. *Let X be a Polish space. If $D \subseteq X$ is a nonempty G_δ -set, then there is a measurable map $h : X \rightarrow D$ such that $\lim_{x \rightarrow x_0} h(x) = x_0$ for every $x_0 \in D$.*

PROOF. Suppose $D = \bigcap_{n \in \mathbb{N}} U_n$, where $(U_n)_{n \in \mathbb{N}}$ is a descending sequence of open sets such that $U_n \subseteq \bigcup_{x_0 \in D} B(x_0, 1/(n+1))$. Any measurable retraction $h : X \rightarrow D$ with the property that if $x \in U_n \setminus U_{n+1}$, then $d(h(x), x) < 1/(n+1)$ will be as required. By definition, it is always possible to find a suitable $h(x) \in D$ for each $x \in U_0 \setminus D$. To ensure that h is measurable, fix an enumeration $(d_i)_{i \in \mathbb{N}}$ of a countable dense subset of D and, if $x \in U_0 \setminus D$, then define $h(x)$ to be the first element in this list that matches all the necessary requirements. (We must have $h(x) = x$ for $x \in D$, and it does not matter how $h(x)$ is defined when $x \notin U_0$ so long as the end result is measurable.) $\square$

PROOF OF PROPOSITION 2.9. By a classical result of Kuratowski (Kechris 1995, I.3.B, Thm. 3.8), we may assume (after first possibly changing its value on a μ -null set) that f is continuous on a μ -measure one G_δ set $D \subseteq X$.

Let h be as in Lemma 2.10. Then $g = f \circ h$ is a measurable function that agrees with f on D and

$$\lim_{x \rightarrow x_0} g(x) = f\left(\lim_{x \rightarrow x_0} h(x)\right) = f(x_0) = g(x_0)$$

for all $x_0 \in D$. $\square$

Remark 2.11. Let S and T be computable Polish spaces. A set $X \subseteq S$ is an effective G_δ set (or Π_2^0 class) when it is the intersection of a uniformly computable sequence of c.e. open sets. 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 some such R' can be computed uniformly from a code for the witness $\{U_n\}_{n \in \mathbb{N}}$. For details, see Hoyrup (2008, Thm. 1.6.2.1); this generalizes a classical result of Kuratowski (Kechris 1995, I.3.B, Thm. 3.8). In conclusion, one can always assume that the set R is an effective G_δ set.

We will introduce a weaker notion of computability for functions in Section 2.6.

2.4 Computable and Almost 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 $(\{0, 1\}^\infty, \mathcal{F}, \mathbb{P})$, 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 $\mathbb{P}$ is the product measure formed from the uniform distribution on $\{0, 1\}$. Throughout the rest of the article, we will take $(\{0, 1\}^\infty, \mathcal{F}, \mathbb{P})$ to be the basic probability space. We will use a SANS SERIF font for random variables.

Definition 2.12 (Random Variable and its Distribution). Let S be a Polish space. A **random variable in S** is a measurable function $X : \{0, 1\}^\infty \rightarrow S$. For a measurable subset $A \subseteq S$, we let $\{X \in A\}$ denote the inverse image $X^{-1}[A] = \{\omega \in \{0, 1\}^\infty : X(\omega) \in A\}$, and for $x \in S$, we similarly define the event $\{X = x\}$. We will write $\mathbb{P}_X$ for the **distribution of X** , which is the measure on S defined by $\mathbb{P}_X(\cdot) := \mathbb{P}\{X \in \cdot\}$.

If S is a computable Polish space, then we say a random variable X in S is a **$\mathbb{P}$ -almost computable random variable** if it is $\mathbb{P}$ -almost computable as a measurable function. Intuitively, X is a $\mathbb{P}$ -almost 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 $\mathbb{P}$ -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 $\mathbb{P}$ -almost computable random variables with *continuous* distributions, such as a uniform random variable (gotten by subdividing the unit interval according to the random bit tape) or an i.i.d. sequence of uniformly distributed random variables (by splitting up the given element of $\{0, 1\}^\infty$ into countably many disjoint subsequences and dovetailing the constructions). (For explicit constructions, see, e.g., Freer (2010, Ex. 3, 4).)

It is crucial that we consider random variables that are merely computable on a $\mathbb{P}$ -measure one subset of $\{0, 1\}^\infty$. To see why, consider the following example, which was communicated to us by Martín Escardó. 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 $\mathbb{P}_X\{1\} = \alpha$. 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 whose denominator is 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., $\frac{1}{3}$).

LEMMA 2.13. *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 $\mathbb{P}(Z_1) = 1 - \mathbb{P}(Z_0) = \alpha$. $\square$

However, 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 Mann (1973), 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 $\mathbb{P}$ -almost 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 (1976). Hence, it is natural to focus our attention on random variables that are merely $\mathbb{P}$ -almost computable.

The setting of $\mathbb{P}$ -almost computable random variables is a natural one for probability theory, and the standard operations on random variables preserve $\mathbb{P}$ -almost computability, including, e.g., addition and multiplication of $\mathbb{P}$ -almost computable real random variables, composition with $\mathbb{P}$ -almost computable measurable functions, and Cartesian products.

2.5 Computable Probability Measures

We now introduce the class of computable probability measures on computable Polish spaces.

Let $(S, \delta_S, \mathcal{D}_S)$ be a computable Polish space, and recall that $\mathcal{B}_S$ denotes its Borel sets and $\mathcal{M}_1(S)$ its Borel probability measures. Consider the subset $\mathcal{D}_{P,S} \subseteq \mathcal{M}_1(S)$ 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., $\nu \in \mathcal{D}_{P,S}$ if and only if $\nu = q_1 \delta_{t_1} + \dots + q_k \delta_{t_k}$ for some rationals $q_i \geq 0$ such that $q_1 + \dots + q_k = 1$ and some points $t_i \in \mathcal{D}_S$, where for $t \in S$ the $\{0, 1\}$ -valued Dirac measure δ_t satisfies $\delta_t(A) = 1$ if and only if $t \in A$ for all measurable sets A . It is a standard fact (see, e.g., Gács (2005, §B.6.2)) that $\mathcal{D}_P$ is dense in the Prokhorov metric δ_P given by

$$\delta_P(\mu, \nu) := \inf \{ \varepsilon > 0 : \forall A \in \mathcal{B}_S, \mu(A) \leq \nu(A^\varepsilon) + \varepsilon \}, \quad (2)$$

where

$$A^\varepsilon := \{ p \in S : \exists q \in A, \delta_S(p, q) < \varepsilon \} = \bigcup_{p \in A} B_\varepsilon(p) \quad (3)$$

is the ε -neighborhood of A and $B_\varepsilon(p)$ is the open ball of radius ε about p . Moreover, $(\mathcal{M}_1(S), \delta_P, \mathcal{D}_{P,S})$ is a computable Polish space. (See Hoyrup and Rojas (2009b, Prop. 4.1.1)). We say that $\mu \in \mathcal{M}_1(S)$ is a computable probability measure when μ is a computable point in $\mathcal{M}_1(S)$ as a computable Polish space. Note when the space S is clear from context, we will refer to $\mathcal{D}_{P,S}$ simply as $\mathcal{D}_P$.

One can define computability on the space of probability measures in other natural ways. Early work by Weihrauch (1999) and Müller (1999) formalized the computability of probability measure in terms of the lower semicomputability of the measure as a function on the set of open sets and in terms of the computability of the measure as a linear operator acting on bounded continuous functions; these notions are equivalent. (See Schröder (2007) for a more general setting.) These notions of computability also agree with the notion of computability defined here in terms of the Prokhorov metric.

PROPOSITION 2.14 (HOYRUP AND ROJAS (2009b, THM. 4.2.1)). *Let S be a computable Polish 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 .*

Note that the measure $\mathbb{P}$ on $\{0, 1\}^\infty$ is a computable probability measure.

We can also characterize the class of computable probability measures in terms of the uniform computability of the integrals of bounded continuous functions:

PROPOSITION 2.15 (HOYRUP AND ROJAS (2009b, COR. 4.3.1)). *Let S be a computable Polish space, let μ be a probability measure on S , and let $\mathcal{F}$ be the set of computable functions from S to $\mathbb{R}^+$. Then μ is computable if and only if $\int f d\mu$ is a c.e. real, uniformly in $f \in \mathcal{F}$.*

COROLLARY 2.16. *Let S be a computable Polish space, let μ be a probability measure on S , and let $\mathcal{F}$ be the set of computable functions from S to $[0, 1]$. Then μ is computable if and only if $\int f d\mu$ is computable, uniformly in $f \in \mathcal{F}$.*

PROOF. First observe that both f and $1 - f$ are non-negative functions. Therefore, by Proposition 2.15, the reals $\int f d\mu$ and $\int (1 - f) d\mu$ are both c.e., and hence the real $\int f d\mu$ is computable. $\square$

Having explained the computability of probability measures in terms of integration, we now relate it to the computability of random variables defined on computable Polish spaces.

Definition 2.17 (Computable Probability Space (Galatolo et al. 2010, Def. 2.4.1)). A **computable probability space** is a pair (S, μ) where S is a computable Polish space and μ is a computable probability measure on S .

The distribution of a $\mathbb{P}$ -almost computable random variable in a computable Polish space is computable.

PROPOSITION 2.18 (GALATOLO ET AL. (2010, PROP. 2.4.2)). *Let X be a $\mathbb{P}$ -almost computable random variable in a computable Polish space S . Then its distribution is a computable point in the computable Polish space $\mathcal{M}_1(S)$.*

However, given a computable measure, there is a $\mathbb{P}$ -almost computable random variable with that distribution.

PROPOSITION 2.19 (HOYRUP AND ROJAS (2009B, THM. 5.1.1)). *Let μ be a computable probability measure on a computable Polish space S . Then there is a $\mathbb{P}$ -almost computable random variable in S whose distribution is μ .*

In summary, the computable probability measures on a computable Polish space are precisely the distributions of $\mathbb{P}$ -almost computable random variables in that space. For this result in a more general setting, see Schröder (2007, Prop. 4.3).

Further, if μ is a computable probability measure and f is computable on a μ -measure one set, then the pushforward $\mu \circ f^{-1}$ is a computable distribution. This fact, along with Proposition 2.19, shows that we have lost no generality in taking $(\{0, 1\}^\infty, \mathcal{F}, \mathbb{P})$ to be our basic probability space.

All of the standard distributions (e.g., normal, uniform, geometric, exponential) found in probability textbooks, and then all the transformations of these distributions by $\mathbb{P}$ -almost computable functions, are easily shown to be computable distributions.

2.6 Weaker Notions of Computability for Functions on Probability Spaces

Another important class of functions on a probability space is the class of L^1 -computable functions. For more details, including some of the history of L^1 -computability, see Hoyrup and Rojas (2009a, §3.1) and Miyabe (2013).

Definition 2.20 (The Metric Space of $L^1(\mu)$ functions (Hoyrup and Rojas 2009a, §3.1)). Let μ be a probability measure on a Polish space S , and let $\mathcal{F}$ be the set of μ -integrable functions from S to $\mathbb{R}$. Then $\delta(f, g) := \int |f - g| d\mu$ is a metric on the quotient space of $\mathcal{F}$ defined by the equivalence relation $f \sim g$ iff $\int |f - g| d\mu = 0$. This metric space is called the **space of $L^1(\mu)$ functions on S** , and we will often speak interchangeably of a μ -integrable function $S \rightarrow \mathbb{R}$ and its equivalence class.

We will make use of the following set of L^1 functions:

Definition 2.21 (Ideal Points for L^1 (Gács 2005, §2)). Let $(S, \delta, \mathcal{D})$ be a computable Polish space. Define $\mathcal{E}$ to be the smallest set of functions containing the constant function 1 and the functions $\{g_{u,r,1/n} : u \in S, r \in \mathbb{Q}, n \geq 1\}$, where

$$g_{u,r,\epsilon}(x) := \max(0, 1 - \max(0, \delta(x, u) - r)/\epsilon), \quad (4)$$

and closed under max, min, and rational linear combinations.

Such functions can be thought of as continuous analogues of step functions having a finite number of steps, each step of which corresponds to a basic open ball with rational radius and ideal center.

LEMMA 2.22 (HOYRUP AND ROJAS (2009A, PROP. 3)). *Let μ be a computable probability measure on a computable Polish space $(S, \delta, \mathcal{D})$. The set $\mathcal{E}$ is dense in the $L^1(\mu)$ functions on S , and the distances between points in $\mathcal{E}$ are computable under the standard enumeration, making this space into a computable Polish space.*

We say that an $L^1(\mu)$ function on a computable Polish space S is **$L^1(\mu)$ -computable** when it is a computable point in the $L^1(\mu)$ functions on S .

LEMMA 2.23 (HOYRUP AND ROJAS (2009A, THM. 4 CLAIM 2 AND THM. 5 CLAIM 2)). *Let (S, μ) be a computable probability space and let T be a computable Polish space. A function $f : S \rightarrow T$ is $L^1(\mu)$ -computable if and only if $\int f d\mu$ is a computable real and for each $r \in \mathbb{N}$, the function f is computable on some set of $\mathbb{P}_X$ -measure at least $1 - 2^{-r}$, uniformly in r .*

In particular, note that every integrable μ -almost computable function is $L^1(\mu)$ -computable.

We obtain the following immediate corollary of Lemma 2.23 using the fact that if a function is μ -almost computable with a computable μ -integral, then we can uniformly find a collection of ideal points that converge to it in $L^1(\mu)$:

COROLLARY 2.24. *Let (S, μ) be a computable probability space and let T be a computable Polish space. Let $f_0, f_1, \dots : S \rightarrow T$ be a sequence of uniformly μ -almost computable functions taking values in a computable Polish space T that converge effectively in $L^1(\mu)$ to a function $f \in L^1(\mu)$. Then f is $L^1(\mu)$ -computable.*

2.7 Almost Decidable Sets and Bases

Let (S, μ) be a computable probability space. We know that the μ -measure of a c.e. open set $A \subseteq S$ is a c.e. real. In general, the measure of a c.e. open set is not a computable real. However, if A is a decidable subset (i.e., $S \setminus A$ is c.e. open), then $\mu(S \setminus A)$ is 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 2.25 (Almost Decidable Set (Galatolo et al. 2010, Def. 3.1.3)). Let (S, μ) be a computable probability space. A 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 $\mu(U) + \mu(V) = 1$. In this case, we say that (U, V) witnesses the μ -almost decidability of A .

The following lemma is immediate.

LEMMA 2.26 (GALATOLO ET AL. (2010, 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 (2008, Lem. 2.15)).

LEMMA 2.27 (GALATOLO ET AL. (2010, 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, which we call a **μ -almost decidable basis**.*

We now show that every c.e. open set of a computable probability space (S, μ) is the union of a computable sequence of μ -almost decidable subsets.

LEMMA 2.28 (ALMOST DECIDABLE SUBSETS). *Let (S, μ) be a computable probability space with ideal points $\{d_i\}_{i \in \mathbb{N}}$, and let $\{r_j\}_{j \in \mathbb{N}}$ be a computable sequence of reals such that $\{B(d_i, r_j)\}_{i, j \in \mathbb{N}}$ is a μ -almost decidable basis. Let V be a c.e. open set. Then, uniformly in $\{r_j\}_{j \in \mathbb{N}}$ and V , we can compute a sequence of μ -almost decidable sets $\{V_k\}_{k \in \mathbb{N}}$ such that $V_k \subseteq V_{k+1}$ for each k , 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 $V = \bigcup_{k \in E} B_k$. 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 computable enumeration of F and let $V_k := \bigcup_{n \leq k} B(d_{i_n}, r_{j_n})$, which is μ -almost decidable. By construction, $V_k \subseteq V_{k+1}$ for each k , 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 2.29. *Let (S, μ) be a computable probability space with ideal points $\{d_i\}_{i \in \mathbb{N}}$, and let $\{r_j\}_{j \in \mathbb{N}}$ be a computable sequence of reals such that $\{B(d_i, r_j)\}_{i,j \in \mathbb{N}}$ is a μ -almost decidable basis. Let $\nu \in \mathcal{M}_1(S)$ be a probability measure on S that is absolutely continuous with respect to μ . Then ν is computable uniformly in the sequence $\{\nu(B(d_i, r_j))\}_{i,j \in \mathbb{N}}$.*

PROOF. Let V be a c.e. open set of S . By Proposition 2.14, it suffices to show that $\nu(V)$ is a c.e. real, uniformly in V . By Lemma 2.28, we can compute a nested sequence $\{V_k\}_{k \in \mathbb{N}}$ of μ -almost decidable sets whose union is V . By the absolute continuity of ν with respect to μ , these sets are also ν -almost decidable. Because V is open, $\nu(V) = \sup_{k \in \mathbb{N}} \nu(V_k)$, which is the supremum of a sequence of reals that is computable uniformly in the sequence $\{\nu(B(d_i, r_j))\}_{i,j \in \mathbb{N}}$. $\square$

We close with the following extension of Corollary 2.16:

PROPOSITION 2.30. *Let S, T be computable Polish spaces, μ a probability measure on T , B a μ -almost decidable subset of $\mathbb{R}$, and $f : S \times T \rightarrow \mathbb{R}$ a bounded function, computable on $R \times T$ with $R \subseteq S$. Then the map $s \mapsto \int_B f(s, t) \mu(dt)$ is a computable function, uniformly in f and B .*

PROOF. This follows immediately from Propositions 3.2.3 and 4.3.1 of Hoyrup and Rojas (2009b). $\square$

3 CONDITIONAL PROBABILITIES AND DISTRIBUTIONS

Let μ be a probability measure on a measurable space of outcomes S , and let $A, B \subseteq S$ be events. Informally, given that event A has occurred, the probability that event B also occurs, written $\mu(B|A)$, must satisfy $\mu(A) \mu(B|A) = \mu(A \cap B)$. Clearly $\mu(B|A)$ is uniquely defined if and only if $\mu(A) > 0$, which leads to the following definition:

Definition 3.1 (Conditioning on Positive-measure Events). Suppose that $\mu(A) > 0$. Then the **conditional probability of B given A** , written $\mu(B|A)$, is defined by

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

It is straightforward to check that, for any fixed event $A \subseteq S$ with $\mu(A) > 0$, the set function $\mu(\cdot | A)$ is a probability measure.

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

$$\mathbb{P}\{Y \in D \mid X \in C\} := \mathbb{P}\left(\{Y \in D\} \middle| \{X \in C\}\right). \quad (7)$$

Again, this is well defined when $\mathbb{P}\{X \in C\} > 0$. When $\mathbb{P}\{X = x\} > 0$, we may simply write

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

for $\mathbb{P}\{Y \in D \mid X \in \{x\}\}$.

This elementary notion of conditioning is undefined when the conditioning event has zero measure, such as when a *continuous* random variable takes a particular value. In the modern formulation of conditional probability due to Kolmogorov (1933), one defines conditioning with respect to (the σ -algebra generated by) a *random variable* rather than an individual event. In theory, this yields a consistent solution to the problem of conditioning on the value of general (and in particular, continuous) random variables, although we will see that other issues arise. (See Kallenberg (2002, Ch. 6) for a rigorous treatment.)

To bridge the divide between the elementary notion of conditioning on events and the abstract approach of conditioning on random variables, consider the case of conditioning on a random variable X taking values in a countable discrete set S and satisfying $\mathbb{P}\{X = x\} > 0$ for all $x \in S$. Let $\{Y \in B\}$ be an event. Then the conditional probability that $Y \in B$ given X , written $\mathbb{P}[Y \in B|X]$, is the *random variable* satisfying $\mathbb{P}[Y \in B|X] = \mathbb{P}\{Y \in B \mid X = x\}$ when $X = x$. Note that there is a measurable function $f_B : S \rightarrow [0, 1]$ satisfying

$$\mathbb{P}\{Y \in B, X \in A\} = \int_A f_B(x) \mathbb{P}_X(dx) \quad (9)$$

for all measurable subsets $A \subseteq S$. For sets A of the form $\{x\}$, for $x \in S$, we have $\mathbb{P}\{Y \in B, X = x\} = f_B(x)\mathbb{P}\{X = x\}$, hence $f_B(x) = \mathbb{P}\{Y \in B|X = x\}$. In summary, $\mathbb{P}[Y \in B|X] = f_B(X)$, and so Equation (9) yields a more abstract characterization of elementary conditional probability for positive-measure events.

The general case is captured by the same defining property. Let X be a random variable in a measurable space S . Then the **conditional probability that $Y \in B$ given X** , written $\mathbb{P}[Y \in B|X]$, is defined to be a random variable in $[0, 1]$ of the form $f_B(X)$ where again $f_B : S \rightarrow [0, 1]$ is such that Equation (9) holds for all measurable subsets $A \subseteq S$. In many situations, such a function f_B is itself the object of interest and so we will let $\mathbb{P}[Y \in B|X = \cdot]$ denote an arbitrary such function. We may then re-express its defining property in the following more intuitive form:

$$\mathbb{P}\{Y \in B, X \in A\} = \int_A \mathbb{P}[Y \in B|X = x] \mathbb{P}_X(dx) \quad (10)$$

for all measurable subsets $A \subseteq S$.

The existence of the conditional probability $\mathbb{P}[Y \in B|X]$, or equivalently, the existence of $\mathbb{P}[Y \in B|X = \cdot]$, follows from the Radon–Nikodym theorem. Recall that a measure μ on a measurable space S is **absolutely continuous** with respect to another measure ν on the same space, written $\mu \ll \nu$, if $\nu(A) = 0$ implies $\mu(A) = 0$ for all measurable sets $A \subseteq S$.

THEOREM 3.2 (RADON–NIKODYM). *Let S be a measurable space and let μ and ν be σ -finite measures on S such that $\mu \ll \nu$. Then there exists a nonnegative measurable function $\frac{d\mu}{d\nu} : S \rightarrow \mathbb{R}_+$ such that*

$$\mu(A) = \int_A \frac{d\mu}{d\nu} d\nu \quad (11)$$

for all measurable subsets $A \subseteq S$.

We call any function $\frac{d\mu}{d\nu}$ satisfying Equation (11) for all measurable subsets $A \subseteq S$ a **Radon–Nikodym derivative** (of μ with respect to ν).

Note that if g is also a Radon–Nikodym derivative of μ with respect to ν , then $g = \frac{d\mu}{d\nu}$ outside a ν -null set, and so *Radon–Nikodym derivatives are unique up to a null set*. (Functions that agree a.e. are called **versions**.) We may safely refer to *the* Radon–Nikodym derivative when we want to ignore such differences, but in some cases these differences are important.

It is straightforward to verify that the function $\mathbb{P}[Y \in B|X = \cdot]$ is a Radon–Nikodym derivative of $\mathbb{P}\{Y \in B, X \in \cdot\}$ with respect to $\mathbb{P}\{X \in \cdot\} = \mathbb{P}_X$, both considered as measures on S , and so a function f_B satisfying Equation (9) for all measurable subsets $A \subseteq S$ always exists, but it is only defined up to a null set. This is inconsequential when the conditional probability $\mathbb{P}[Y \in B|X]$ is the object of interest. In applications, especially statistical ones, however, the function $\mathbb{P}[Y \in B|X = \cdot]$ mapping values in S to probabilities is the object of interest, and, moreover, one typically wants to evaluate this function at particular observed values $x \in S$. Because $\mathbb{P}[Y \in B|X = \cdot]$ is merely determined up to a $\mathbb{P}_X$ -null set, interpreting its values at individual points is problematic.

As mentioned in the introduction, the fact that general conditional probabilities are not uniquely defined at points is the subject of a large body of literature. However, in some circumstances, two versions of $\mathbb{P}[Y \in B|X = \cdot]$ must agree at individual points. In particular, if two versions are continuous at a point in the support of the distribution $\mathbb{P}_X$, then they agree on the value at that point. To state this claim formally, we first recall the definition of the support of a distribution:

Definition 3.3. 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.,

$$\text{supp}(\mu) := \{x \in S : \forall B \in \mathcal{S} (x \in B \Rightarrow \mu(B) > 0)\}. \quad (12)$$

Note that the support of μ can equivalently be defined as the smallest closed set of μ -measure one. We now state our claim formally:

LEMMA 3.4. Let S be a Polish space. Suppose $f_1, f_2 : S \rightarrow [0, 1]$ satisfy $\mathbb{P}[Y \in B|X] = f_1(X) = f_2(X)$ a.s. If $x \in S$ is a point of continuity of f_1 and f_2 , and $x \in \text{supp}(\mathbb{P}_X)$, then $f_1(x) = f_2(x)$. In particular, if f_1 and f_2 are continuous on a $\mathbb{P}_X$ -measure one set $D \subseteq S$, then they agree everywhere in $D \cap \text{supp}(\mathbb{P}_X)$.

The proof is immediate from the following elementary result:

LEMMA 3.5. Let $f_1, f_2 : S \rightarrow T$ be two measurable functions between Polish spaces S and T , and suppose that $f_1 = f_2$ almost everywhere with respect to some measure μ on S . Let $D \subseteq S$ be a set of μ -measure one. If $x \in S$ is a point of continuity of f_1 and f_2 on D , and $x \in \text{supp}(\mu)$, then $f_1(x) = f_2(x)$. In particular, if f_1 and f_2 are continuous on D , then they agree everywhere in $D \cap \text{supp}(\mu)$.

PROOF. Let δ_T be any metric under which T is complete. Define the measurable function $g : S \rightarrow \mathbb{R}$ by

$$g(x) = \delta_T(f_1(x), f_2(x)). \quad (13)$$

We know that $g = 0$ μ -a.e., and also that g is continuous at x on D , because f_1 and f_2 are continuous at x on D and δ_T is continuous (on all of T). Assume, for the purpose of contradiction, that $g(x) = \varepsilon > 0$. By the continuity of g on D , there is an open neighborhood B of x such that $g(B \cap D) \subseteq (\frac{\varepsilon}{2}, \frac{3\varepsilon}{2})$. But $x \in \text{supp}(\mu)$, hence $\mu(B \cap D) = \mu(B) > 0$, contradicting $g = 0$ μ -a.e. $\square$

The observation that continuity gives a unique answer to conditioning on zero-measure events of the form $\{X = x\}$ is an old one, going back to at least Tjur (1974).

3.1 Conditional Distributions

For a pair of random variables X and Y taking values in a pair of measurable space S and T , respectively, it is natural to consider not just individual conditional probabilities $\mathbb{P}[Y \in B|X]$, for measurable subsets $B \subseteq T$, but the entire **conditional distribution** $\mathbb{P}[Y|X] := \mathbb{P}[Y \in \cdot | X]$. Unfortunately, the fact that Radon–Nikodym derivatives are only defined up to a null set can cause problems. In particular, while it is the case that

$$\sum_j \mathbb{P}[Y \in B_j|X] = \mathbb{P}[Y \in B|X] \text{ a.s.} \quad (14)$$

for every countable measurable partition $B_0, B_1, \dots$ of a measurable set $B \subseteq T$, the random set function given by $B \mapsto \mathbb{P}[Y \in B|X]$ need not be a measure in general because the exceptional null set may depend on the sequence. However, when T is Polish, we can construct versions of the conditional probabilities that combine to produce a measure. To make this definition precise, we recall the notion of a probability kernel.

Definition 3.6 (Probability Kernel). Let S and T be Polish 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.

For every $\kappa : S \times \mathcal{B}_T \rightarrow [0, 1]$, let $\bar{\kappa}$ be the map $s \mapsto \kappa(s, \cdot)$. It can be shown that κ is a probability kernel from S to T if and only if $\bar{\kappa}$ is a (Borel) measurable function from S to $\mathcal{M}_1(T)$ (Kallenberg 2002, Lem. 1.40), where we adopt the weak topology on $\mathcal{M}_1(T)$, which is Polish because T is.

We say that a conditional distribution $\mathbb{P}[Y|X]$ **has a regular version** when, for some probability kernel κ from S to T ,

$$\mathbb{P}[Y \in B|X] = \kappa(X, B) \text{ a.s.} \quad (15)$$

for every measurable subset $B \subseteq T$. In this case, we would say that $\bar{\kappa}(X)$ is a regular version of the conditional distribution.

PROPOSITION 3.7 (REGULAR VERSIONS (KALLENBERG 2002, LEM. 6.3)). *Let X and Y be random variables in a Polish space S and a measurable space T , respectively. Then there is a regular version of the conditional distribution $\mathbb{P}[Y|X]$, which is, moreover, determined by the joint distribution of X and Y .*

As with the derivatives underlying conditional probabilities, $\bar{\kappa}$ is only defined up to a $\mathbb{P}_X$ -null set. When such a kernel κ exists, i.e., when there is a regular version of the conditional distribution $\mathbb{P}[Y|X]$, we define $\mathbb{P}[Y|X = \cdot]$ to be equal to some arbitrary version of $\bar{\kappa}$.

Despite the fact that the kernels underlying regular versions of conditional distributions are defined only up to sets of measure zero, it follows immediately from Lemma 3.5 that when S and T are Polish, any two versions of $\mathbb{P}[Y|X = \cdot]$ that are continuous on some subset of the support of $\mathbb{P}_X$ must agree on that subset. More carefully, let $\bar{\kappa}_1(X)$ and $\bar{\kappa}_2(X)$ be regular versions of the conditional distribution $\mathbb{P}[Y|X]$. If $x \in S$ is a point of continuity of $\bar{\kappa}_1$ and $\bar{\kappa}_2$, and $x \in \text{supp}(\mathbb{P}_X)$, then $\bar{\kappa}_1(x) = \bar{\kappa}_2(x)$. In particular, if both maps are continuous on a set $D \subseteq S$, then they agree everywhere in $D \cap \text{supp}(\mathbb{P}_X)$.

When conditioning on a random variable whose distribution concentrates on a countable set, it is well known that a regular version of the conditional distribution can be built by elementary conditioning with respect to single events. This includes the special case of conditioning on **discrete random variables**, i.e., those concentrating on a countable discrete subspace.

LEMMA 3.8. *Let X and Y be random variables in Polish spaces S and T , respectively. Suppose the distribution of X concentrates on a countable set $R \subseteq S$, i.e., $\mathbb{P}_X(R) = 1$ and $x \in R$ implies $\mathbb{P}_X\{x\} > 0$. Let ν be an arbitrary probability measure on T . Define the function $\kappa : S \times \mathcal{B}_T \rightarrow [0, 1]$ by*

$$\kappa(x, B) := \mathbb{P}\{Y \in B \mid X = x\} \quad (16)$$

for all $x \in R$ and $\bar{\kappa}(x) = \nu$ for $x \notin R$. Then κ is a probability kernel and $\bar{\kappa}(X)$ is a regular version of the conditional distribution $\mathbb{P}[Y|X]$.

PROOF. The function κ is well defined because $\mathbb{P}\{X = x\} > 0$ for all $x \in R$. It follows that $\bar{\kappa}(x)$ is a probability measure for every x . Because R is countable, $\bar{\kappa}$ is also measurable and so κ is a

probability kernel from S to T . Note that $\mathbb{P}\{X \in R\} = 1$ and so, for all measurable sets $A \subseteq S$ and $B \subseteq T$, we have

$$\int_A \kappa(x, B) \mathbb{P}_X(dx) = \sum_{x \in R \cap A} \mathbb{P}\{Y \in B \mid X = x\} \mathbb{P}\{X = x\} \quad (17)$$

$$= \sum_{x \in R \cap A} \mathbb{P}\{Y \in B, X = x\} \quad (18)$$

$$= \mathbb{P}\{Y \in B, X \in A\}. \quad (19)$$

That is, $\kappa(X, B)$ is the conditional probability of the event $\{Y \in B\}$ given X , and so $\bar{\kappa}(X)$ is a regular version of the conditional distribution $\mathbb{P}[Y|X]$. $\square$

3.2 Dominated Families

Beyond the setting of conditioning on discrete random variables, explicit formulas for conditional distributions are also available when Bayes' rule applies. We begin by introducing the notion of a dominated kernel. (The usual terms, such as dominated families or models, refers to measurable families of probability measures, i.e., probability kernels.)

Definition 3.9 (Dominated Kernel). A probability kernel κ from T to S is **dominated** when there is a σ -finite measure ν on S such that $\bar{\kappa}(t) \ll \nu$ for every $t \in T$.

Let X and Y be random variables in Polish spaces S and T , respectively, and let $\bar{\kappa}_{X|Y}(X)$ be a regular version of $\mathbb{P}[X|Y]$ such that $\kappa_{X|Y}$ is dominated. Then there exists a (product) measurable function $p_{X|Y}(\cdot | \cdot) : S \times T \rightarrow \mathbb{R}^+$ such that $p_{X|Y}(\cdot | y)$ is a Radon–Nikodym derivative of $\bar{\kappa}_{X|Y}(y)$ with respect to ν for every $y \in T$, i.e.,

$$\kappa_{X|Y}(y, A) = \int_A p_{X|Y}(x|y) \nu(dx) \quad (20)$$

for every measurable set $A \subseteq S$ and every $y \in T$.

Definition 3.10 (Conditional Density). We call any such function $p_{X|Y}$ a **conditional density of X given Y (with respect to ν)**.

Common finite-dimensional, parametric families of distributions (e.g., exponential families such as Gaussian, gamma, etc.) are dominated, and so, in probabilistic models composed from these families, conditional densities exist and Bayes' rule gives a formula for expressing the conditional distribution. We give a proof of this classic result for completeness.

LEMMA 3.11 (BAYES' RULE (SCHERVISH 1995, THM. 1.13)). *Let X and Y be random variables as in Proposition 3.7, and assume that there exists a conditional density $p_{X|Y}$ of X given Y with respect to a σ -finite measure ν . Let μ be an arbitrary distribution on T and define $\kappa : S \times \mathcal{B}_T \rightarrow [0, 1]$ by*

$$\kappa(x, B) = \frac{\int_B p_{X|Y}(x|y) \mathbb{P}_Y(dy)}{\int p_{X|Y}(x|y) \mathbb{P}_Y(dy)}, \quad B \in \mathcal{B}_T, \quad (21)$$

for those points $x \in S$ where the denominator is positive and finite, and by $\bar{\kappa}(x) = \mu$ otherwise. Then κ is a probability kernel and $\bar{\kappa}(X)$ is a regular version of the conditional distribution $\mathbb{P}[Y|X]$.

PROOF. Let $\bar{\kappa}_{X|Y}(Y)$ be a regular version of the conditional distribution $\mathbb{P}[X|Y]$. By hypothesis, $\kappa_{X|Y}$ is dominated by ν and $p_{X|Y}$ is a conditional density with respect to ν . By Proposition 3.7 and Fubini's theorem, for measurable sets $A \subseteq S$ and $B \subseteq T$, we have that

$$\mathbb{P}\{X \in A, Y \in B\} = \int_B \kappa_{X|Y}(y, A) \mathbb{P}_Y(dy) \quad (22)$$

$$= \int_B \left(\int_A p_{X|Y}(x|y) \nu(dx) \right) \mathbb{P}_Y(dy) \quad (23)$$

$$= \int_A \left(\int_B p_{X|Y}(x|y) \mathbb{P}_Y(dy) \right) \nu(dx). \quad (24)$$

Taking $B = T$, we have

$$\mathbb{P}_X(A) = \int_A \left(\int p_{X|Y}(x|y) \mathbb{P}_Y(dy) \right) \nu(dx). \quad (25)$$

Because $\mathbb{P}_X(S) = 1$, this implies that the set of points x for which the denominator of the right-hand side of Equation (21) is infinite has ν -measure zero, and thus $\mathbb{P}_X$ -measure zero. Taking A to be the set of points x for which the denominator is zero, we see that $\mathbb{P}_X(A) = 0$. It follows that Equation (21) characterizes κ up to a $\mathbb{P}_X$ -null set.

By Equation (25), we see that the denominator is a density of $\mathbb{P}_X$ with respect to ν , and so we have

$$\int_A \kappa(x, B) \mathbb{P}_X(dx) = \int_A \kappa(x, B) \left(\int p_{X|Y}(x|y) \mathbb{P}_Y(dy) \right) \nu(dx), \quad (26)$$

for all measurable sets $A \subseteq S$ and $B \subseteq T$. Finally, by the definition of κ , Equation (24), and the fact that the denominator is positive and finite for $\mathbb{P}_X$ -almost every x , we see that $\bar{\kappa}(X)$ is a regular version of the conditional distribution $\mathbb{P}[Y|X]$. $\square$

Comparing Bayes' rule in Equation (21) to the definition of conditional density in Equation (20), we see that any conditional density of Y given X (with respect to $\mathbb{P}_Y$) satisfies

$$p_{Y|X}(y|x) = \frac{p_{X|Y}(x|y)}{\int p_{X|Y}(x|y) \mathbb{P}_Y(dy)}, \quad (27)$$

for $\mathbb{P}_{(X,Y)}$ -almost every (x, y) .

The following result suggests why the mere a.e. definedness of conditional distributions can be ignored by those working entirely within the framework of dominated families.

PROPOSITION 3.12. *Let X and Y be random variables on Polish spaces S and T , respectively, let $\bar{\kappa}(X)$ be a regular version of the conditional distribution $\mathbb{P}[Y|X]$, 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 κ as defined in Equation (21) is a version of $\bar{\kappa}$ that is continuous on R . In particular, if R is a $\mathbb{P}_X$ -measure one subset, then κ is a $\mathbb{P}_X$ -almost continuous version.*

We defer the proof to Section 9.2. We will use this result in the proof of Lemma 7.3, towards our central result.

4 COMPUTABLE CONDITIONAL PROBABILITIES AND DISTRIBUTIONS

Before we lay the foundations for the remainder of the article and define notions of computability for conditional probability and conditional distributions in the abstract setting, we address the computability of distributions conditioned on positive-measure sets. In order for the distributions obtained from positive measure sets to be computable, we will need the conditioning events to be almost decidable sets.

LEMMA 4.1 (GALATOLO ET AL. (2010, PROP. 3.1.2)). *Let (S, μ) be a computable probability space and let A be a μ -almost decidable subset of S satisfying $\mu(A) > 0$. Then $\mu(\cdot | A)$ is a computable probability measure, uniformly in a witness to the μ -almost decidability of A .*

PROOF. By Lemma 2.27 there is a μ -almost decidable basis for S . Note that $\mu(\cdot | A)$ is absolutely continuous with respect to μ . Hence, by Corollary 2.29, it suffices to show that $\frac{\mu(B \cap A)}{\mu(A)}$ is computable for a μ -almost decidable set B , uniformly in witnesses to the μ -almost decidability of A and B . All subsequent statements in this proof are uniform in both. Now, $B \cap A$ is μ -almost decidable with computable witness, and so its measure, the numerator, is a computable real. The denominator is likewise the measure of a set that is almost decidable with computable witness, hence is a computable real. Finally, the ratio of two computable reals is itself computable. $\square$

In the abstract setting, conditional probabilities are random variables. In many applications of probability, including statistics, the conditional probability map, or some version of it, is the actual object of interest, and so the computability of this map is our focus.

Let $B \subseteq T$ be a measurable set. Viewing $\mathbb{P}[Y \in B | X = \cdot]$ as a function from S to $[0, 1]$, recall that we can speak formally as to whether this function is everywhere computable, $\mathbb{P}_X$ -almost computable, and/or L^1 -computable. Recall also that the function $\mathbb{P}[Y \in B | X = \cdot]$ may have many versions that agree only up to a null set. Despite this, their almost computability does not differ (up to a change in domain by a null set).

LEMMA 4.2. *Let f be a measurable function from a computable probability space (S, μ) to a computable Polish space T . If any version of f is computable on a μ -measure p set, then every version of f is computable on a μ -measure p set. In particular, if one version is μ -almost computable, then all version are.*

PROOF. Let f be computable on a μ -measure p set D , and let g be a version of f , i.e., $Z := \{s \in S : f(s) \neq g(s)\}$ is a μ -null set. Therefore, $f = g$ on $D \setminus Z$. Hence, g is computable on the μ -measure p set $D \setminus Z$. If f is μ -almost computable, then it is computable on a μ -measure one set, and so g is as well. $\square$

We can develop notions of computability for conditional distributions in a similar way. We begin by characterizing the computability of probability kernels.

Definition 4.3 (Computable Probability Kernel). Let S and T be computable Polish 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 $\bar{\kappa} : S \rightarrow \mathcal{M}_1(T)$ given by $\bar{\kappa}(s) := \kappa(s, \cdot)$ is a computable function in the ordinary sense between S and the computable Polish space $\mathcal{M}_1(T)$ induced by T . Similarly, we say that κ is computable on a subset $D \subseteq S$ when $\bar{\kappa}$ is computable on D .

As we will see, this notion of computability corresponds with a more direct notion of computability for κ , which we now develop. We begin by noting that the collection of sets of the form

$$P_T(A, q) := \{\mu \in \mathcal{M}_1(T) : \mu(A) > q\} \quad (28)$$

for A open and q rational, form a subbasis for the weak topology on $\mathcal{M}_1(T)$ (which is the topology induced by the Prokhorov metric). Indeed, it suffices for A to range over finite unions of some countable basis of T . We will also omit mention of T when the ambient space is clear from context.

The next result relates balls in the Prokhorov metric to the subbasis elements above. Recall that δ_p denotes the Prokhorov metric and that the collection $\mathcal{D}_P$ of measures with finitely many point masses on elements $\mathcal{D}_T$, each assigned rational mass, form a dense set.

PROPOSITION 4.4 (GÁCS (2005, PROP. B.17)). *Let $v, \mu \in \mathcal{M}_1(T)$, and assume that v is supported on a finite set S . Then the condition $\delta_p(v, \mu) < \epsilon$ is equivalent to the finite set of conditions*

$$\mu(A^\epsilon) > v(A) - \epsilon \quad (29)$$

for all $A \subseteq S$.

The next corollary states that we can compute a representation for a Prokhorov ball in terms of the subbasis elements. The sets are easily defined from those in Proposition 4.4.

COROLLARY 4.5. *Uniformly in $\nu \in \mathcal{D}_P$ and $\epsilon \in \mathbb{Q}$, we can compute a finite collection of pairs $(A_i, q_i)_{i \leq n}$, each A_i a finite union of open balls of radius ϵ around elements of $\mathcal{D}_T$ and each q_i a rational, such that*

$$\{\mu \in \mathcal{M}_1(T) : \delta_p(\mu, \nu) < \epsilon\} = \bigcap_{i \leq n} P(A_i, q_i). \quad (30)$$

Finally, as a direct consequence of Hoyrup and Rojas (2009b, Prop. 4.2.1), these subbasis elements are c.e. open.

PROPOSITION 4.6. *Let A be a c.e. open subset of T and q be a rational. Then the set $P(A, q)$ is c.e. open in the Prokhorov metric, uniformly in A and q .*

Recall that a lower semicomputable function from a computable Polish 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 Polish 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

$$f^{-1}[(q, 1]] \cap D = U_q \cap D. \quad (31)$$

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 4.7. *Let S and T be computable Polish spaces, let κ be a probability kernel from S to T , and let $D \subseteq S$. If $\bar{\kappa}$ is computable on D , then $\kappa(\cdot, A)$ is lower semicomputable on D uniformly in the c.e. open set A , and conversely.*

PROOF. Let $q \in (0, 1)$ be rational, let $A \subseteq T$ be c.e. open, and define $I := (q, 1]$. Then

$$\kappa^{-1}(\cdot, A)[I] = \{x : \bar{\kappa}(x)(A) \in I\} = \bar{\kappa}^{-1}[P(A, q)], \quad (32)$$

where $P(A, q)$ is as in Equation (28). By Proposition 4.6, $P(A, q)$ is even c.e. open.

Suppose $\bar{\kappa}$ is computable on D . Then there is a c.e. open set $V_{A,q}$, uniformly computable in q and A , such that

$$V_{A,q} \cap D = \bar{\kappa}^{-1}[P(A, q)] \cap D = \kappa(\cdot, A)^{-1}[I] \cap D, \quad (33)$$

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

Conversely, suppose $\kappa(\cdot, A)$ is lower semicomputable on D , uniformly in A . Then by Equation (32), uniformly in A and q , we can find a c.e. open $V_{A,q}$ such that Equation (33) holds.

By Corollary 4.5, every basic open ball in the Prokhorov metric is the finite intersection of sets of the form $P(A, q)$, which are c.e. open themselves because a finite intersection of c.e. open sets is c.e. open. Therefore, uniformly in a c.e. open set U in the Prokhorov metric, we can find a c.e. open set V in S such that

$$V \cap D = \bar{\kappa}^{-1}[U] \cap D. \quad (34)$$

Hence, $\bar{\kappa}$ is computable on D . □

Let X and Y be random variables in computable Polish spaces S and T , respectively, and let $\bar{\kappa}(X)$ be a regular version of the conditional distribution $\mathbb{P}[Y|X]$. The above notions of computability are suitable for talking about the computability of κ or any other version of it, and are appropriate notions of computability for statistical applications.

Intuitively, a probability kernel κ is computable when, for some (and hence for any) version of κ , there is a program that, given as input a representation of a point $s \in S$, outputs a representation of the measure $\bar{\kappa}(s)$ for $\mathbb{P}_X$ -almost every input s .

5 DISCONTINUOUS CONDITIONAL DISTRIBUTIONS

Our study of the computability of conditional distributions begins at 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 will rule out almost computability (though not L^1 -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 Polish 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 when $\mathbb{P}\{C = 1\} = 1 - \mathbb{P}\{C = 0\} = p$. A random variable N is a **geometric**(p) random variable when it takes values in $\mathbb{N} = \{0, 1, 2, \dots\}$ and satisfies

$$\mathbb{P}\{N = n\} = p^n (1 - p) \quad (35)$$

for all $n \in \mathbb{N}$. A random variable that takes values in a finite set is **uniformly distributed** when it assigns equal probability to each element. A continuous random variable U on the unit interval is **uniformly distributed** 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, provided p , ℓ , and r are computable reals and $p \in [0, 1]$.

Let N , C , and U be independent $\mathbb{P}$ -almost computable random variables such that N is a geometric($\frac{1}{2}$) random variable, C is a Bernoulli($\frac{1}{2}$) random variable, and U is a uniformly distributed random variable in $[0, 1]$. Fix a computable enumeration $\{r_i\}_{i \in \mathbb{N}}$ of the rational numbers (without repetition) in $(0, 1)$, and consider the random variable

$$X := \begin{cases} r_N, & \text{if } C = 1; \\ U, & \text{otherwise,} \end{cases} \quad (36)$$

which is also $\mathbb{P}$ -almost computable because it is a computable function of C , U , and N .

PROPOSITION 5.1. *Every version of $\mathbb{P}[C = 1|X = \cdot]$ is discontinuous everywhere on every $\mathbb{P}_X$ -measure one set. In particular, no version is $\mathbb{P}_X$ -almost computable.*

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

$$\mathbb{P}\{U \in A\} > 0 \Rightarrow \mathbb{P}\{X \in A\} > 0 \quad (37)$$

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

$$\mathbb{P}\{C = 1 \mid X \text{ rational}\} = 1, \quad (38)$$

while

$$\mathbb{P}\{C = 1 \mid X \text{ irrational}\} = 0. \quad (39)$$

It is also straightforward to verify that C and X are conditionally independent, given an indicator for the event $\{X \text{ rational}\}$. Therefore, all versions f of $\mathbb{P}[C = 1|X = \cdot]$ satisfy, for $\mathbb{P}_X$ -almost

every x ,

$$f(x) = \begin{cases} 1, & x \text{ rational;} \\ 0, & x \text{ irrational.} \end{cases} \quad (40)$$

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

Suppose some version of f were continuous at a point y on a $\mathbb{P}_X$ -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 Lebesgue-almost every irrational in I . Furthermore, the image of every rational in $I \cap R$ is 1, and the image of Lebesgue-almost every irrational in $I \cap R$ is 0, a contradiction. $\square$

Although we cannot hope to compute $\mathbb{P}[C = 1|X = \cdot]$ on a $\mathbb{P}_X$ -measure one set, we can compute it in a weaker sense.

PROPOSITION 5.2. $\mathbb{P}[C = 1|X = \cdot]$ is $L^1(\mathbb{P}_X)$ -computable.

PROOF. By Corollary 2.24, it suffices to construct a sequence of uniformly $\mathbb{P}_X$ -almost computable functions that converge effectively in $L^1(\mathbb{P}_X)$ to $\mathbb{P}[C = 1|X = \cdot]$. Let $\phi_k = \frac{1}{2} \min_{m < n \leq k} |r_m - r_n|$ be half the minimum distance between any pair among $r_0, \dots, r_k$, and define, for every $k \in \mathbb{N}$,

$$f_k(x) := \begin{cases} 1, & \text{if } |x - r_n| < \frac{1}{(k+1)\sqrt{2}} \min(2^{-k-2}, \phi_k) \text{ holds for some } n \leq k; \\ 0, & \text{otherwise.} \end{cases} \quad (41)$$

Note that the set on which f_k takes the value 1 is uniformly $\mathbb{P}_X$ -almost decidable in part because its boundary points are irrationals, a null set. It is then clear that the functions f_k , for $k \in \mathbb{N}$, are uniformly $\mathbb{P}_X$ -almost computable. For every $k \in \mathbb{N}$, we have that $\mathbb{P}\{X = r_n \text{ for some } n > k\} = 2^{-k-2}$. Therefore,

$$\int \left| \mathbb{P}[C = 1|X = x] - f_k(x) \right| \mathbb{P}_X(dx) \quad (42)$$

$$\leq \mathbb{P}\{X = r_n \text{ for some } n > k\} + \int_{[0,1] \setminus \mathbb{Q}} \left| \mathbb{P}[C = 1|X = x] - f_k(x) \right| \mathbb{P}_X(dx) \quad (43)$$

$$\leq 2^{-k-2} + \frac{1}{(k+1)\sqrt{2}} \sum_{n=0}^k 2^{-k-2} < 2^{-k-1}, \quad (44)$$

completing the proof. $\square$

6 CONDITIONING IS DISCONTINUOUS

Conditioning in general can produce discontinuous conditional distributions, which is an obstruction to a conditioning operator being computable. But even if we restrict our attention to distributions that admit conditional distributions that are continuous on their support, the operation of conditioning cannot be computable because, as we will show, it is discontinuous. Indeed, conditioning is discontinuous in a rather strong way. We use the recursion theorem to explain the computational consequences. Namely, for any potential program analysis that aims to perform conditioning on an arbitrary given distribution, there is a representation of that distribution such that the program analysis cannot identify a single nontrivial fact about its conditional distribution.

To begin, we formalize the notion of a conditioning operator.

Definition 6.1. Let $\mathcal{F} \subseteq \mathcal{M}_1([0, 1]^2)$ be a set of probability measures. A map $\Phi : \mathcal{M}_1([0, 1]^2) \times [0, 1] \rightarrow \mathcal{M}_1([0, 1])$ is a **conditioning operator** (for $\mathcal{F}$) if, for all distributions $\mu \in \mathcal{F}$ and random variables X and Y with joint distribution μ , we have $\Phi(\mu, x) = \mathbb{P}[Y|X = x]$ for $\mathbb{P}_X$ -almost all x .

Observe, by Proposition 3.7, that there is a conditioning operator for all $\mathcal{F} \subseteq \mathcal{M}_1([0, 1]^2)$.

Definition 6.2. Let $\mathcal{F} \subseteq \mathcal{M}_1([0, 1]^2)$. A conditioning operator for $\mathcal{F}$ is **computable** if it is computable on $\mathcal{F} \times [0, 1]$, considered as a function $\mathcal{M}_1([0, 1]^2) \times [0, 1] \rightarrow \mathcal{M}_1([0, 1])$, where both $\mathcal{M}_1([0, 1]^2) \times [0, 1]$ and $\mathcal{M}_1([0, 1])$ are taken to be the canonical computable Polish spaces.

The previous section motivates restricting one's attention to conditioning operators for the set $\mathcal{F}_0 \subseteq \mathcal{M}_1([0, 1]^2)$ of probability distributions on pairs (X, Y) of random variables in $[0, 1]$ such that there exists a $\mathbb{P}_X$ -almost continuous version of the conditional distribution map $\mathbb{P}[Y|X = \cdot]$. We will show that conditioning operators for $\mathcal{F}_0$ are not computable, simply on grounds of continuity.

Recall that the name of a probability measure $\mu \in \mathcal{M}_1(T)$ on a computable Polish space T is given by a Cauchy sequence in the dense elements $\mathcal{D}_{P,T}$ of the associated Prokhorov metric. Note that $\mathcal{F}_0$ contains $\mathcal{D}_{P,[0,1]^2}$. Further recall that $P_T(A, q)$ is defined to be the set $\{\eta \in \mathcal{M}_1(T) : \eta(A) > q\}$, for any open set $A \subseteq T$ and rational $q \in \mathbb{Q}$. Let $\mathcal{A}(T) := \{(A, q) : A \text{ is a finite union of open balls in } T, \text{ and } q \in \mathbb{Q}\}$.

Given a computable Cauchy sequence in the Prokhorov metric that converges to a measure $\mu \in \mathcal{M}_1(T)$, by Lemma 4.5, we can compute a sequence $\langle A_i, q_i \rangle_{i \in \mathbb{N}}$ in $\mathcal{A}(T)$ such that $\bigcap_{i \in \mathbb{N}} P_T(A_i, q_i) = \{\mu\}$. Further, by Proposition 4.6, given a finite sequence $\langle A_i, q_i \rangle_{i \leq n}$ in $\mathcal{A}(T)$, we can compute, uniformly in $\langle A_i, q_i \rangle_{i \leq n}$, a ball in the Prokhorov metric contained in $\bigcap_{i \leq n} P_T(A_i, q_i)$. Therefore, uniformly in a probability measure μ and a collection $\langle A_i, q_i \rangle_{i \in \mathbb{N}}$ with $\bigcap_{i \in \mathbb{N}} P_T(A_i, q_i) = \{\mu\}$, we can computably recover a name for μ . Conversely, from a name for μ , we can uniformly compute such a collection.

LEMMA 6.3. *Let $\mathcal{F} \subseteq \mathcal{M}_1([0, 1]^2)$ contain $\mathcal{D}_{P,[0,1]^2}$. For every $\alpha \in \mathcal{D}_{P,[0,1]}$, computable representation $\{v_i\}_{i \in \mathbb{N}}$ of $v \in \mathcal{M}_1([0, 1]^2)$, computable representation $\{x_i\}_{i \in \mathbb{N}}$ of $x \in [0, 1]$, and rational $\epsilon \in (0, 1)$, we can uniformly find a measure $\mu \in \mathcal{D}_{P,[0,1]^2}$ such that $\delta_P(\mu, v) < \epsilon$ and $\Phi(\mu, x) = \alpha$ for every conditioning operator Φ for $\mathcal{F}$.*

PROOF. Relative to $\{v_i\}_{i \in \mathbb{N}}$ and $\{x_i\}_{i \in \mathbb{N}}$, we can computably find an element $p^* \in \mathcal{D}_{P,[0,1]^2}$ such that $p^*({x} \times [0, 1]) = 0$ and $\delta_P(p^*, v) < \frac{\epsilon}{2}$. Let δ_x denote the Dirac measure on $[0, 1]$ concentrating on x and let $\tau \otimes \tau'$ denote the product measure on $[0, 1]^2$ with respective marginal distributions $\tau, \tau' \in \mathcal{M}_1([0, 1])$. Defining $p := \frac{\epsilon}{2}(\delta_x \otimes \alpha) + (1 - \frac{\epsilon}{2})p^*$, it is easy to check that $\delta_P(p, p^*) \leq \frac{\epsilon}{2}$, hence $\delta_P(v, p) < \epsilon$. Because $p^*({x} \times [0, 1]) = 0$ and p^* is a finite mixture of point masses, every conditioning operator Φ for $\mathcal{F}$ must satisfy $\Phi(\mu, x) = \alpha$. $\square$

PROPOSITION 6.4. *Let $\mathcal{F} \subseteq \mathcal{M}_1([0, 1]^2)$ contain $\mathcal{D}_{P,[0,1]^2}$. Every conditioning operator on $\mathcal{F}$ is discontinuous everywhere, hence noncomputable.*

PROOF. On $\mathcal{M}_1([0, 1])$, adopt the weak topology (induced by the standard topology on $[0, 1]$). On $\mathcal{M}_1([0, 1]^2) \times [0, 1]$, adopt the product topology induced by the weak and standard topologies, respectively. Then Lemma 6.3 implies that every conditioning operator Φ for $\mathcal{F}$ is discontinuous everywhere. Hence, every conditioning operator is noncomputable. $\square$

The above definitions and proposition capture the essential difficulty of conditioning: a finite approximation to the joint distribution determines nothing about the result of conditioning on a particular point.

We now establish a stronger notion of noncomputability, namely that it is not even possible to always produce *some* nontrivial fact about a conditional distribution. For $e \in \mathbb{N}$, let φ_e denote the partial computable function defined by code e . The recursion theorem, due to Kleene (1938), states that when F is a total computable function, there is some integer i for which the partial computable functions φ_i and $\varphi_{F(i)}$ are equal partial functions (i.e., they are defined on the same inputs and are equal where they are defined). For more details, see Rogers (1987, Ch. 11).

Definition 6.5. A program $\varphi_e : \mathbb{N} \rightarrow \mathbb{N}$ **represents a distribution** μ_e on a computable Polish space T if it is total and on input k , the output value $\varphi_e(k)$ is a code for a pair $(A_k, q_k) \in \mathcal{A}(T)$ such that $\{\mu\} = \bigcap_{k \in \mathbb{N}} P_T(A_k, q_k)$.

Definition 6.6. A program $\varphi_a : \mathbb{N}^3 \rightarrow \mathbb{N}$ is a **conditioning program** for $\mathcal{F} \subseteq \mathcal{M}_1([0, 1]^2)$ if it is total and whenever e represents a computable distribution $\mu_e \in \mathcal{F}$, there exists a conditioning operator Φ for $\mathcal{F}$ such that, for every code $j \in \mathbb{N}$ for a computable real $x \in [0, 1]$, and for every $k \in \mathbb{N}$, the return value $\varphi_a(e, j, k)$ is a code for either the empty string or an element $(A, q) \in \mathcal{A}([0, 1])$ such that $\Phi(\mu_e, x) \in P_{[0, 1]}(A, q)$ and $P_{[0, 1]}(A, q) \neq \mathcal{M}_1([0, 1])$.

THEOREM 6.7 (NONAPPROXIMABLE CONDITIONAL DISTRIBUTIONS). *Suppose that φ_a is a conditioning program for a set $\mathcal{F} \subseteq \mathcal{M}_1([0, 1]^2)$ containing $\mathcal{D}_{P, [0, 1]^2}$. Let $e \in \mathbb{N}$ be a code for a computable distribution μ_e on $[0, 1]^2$, and let $j \in \mathbb{N}$ be a code for a computable real $x \in [0, 1]$. Then uniformly in a, e , and j , we can compute an $i \in \mathbb{N}$ such that $\mu_e = \mu_i$ and $\varphi_a(i, j, k)$ is a code for the empty string for every $k \in \mathbb{N}$.*

PROOF. Uniformly in a , we can compute some $b \in \mathbb{N}$ such that for all $n, m, r \in \mathbb{N}$: if the value $\varphi_a(n, m, r')$ is a code for the empty string for all $r' \leq r$, then $\varphi_b(n, m, r)$ is also a code for the empty string; and otherwise $\varphi_b(n, m, r) = \varphi_a(n, m, r')$, where r' is the least index such that $\varphi_a(n, m, r')$ is not a code for the empty string. Note that for each $n, m \in \mathbb{N}$, the value $\varphi_b(n, m, k)$ is a code for the empty string for all $k \in \mathbb{N}$ if and only if $\varphi_a(n, m, k)$ is a code for the empty string for all k . Let $\eta_{n,j}$ denote the least index $k \in \mathbb{N}$ such that $\varphi_b(n, j, k)$ is not a code for the empty string, if such k exists, and ∞ otherwise. Note that for each $k \in \mathbb{N}$, we can compute (uniformly in n, e, a , and j) whether or not $k < \eta_{n,j}$ (even though the finiteness of $\eta_{n,j}$ may not be computable). For $k \in \mathbb{N}$, let (A_k, q_k) be the pair coded by $\varphi_e(k)$.

Define the total computable function $F : \mathbb{N} \rightarrow \mathbb{N}$ such that for $n, k \in \mathbb{N}$,

$$\varphi_{F(n)}(k) = \begin{cases} \varphi_e(k) & \text{if } k < \eta_{n,j}; \text{ and} \\ \varphi_{e'}(k - \eta_{n,j}) & \text{if } k \geq \eta_{n,j}, \end{cases} \quad (45)$$

where $e' \in \mathbb{N}$ is defined as follows: Let (A', q') be the pair coded by $\varphi_b(n, j, \eta_{n,j})$. First, compute a Prokhorov ball $B \subseteq \mathcal{M}_1([0, 1]^2)$ contained within $\bigcap_{\ell \leq \eta_{n,j}} P_{[0, 1]^2}(A_\ell, q_\ell)$. Next, compute some $\alpha \in \mathcal{M}_1([0, 1])$ such that $\alpha(A') = 0$, and hence $\alpha \notin P_{[0, 1]}(A', q')$. Then, by Lemma 6.3, compute a code e' for a distribution $\nu \in B$ such that $\Phi(\nu, x) = \alpha$ for every conditioning operator Φ for $\mathcal{F}$.

By the recursion theorem, we can compute an index i , uniformly in a, e , and j , such that $\varphi_{F(i)} = \varphi_i$. We now argue that $\eta_{i,j} = \infty$, which implies that $\varphi_i = \varphi_e$ by Equation (45).

Suppose, for a contradiction, that $\eta_{i,j} \in \mathbb{N}$. Then $\varphi_b(i, j, \eta_{i,j}) = (A', q')$ for some $(A', q') \in \mathcal{A}([0, 1])$. Hence, as φ_a is a conditioning program, there is some conditioning operator Φ for $\mathcal{F}$, such that $\Phi(\mu_i, x) \in P_{[0, 1]}(A', q')$, where μ_i is the measure represented by φ_i . By construction, for every conditioning operator Φ for $\mathcal{F}$, we have $\Phi(\mu_i, x) \notin P_{[0, 1]}(A', q')$, a contradiction. $\square$

These results rely on the density of the finitely supported discrete probability distributions $\mathcal{D}_{P, [0, 1]^2}$. However, analogous results can be established if we restrict ourselves to absolutely continuous distributions admitting continuous joint density functions. In this case, the role of the finitely supported continuous distributions would be played by absolutely continuous distributions with sharp but continuous bump functions concentrating on small sets. The fundamental obstruction is the same: partial information in the weak topology does not suffice to condition continuously.

7 NONCOMPUTABLE ALMOST-CONTINUOUS CONDITIONAL DISTRIBUTIONS

In this section, we construct a pair of $\mathbb{P}$ -almost computable random variables X in $[0, 1]$ and N in $\mathbb{N}$ such that the conditional probability map $\mathbb{P}[N = k | X = \cdot]$ is not even $L^1(\mathbb{P}_X)$ -computable,

despite the existence of a $\mathbb{P}_X$ -almost continuous version. Our construction in this section can be thought of as providing a single witness to the noncomputability of the conditioning operator.

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. We may then take $\emptyset' : \mathbb{N} \rightarrow \{0, 1\}$ to denote the halting set

$$\{\ell : M_\ell \text{ halts on input } 0\}, \quad (46)$$

which is computably enumerable but not computable. The set $\emptyset'$ and the function h are computable from each other because

$$\emptyset' = \{n \in \mathbb{N} : h(n) < \infty\}. \quad (47)$$

We now use h to define a pair of $\mathbb{P}$ -almost computable random variables (N, X) such that $\emptyset'$ is computable from $\mathbb{P}[N|X = \cdot]$.

Let N , C , U , and V be independent $\mathbb{P}$ -almost computable random variables such that N is a geometric($\frac{1}{3}$) random variable, C is a Bernoulli($\frac{1}{3}$) random variable, and U and V are uniformly distributed random variables in $[0, 1]$.

Let $\lfloor x \rfloor$ denote the greatest integer $y \leq x$, and note that $\lfloor 2^k V \rfloor$ is uniformly distributed in $\{0, 1, 2, \dots, 2^k - 1\}$ and is $\mathbb{P}$ -almost computable. For each $k \in \mathbb{N}$, consider the derived random variable

$$X_k := \frac{2\lfloor 2^k V \rfloor + C + U}{2^{k+1}}. \quad (48)$$

Note that $\lim_{k \rightarrow \infty} X_k$ almost surely exists. Define $X_\infty := \lim_{k \rightarrow \infty} X_k$, and observe that $X_\infty = V$ a.s. Finally, define $X := X_{h(N)}$.

PROPOSITION 7.1. *The random variable X is $\mathbb{P}$ -almost computable.*

PROOF. Because U and V are computable on a $\mathbb{P}$ -measure one set and a.s. nondyadic, their binary expansions $\{U_n : n \in \mathbb{N}\}$ and $\{V_n : n \in \mathbb{N}\}$ (which are uniquely determined with probability 1) are themselves $\mathbb{P}$ -almost computable random variables in $\{0, 1\}$, uniformly in n .

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

$$D_k = \begin{cases} V_k, & h(N) > k; \\ C, & h(N) = k; \\ U_{k-h(N)-1}, & h(N) < k. \end{cases} \quad (49)$$

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

Let D denote the $\mathbb{P}$ -almost computable random real whose binary expansion is $\{D_k\}_{k \geq 0}$. There are two cases to consider.

First, conditioned on the event $\{h(N) = \infty\}$, we have that $D_k = V_k$ for all $k \geq 0$, and so $D = V = X_\infty = X$ almost surely.

In the second case, let $m \in \mathbb{N}$, and condition on the event $\{h(N) = m\}$. We must then show that $D = X_m$ a.s. Note that

$$\lfloor 2^m X_m \rfloor = \lfloor 2^m V \rfloor = \sum_{k=0}^{m-1} 2^{m-1-k} V_k = \lfloor 2^m D \rfloor, \quad (50)$$

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 $\mathbb{P}[N|X = \cdot]$ is $\mathbb{P}_X$ -almost continuous. We begin by characterizing the conditional density of X given N .

LEMMA 7.2. *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 $[0, 1]$ given by $p_{X_\infty}(x) = 1$ and*

$$p_{X_k}(x) = \begin{cases} \frac{4}{3}, & \lfloor 2^{k+1}x \rfloor \text{ even}; \\ \frac{2}{3}, & \lfloor 2^{k+1}x \rfloor \text{ odd}, \end{cases} \quad (51)$$

for $k < \infty$.

PROOF. We have $X_\infty = V$ a.s., and so the constant function taking the value 1 is a 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,

$$\mathbb{P}\{\lfloor 2^{k+1}X_k \rfloor = \ell\} = 2^{-k} \cdot \begin{cases} \frac{2}{3}, & \ell \text{ even}; \\ \frac{1}{3}, & \ell \text{ odd}, \end{cases} \quad (52)$$

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

$$p(x) = 2^{-k} \cdot \begin{cases} \frac{2}{3}, & \lfloor x \rfloor \text{ even}; \\ \frac{1}{3}, & \lfloor x \rfloor \text{ odd}. \end{cases} \quad (53)$$

A density of X_k is then obtained by the 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

$$p_{X|N}(x|n) = p_{X_{h(n)}}(x) \quad (54)$$

for $x \in [0, 1]$ and $n \in \mathbb{N}$. By Bayes' rule (Lemma 3.11), for every $B \subseteq \mathbb{N}$ and $\mathbb{P}_X$ -almost every $x \in [0, 1]$,

$$\mathbb{P}[N \in B|X = x] = \frac{\phi(x, B)}{\phi(x, \mathbb{N})}, \quad (55)$$

where

$$\phi(x, B) := \sum_{n \in B} p_{X|N}(x|n) \cdot \mathbb{P}\{N = n\}. \quad (56)$$

LEMMA 7.3. $\mathbb{P}[N|X = \cdot]$ is $\mathbb{P}_X$ -almost continuous.

PROOF. For every $k \in \mathbb{N} \cup \{\infty\}$, it holds that p_{X_k} is bounded by $4/3$, positive, and continuous on the $\mathbb{P}_X$ -measure one set R of nondyadic reals in the unit interval. Therefore, $p_{X|N}$ is positive, bounded, and continuous on $R \times \mathbb{N}$, and so, by Proposition 3.12, $\mathbb{P}[N|X = \cdot]$ is $\mathbb{P}_X$ -almost continuous. $\square$

LEMMA 7.4. $\mathbb{P}[N|X = \cdot]$ is $\emptyset'$ -computable on a $\mathbb{P}_X$ -measure one set.

PROOF. Let $B \subseteq \mathbb{N}$ be a computable set. By Lemma 4.7, Equation (55), and the computability of division, it suffices to show that $\phi_B := \phi(\cdot, B)$ is $\mathbb{P}_X$ -almost computable from $\emptyset'$. Note that $\mathbb{P}_X$ is absolutely continuous with respect to Lebesgue measure (on $[0, 1]$) and vice versa, and so $\mathbb{P}_X$ -almost and Lebesgue-almost computability coincide. We will therefore drop the measure when referring to almost computability for the remainder of the proof. Define, for each $n \in \mathbb{N}$, the function a_n on $[0, 1]$ given by

$$a_n(x) := \begin{cases} 3, & h(n) = \infty; \\ 2, & h(n) < \infty \text{ and } \lfloor 2^{h(n)}x \rfloor \text{ even; and} \\ 4, & h(n) < \infty \text{ and } \lfloor 2^{h(n)}x \rfloor \text{ odd.} \end{cases} \quad (57)$$

Clearly the functions $\{a_n(x)\}_{n \in \mathbb{N}}$ are almost computable from $\emptyset'$, uniformly in n . Observe that, for all $n \in \mathbb{N}$ and $x \in [0, 1]$,

$$a_n(x) = 3 p_{X|N}(x|n) \quad (58)$$

by Lemma 7.2. Also recall that $\mathbb{P}\{N = n\} = \frac{4}{5} \cdot 5^{-n}$ for all $n \in \mathbb{N}$. Hence, for any finite set $F \subseteq \mathbb{N}$, the function

$$\phi_F(x) := \phi(x, F) = \sum_{n \in F} p_{X|N}(x|n) \cdot \mathbb{P}\{N = n\} = \frac{4}{15} \sum_{n \in F} a_n(x) \cdot 5^{-n} \quad (59)$$

is almost computable from $\emptyset'$, uniformly in F . However, for every $k \in \mathbb{N}$ and $x \in [0, 1]$,

$$|\phi_B(x) - \phi_{F_k}(x)| \leq \sum_{n > k} p_{X|N}(x|n) \cdot \mathbb{P}\{N = n\} \leq \frac{4}{15} \cdot 5^{-k}, \quad (60)$$

where $F_k = B \cap \{1, \dots, k\}$. It follows that ϕ_B is almost computable from $\emptyset'$. $\square$

PROPOSITION 7.5. Let $R \subseteq [0, 1]$ be a measurable subset of $\mathbb{P}_X$ -measure greater than $\frac{5}{6}$. For each $k \in \mathbb{N}$, the conditional probability map $\mathbb{P}[N = k|X = \cdot]$ is neither lower nor upper semicomputable on R .

PROOF. First note that if R had Lebesgue measure no greater than $\frac{3}{4}$, then, for $k \in \mathbb{N}$,

$$\mathbb{P}_{X_k}(R) \leq \frac{1}{2} \cdot \frac{4}{3} + \frac{1}{4} \cdot \frac{2}{3} = \frac{5}{6}, \quad (61)$$

and so $\mathbb{P}_X(R) \leq \frac{5}{6}$, a contradiction. Hence, R has Lebesgue measure greater than $\frac{3}{4}$.

Fix $k \in \mathbb{N}$. By (55) and the definition of ϕ ,

$$\mathbb{P}[N = k|X = x] = \frac{p_{X|N}(x|k) \cdot \mathbb{P}\{N = k\}}{\phi(x, \mathbb{N})}. \quad (62)$$

for a.e. x . The density $p_{X|N}(\cdot|k) = p_{X_{h(k)}}$, given by Lemma 7.2, is a piecewise constant function and computable on a measure one set. Furthermore, $\mathbb{P}\{N = k\} = 4 \cdot 5^{-k-1}$ is a computable real. Hence, it remains to show that, as a function of x , the denominator $\phi(x, \mathbb{N}) := \sum_{n \in \mathbb{N}} p_{X|N}(x|n) \cdot \mathbb{P}\{N = n\}$ of the right-hand side is neither lower nor upper semicomputable on R . We will show the former; the latter follows in a similar fashion.

Recall the sequence of functions $\{a_n\}_{n \in \mathbb{N}}$ on $[0, 1]$ from Equation (57), and define the function $\tau : [0, 1] \rightarrow \mathbb{R}$ by

$$\tau(x) := \sum_{n \in \mathbb{N}} a_n(x) \cdot 5^{-n}, \quad (63)$$

which furthermore satisfies

$$\tau(x) = \sum_{n \in \mathbb{N}} 3 p_{X|N}(x|n) \cdot \frac{5}{4} \mathbb{P}\{N = n\} \quad (64)$$

for all $x \in [0, 1]$. Observe that for every $x \in [0, 1]$, the real $\tau(x)$ has a unique base-5 expansion with all the digits contained in $\{2, 3, 4\}$, which must necessarily be given by the sequence $a_0(x), a_1(x), \dots$. For each $\ell \in \mathbb{N}$, define

$$\tau_\ell(x) := 5^\ell \left(\tau(x) - \sum_{n < \ell} a_n(x) \cdot 5^{-n} \right) = \sum_{n \in \mathbb{N}} a_{n+\ell}(x) \cdot 5^{-n}. \quad (65)$$

Note that $\tau_0 (= \tau)$ is lower semicomputable on R precisely when $\phi(x, \mathbb{N})$ is lower semicomputable on R . Further note that if $a_\ell(x) = 2$, then $\tau_\ell(x) < 3$ (where the inequality is strict because for every n , there exists an $m > n$ such that $h(m) = \infty$). However, if $a_\ell(x) \geq 3$, then $\tau_\ell(x) > 3$.

We now prove by induction, uniformly in ℓ , that, if τ_ℓ is lower semicomputable on R , then $\tau_{\ell+1}$ is lower semicomputable on R , and we can compute whether or not $h(\ell)$ is finite. This then implies that if τ were lower semicomputable on R , then $\emptyset'$ would be computable, a contradiction.

Suppose τ_ℓ is lower semicomputable on R . Then, we can compute (as a c.e. open subset of $[0, 1]$) a set S such that $S \cap R = \tau_\ell^{-1}([3, \infty)) \cap R$. Simultaneously run M_ℓ (on input 0). If $h(\ell) < \infty$, then we will eventually notice this fact by observing that M_ℓ halts. It is also the case that $S \cap R \subseteq \{x \in [0, 1] : a_\ell(x) = 4\}$, and so, by the definition of the functions $\{a_n\}$, the set $S \cap R$ has Lebesgue measure at most $\frac{1}{2}$. Therefore, S has measure at most $\frac{1}{2}$ plus the measure of $[0, 1] - R$, i.e., S has measure less than $\frac{1}{2} + (1 - \frac{3}{4}) = \frac{3}{4}$. However, if $h(\ell) = \infty$, then $S \supseteq R$, and hence S has Lebesgue measure greater than $\frac{3}{4}$, which we will eventually notice (and which rules out the first case). Hence, we can compute whether or not $h(\ell)$ is finite, which in turn implies that a_ℓ is computable on a measure one set. Therefore, as $\tau_{\ell+1}(x) = 5 \cdot (\tau_\ell(x) - a_\ell(x))$, the function $\tau_{\ell+1}$ is lower semicomputable on R . $\square$

We may summarize our central result as follows:

THEOREM 7.6. *There are $\mathbb{P}$ -almost computable random variables W and Y on $[0, 1]$ such that the conditional distribution map $\mathbb{P}[Y|W = \cdot]$ is $\mathbb{P}_W$ -almost continuous but not $\mathbb{P}_W$ -almost computable.*

PROOF. Let X and N be as above, let Y be uniformly distributed on $[0, 1]$, and let M be the geometric($\frac{1}{2}$) random variable given by $M := \lfloor -\log_2 Y \rfloor$. Finally, let W be such that $\mathbb{P}[W|Y](\omega) = \mathbb{P}[X|N = M(\omega)]$. Note that we have $\mathbb{P}[Y \in (2^{-n-1}, 2^{-n})|W = x] = \mathbb{P}[Y \in [2^{-n-1}, 2^{-n})|W = x] = \mathbb{P}[N = n|X = x]$ for $n \in \mathbb{N}$ and $x \in [0, 1]$. The result then follows from Lemma 7.3 and Proposition 7.5. $\square$

For a visualization of the distribution of (W, Y) , see Figure 1.

The proof of Proposition 7.5 shows that not only is the conditional distribution map $\mathbb{P}[N|X = \cdot]$ not $\mathbb{P}_X$ -almost computable, but that, in fact, it computes the halting set $\emptyset'$; to make this precise, we would define the notion of an oracle that encodes $\mathbb{P}[N|X = \cdot]$ using, e.g., infinite strings as in the Type-2 Theory of Effectivity. Despite not having a definition of computability from the conditional distribution map, we can easily relativize the notion of computability for the conditional distribution map to obtain the following:

COROLLARY 7.7. *If $\mathbb{P}[N|X = \cdot]$ is A -computable on a set of $\mathbb{P}_X$ -measure greater than $\frac{5}{6}$ for an oracle $A \subseteq \mathbb{N}$, then A computes the halting set, i.e., $A \geq_T \emptyset'$. In particular, $\mathbb{P}[N|X = \cdot]$ is not computable on any set of $\mathbb{P}_X$ -measure greater than $\frac{5}{6}$, and hence no version of $\mathbb{P}[N|X = \cdot]$ is $\mathbb{P}_X$ -almost computable.*

Fig. 1. A visualization of the distribution of (W, Y) , as defined in Theorem 7.6. In this plot, the darkness of each pixel is proportional to the mass contained in that region. Note that this is not a plot of the (noncomputable) density, but rather of a discretization to a given pixel size. 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 (as opposed to nonuniform but at a finer granularity than the resolution of this printer/display) is tantamount to solving the halting problem, but it is possible to sample from this distribution nonetheless.

PROOF. Suppose A is such that $\mathbb{P}[N|X = \cdot]$ is A -computable on a set of $\mathbb{P}_X$ -measure greater than $5/6$. Then $\mathbb{P}[N = 1|X = \cdot]$ is A -computable on the same set. But then, by the argument in the proof of Proposition 7.5, the function h , and hence the halting set $\emptyset'$, is computable from A . $\square$

However, by Lemma 7.4, the conditional distribution map $\mathbb{P}[N|X = \cdot]$ is in fact $\emptyset'$ -computable on a $\mathbb{P}_X$ -measure one set, and so the bound in Corollary 7.7 is the best possible.

Computable operations map computable points to computable points, and so Corollary 7.7 provides another context in which conditioning operators are noncomputable (cf. Proposition 6.4).

The next result shows that Proposition 7.5 also rules out the computability of conditional probability maps in the weaker sense of $L^1(\mathbb{P}_X)$ -computability. This extends Hoyrup and Rojas (2011, Prop. 3), which states that there is a pair of computable measures $\mu \ll \nu$ on a computable Polish space such that the Radon–Nikodym derivative $d\mu/d\nu$ is not L^1 -computable. Namely, we show that in the case of measures on $[0, 1]$, the measures can be taken to be of the form $\mu = \mathbb{P}\{X \in \cdot, N = k\}$ and $\nu = \mathbb{P}\{X \in \cdot\} = \mathbb{P}_X$.

PROPOSITION 7.8. *For every $k \in \mathbb{N}$, the map $\mathbb{P}[N = k|X = \cdot]$ is not $L^1(\mathbb{P}_X)$ -computable.*

PROOF. Let $k \in \mathbb{N}$. By Proposition 7.5, the conditional probability map $\mathbb{P}[N = k|X = \cdot]$ is not computable on any measurable set R of $\mathbb{P}_X$ -measure greater than $5/6$. However, by Lemma 2.23, the conditional probability map is $L^1(\mathbb{P}_X)$ -computable only if, for each $r \in \mathbb{N}$, the map is computable on some set of $\mathbb{P}_X$ -measure at least $1 - 2^{-r}$, uniformly in r . This does not hold, and so every conditional probability map $\mathbb{P}[N = k|X = \cdot]$ is not $L^1(\mathbb{P}_X)$ -computable. $\square$

It is natural to ask whether this construction can be modified to produce a pair of $\mathbb{P}$ -almost computable random variables, like N and X , such that the corresponding conditional distribution map is not $\mathbb{P}_X$ -almost computable even though it has an *everywhere continuous* version. We provide such a strengthening in the next section.

8 NONCOMPUTABLE EVERYWHERE CONTINUOUS CONDITIONAL DISTRIBUTIONS

In Section 7, we demonstrated a pair of computable random variables that admit a conditional distribution that is continuous on a measure one set but still noncomputable on every measure one set. It is therefore 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 map $\mathbb{P}[N|Z = \cdot]$, which is itself nonetheless not computable. In fact, we do so now, using a construction similar to that of (X, N) in Section 7.

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 Bernoulli random variable 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 begin the construction. Let N, U, V , and C be as in the first construction. We next define several random variables from which we will construct S and then Z .

LEMMA 8.1. *There is a random variable F in $[0, 1]$ with the following properties:*

- (1) F is $\mathbb{P}$ -almost computable.
- (2) $\mathbb{P}_F$ admits a computable 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 8.1, and independent of all earlier random variables mentioned. Note that F is almost surely nondyadic and so the r th bit F_r of F is a $\mathbb{P}$ -almost computable random variable, uniformly in r .

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

$$S = \frac{1}{8} \times \begin{cases} F, & \text{if } D = 0; \\ 4 + (1 - F), & \text{if } D = 4; \\ 4C + (D \bmod 4) + U, & \text{otherwise.} \end{cases} \quad (66)$$

It is clear that S is also $\mathbb{P}$ -almost computable, and straightforward to show that

Fig. 2. Left: The graph of the function defined by $f(x) = \exp\{-(1-x^2)^{-1}\}$, for $x \in (-1, 1)$, and 0 otherwise, a C^∞ bump function whose derivatives at ± 1 are all 0. Right: A density $p(y) = 2/3 (1 + \Phi(2y-1)/\Phi(1))$, for $y \in (0, 1)$, of a random variable satisfying Lemma 8.1, where $\Phi(y) = \int_{-1}^y f(x) dx$ is the integral of the bump function.

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

(i) $\mathbb{P}_S$ admits an infinitely differentiable and computable 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.)

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

$$Z_k := \frac{\lfloor 2^k V \rfloor + S}{2^k}, \quad (67)$$

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

$$(Z_k)_n = \begin{cases} V_n, & n < k; \\ S_{n-k}, & n \geq k \end{cases} \quad \text{a.s.}, \quad (68)$$

where S_ℓ denotes the ℓ th bit of S . It is straightforward to show from (i) and (ii) above that $\mathbb{P}_{Z_k}$ admits an infinitely differentiable density p_{Z_k} with respect to Lebesgue measure on $[0, 1]$.

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

LEMMA 8.2. *The random variable Z is $\mathbb{P}$ -almost computable.*

LEMMA 8.3. *There is an everywhere continuous version of $\mathbb{P}[N|Z = \cdot]$.*

PROOF. By construction, the conditional density of Z given N is everywhere continuous, bounded, and positive. The result follows from Proposition 3.12 for $R = [0, 1]$. $\square$

We next show that, for each $k \in \mathbb{N}$, the conditional probability map $\mathbb{P}[N = k|Z = \cdot]$ is not computable. Our proof relies on the fact that, for each $k \in \mathbb{N}$, there is a large set of points $x \in [0, 1]$ for which the density $p_{Z|N}(x|k) := p_{Z_{h(k)}}(x)$ agrees with $p_{X|N}(x|k)$. This will then allow us to use techniques similar to those of Proposition 7.5.

We say a real $x \in [0, 1]$ is **valid for $\mathbb{P}_S$** 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 $\mathbb{P}_S$. The following are then consequences of the construction of S and the definition of valid points:

- (iii) If x is valid for $\mathbb{P}_S$, then $p_S(x) \in \{\frac{2}{3}, \frac{4}{3}\}$. In particular, $p_S(x) = \frac{4}{3}$ for $x \in (\frac{1}{8}, \frac{1}{2})$, and $p_S(x) = \frac{2}{3}$ for $x \in (\frac{5}{8}, 1)$.
- (iv) The Lebesgue measure of points valid for $\mathbb{P}_S$ is $(\frac{1}{2} - \frac{1}{8}) + (1 - \frac{5}{8}) = \frac{3}{4}$.

For $k < \infty$, we say that $x \in [0, 1]$ is **valid for $\mathbb{P}_{Z_k}$** if the fractional part of $2^k x$ is valid for $\mathbb{P}_S$, and we say that x is **valid for $\mathbb{P}_{Z_\infty}$** for all x . Let A_k be the collection of x valid for $\mathbb{P}_{Z_k}$, and let $A_\infty = [0, 1]$. Note that, at all points x that are valid for $\mathbb{P}_{Z_k}$, the density $p_{Z_k}(x)$ agrees with $p_{X_k}(x)$, and, at all points x that are valid for $\mathbb{P}_{Z_\infty}$, the density $p_{Z_\infty}(x)$ agrees with $p_{X_\infty}(x)$.

For each $k < \infty$, the set A_k consists of the set of points that are valid for $\mathbb{P}_S$ first dilated by a factor of 2^{-k} and then tiled 2^k -many times, and so A_k has Lebesgue measure $\frac{3}{4}$; the set $A_\infty = [0, 1]$ has Lebesgue measure one.

PROPOSITION 8.4. *Let $R \subseteq [0, 1]$ be a measurable subset of $\mathbb{P}_Z$ -measure greater than $11/12$. Then for each $k \in \mathbb{N}$, the conditional probability map $\mathbb{P}[N = k|Z = \cdot]$ is neither lower semicomputable nor upper semicomputable on R .*

PROOF. We proceed analogously to the proof of Proposition 7.5. First note that if R had Lebesgue measure no greater than $\frac{7}{8}$, then, for $k \in \mathbb{N}$,

$$\mathbb{P}_{Z_k}(R) \leq \frac{1}{2} \cdot \frac{4}{3} + \frac{3}{8} \cdot \frac{2}{3} = \frac{11}{12} \quad (69)$$

and so $\mathbb{P}_Z(R) \leq \frac{11}{12}$, a contradiction. Hence, R has Lebesgue measure greater than $\frac{7}{8}$.

Fix $k \in \mathbb{N}$. By Bayes' rule (Lemma 3.11),

$$\mathbb{P}[N = k|Z = x] = \frac{p_{Z|N}(x|k) \cdot \mathbb{P}\{N = k\}}{\sum_{n \in \mathbb{N}} p_{Z|N}(x|n) \cdot \mathbb{P}\{N = n\}} \quad (70)$$

for $\mathbb{P}_Z$ -a.e. x , and hence for Lebesgue-a.e. x .

Now, $p_{Z|N}(\cdot|k)$ is piecewise (on a finite number of pieces with rational endpoints) the computable dilation and translation of the computable function p_S . By (ii) above, the endpoints of the pieces line up, and so $p_{Z|N}(\cdot|k)$ is computable.

Because $\mathbb{P}\{N = k\} = \frac{4}{5} \cdot 5^{-k}$ is a computable real, it remains to show that, as a function of x , the denominator $\sum_{n \in \mathbb{N}} p_{Z|N}(x|n) \cdot \mathbb{P}\{N = n\}$ of the right-hand side of Equation (70) is neither lower nor upper semicomputable on R . We will show the former; the latter follows in a similar fashion.

Define, for each $\ell \in \mathbb{N}$ the function $\psi_\ell : [0, 1] \rightarrow \mathbb{R}$ by

$$\psi_\ell(x) = 5^\ell \sum_{n \geq \ell} 3 p_{Z|N}(x|n) \cdot \frac{5}{4} \mathbb{P}\{N = n\}. \quad (71)$$

Note that ψ_0 is equal to $\frac{15}{4}$ times the denominator. Hence, it suffices to show that ψ_0 is not lower semicomputable on R .

We now prove by induction, uniformly in ℓ , that, if ψ_ℓ is lower semicomputable on R , then $\psi_{\ell+1}$ is lower semicomputable on R , and we can compute whether or not $h(\ell)$ is finite. This then implies that if ψ_0 were lower semicomputable on R , then $\emptyset'$ would be computable, a contradiction.

Suppose that ψ_ℓ is lower semicomputable on R . Then, we can compute (as a c.e. open subset of $[0, 1]$) a set S such that $S \cap R = \psi_\ell^{-1}[(3, \infty)] \cap R$.

If ℓ is such that $h(\ell) = \infty$, then for $x \in [0, 1]$, we have $p_{Z|N}(x | \ell) = 1$, and so

$$\psi_\ell(x) = 3 + \frac{1}{5} \psi_{\ell+1}(x) > 3. \quad (72)$$

In particular, S will contain R , and hence have Lebesgue measure greater than $\frac{7}{8}$.

However, if ℓ is instead such that $h(\ell) < \infty$, then $S \cap R \cap A_{h(\ell)}$ has Lebesgue measure at most $\frac{1}{2}$. Hence,

$$S \subseteq (S \cap R \cap A_{h(\ell)}) \cup ([0, 1] \setminus R) \cup ([0, 1] \setminus A_{h(\ell)}), \quad (73)$$

and so the Lebesgue measure of S is less than $\frac{1}{2} + (1 - \frac{7}{8}) + (1 - \frac{3}{4}) = \frac{7}{8}$.

Simultaneously (1) enumerate a c.e. sequence of basic open sets whose union is S , hence computing arbitrarily good lower bounds on its measure, and (2) run M_ℓ (on input 0). Either S will have measure greater than $\frac{7}{8}$ or M_ℓ will halt, but not both. Hence, we will eventually learn whether or not $h(\ell)$ is finite, which in turn shows that $p_{Z|N}(\cdot | \ell)$ is computable on $[0, 1]$, and so $\psi_{\ell+1}$ is also lower semicomputable on R . $\square$

As before, it follows immediately that $\mathbb{P}[N|Z = \cdot]$ is not computable, even on a $\mathbb{P}_Z$ -measure one set. It is possible to carry on the same development, showing that the conditional probability map is not computable as an element in $L^1(\mathbb{P}_Z)$. We state only the following strengthening of Corollary 7.7:

COROLLARY 8.5. *Let Φ be a conditioning operator for the set of probability distributions on pairs (X, Y) of random variables in $[0, 1]$ such that there exists an everywhere continuous version of the conditional distribution map $\mathbb{P}[Y|X = \cdot]$. Then Φ is noncomputable.*

9 POSITIVE RESULTS

Despite the fact that conditioning is not computable in general, many practitioners view conditioning as a routine operation involving the application of Bayes' rule. Indeed, below we show that conditioning is computable under mild assumptions when the observation is discrete, or when Bayes' rule is applicable, or when the observation is corrupted by independent "smooth" computable noise. We conclude by examining the setting where a notion of symmetry known as *exchangeability* holds.

9.1 Discrete Random Variables

We begin with the problem of conditioning on a random variable that takes values in a discrete set. Given an appropriate notion of computability for discrete sets, conditioning is always possible in this setting, as it reduces to the elementary notion of conditional probability with respect to single events.

Definition 9.1 (Computably Discrete Set). Let S be a computable Polish space. A subset $D \subseteq S$ is **computably discrete** when there exists a function $f : S \rightarrow \mathbb{N}$ that is computable and injective on D . In particular, such a D is countable. We call f the **witness** to the discreteness of D .

PROPOSITION 9.2. Let X and Y be $\mathbb{P}$ -almost computable random variables in computable Polish spaces S and T , respectively, let $D \subseteq S$ be a computably discrete subset with witness f , and assume that $\mathbb{P}\{X = d\} > 0$ for all $d \in D$. Then the conditional distribution map $\mathbb{P}[Y|X = \cdot]$ is computable on D , uniformly in X , Y , and f .

PROOF. Let $A \subseteq T$ be a $\mathbb{P}_Y$ -almost decidable set. By Lemma 4.7, it suffices to show that $\mathbb{P}[Y \in A|X = \cdot]$ is computable on D , uniformly in μ , f , and (the witness for the $\mathbb{P}_Y$ -almost decidability of) A . Let $x \in D$. Uniformly in f and x , we can find a witness to a $\mathbb{P}_X$ -almost decidable set $B_x \subseteq S$ such that $x \in B_x$ and no other element of D is in B_x . It then follows that

$$\mathbb{P}[Y \in A|X = x] = \mathbb{P}[Y \in A|X \in B_x] \text{ a.s.} \quad (74)$$

By Lemma 4.1, $\mathbb{P}[Y \in A|X \in B_x]$ is a computable real uniformly in μ and (the witness for the $\mathbb{P}_X$ -almost decidability of) B_x , hence uniformly in μ , f , and x . $\square$

The proof above relies on the ability to compute probabilities for continuity sets. In practice, computing probabilities can be inefficient. We give an alternative proof of Proposition 9.2 via a rejection-sampling argument, which yields an algorithm that can be much more efficient. We begin with a version of a well-known result.

LEMMA 9.3 (REJECTION SAMPLING). Let X and Y be random variables in computable Polish spaces S and T , respectively, let $B \subseteq S$ be a measurable set with positive $\mathbb{P}_X$ -measure, let $(X_0, Y_0), (X_1, Y_1), \dots$ be a sequence of i.i.d. copies of (X, Y) , and define ν to be the distribution of $(X_0, X_1, \dots)$. Further define $g : S^{\mathbb{N}} \rightarrow \mathbb{N}$ to be the map

$$(x_0, x_1, \dots) \mapsto \inf \{n \in \mathbb{N} : x_n \in B\} \quad (75)$$

and set

$$\zeta := g(X_0, X_1, \dots). \quad (76)$$

Then Y_ζ has distribution $\mathbb{P}\{Y \in \cdot | X \in B\}$ and, if B is a $\mathbb{P}_X$ -almost decidable set, then g is ν -almost computable, uniformly in a witness for B .

PROOF. For all $n \in \mathbb{N}$, we have $\mathbb{P}\{\zeta = n\} = \mathbb{P}\{X_n \in B\} \prod_{i=0}^{n-1} \mathbb{P}\{X_i \notin B\} = \mathbb{P}\{X \in B\} (1 - \mathbb{P}\{X \in B\})^n$. Summing this expression over all $n \in \mathbb{N}$, we obtain

$$\mathbb{P}\{\zeta < \infty\} = 1. \quad (77)$$

For $\mathbb{P}_\zeta$ -almost all $n \in \mathbb{N}$,

$$\mathbb{P}\{Y_\zeta \in \cdot | \zeta = n\} = \mathbb{P}\{Y_n \in \cdot | X_n \in B \text{ and } X_j \notin B \text{ for all } j < n\} \quad (78)$$

$$= \mathbb{P}\{Y_n \in \cdot | X_n \in B\} = \mathbb{P}\{Y \in \cdot | X \in B\}, \quad (79)$$

where the first equality follows from the definition of ζ , the second from the i.i.d. property, and the last by definition. As the right-hand side has no dependence on n , it follows from the chain rule of conditional expectation and Equation (77) that $\mathbb{P}\{Y_\zeta \in \cdot\} = \mathbb{P}\{Y \in \cdot | X \in B\}$.

We now establish the ν -almost computability of g . It follows from $\mathbb{P}\{\zeta < \infty\} = 1$ that g is $\mathbb{N}$ -valued on a ν -measure one set. By definition, the characteristic function for B , written $\chi_B : S \rightarrow \{0, 1\}$, is $\mathbb{P}_X$ -almost computable uniformly in a witness for the almost decidability of B . Hence, $g(x_0, x_1, \dots) = \min\{n \in \mathbb{N} : \chi_B(x_n) = 1\}$ on a ν -measure one set $A \subseteq S^{\mathbb{N}}$ of input sequences. But then $A \cap g^{-1}(n) = A \cap ((\chi_B^{-1}(0))^{n-1} \times \chi_B^{-1}(1) \times S^{\mathbb{N}})$, and so g is ν -almost computable. $\square$

We may now give an alternative proof of Proposition 9.2.

PROOF OF PROPOSITION 9.2 (VIA REJECTION SAMPLING). Let μ be the distribution of (X, Y) , which is computable. Uniformly in μ , we may compute a sequence $(X_0, Y_0), (X_1, Y_1), \dots$ of independent $\mathbb{P}$ -almost computable random variables with distribution μ . Let f be a witness to the computable discreteness of D . Then $(f(X_0), Y_0), (f(X_1), Y_1), \dots$ is also a computable sequence of independent $\mathbb{P}$ -almost computable random variables. Uniformly in f , we can compute a sequence of disjoint $\mathbb{P}_X$ -continuity sets $\langle B_k \rangle_{k \in \mathbb{N}}$ such that $x \in B_{f(x)}$ for $\mathbb{P}_X$ -almost all x . For $k \in \mathbb{N}$, let g_k be the function from Lemma 9.3 with respect to the set B_k . Uniformly in k , the random variable $\zeta_k = g_k(f(X_0), f(X_1), \dots)$ is $\mathbb{P}$ -almost computable, and then so is the random variable Y_{ζ_k} . For $k \in \mathbb{N}$, let γ_k be the distribution of Y_{ζ_k} . By Lemma 9.3, $\gamma_k = \mathbb{P}\{Y \in \cdot \mid f(X) = k\}$ and γ_k is computable, uniformly in k , f , and μ . It follows that the map $\gamma : \mathbb{N} \rightarrow \mathcal{M}_1(T)$ given by $k \mapsto \gamma_k$ is computable, uniformly in f and μ . Note, however, that for $\mathbb{P}_X$ -almost all x , we have $\mathbb{P}[Y|X = x] = \mathbb{P}\{Y \in \cdot \mid f(X) = f(x)\}$. Therefore, $\mathbb{P}[Y|X = \cdot] = \gamma \circ f$ is $\mathbb{P}_X$ -almost computable, uniformly in f and μ . $\square$

Rejection sampling has been used to give informal semantics of conditioning operators for discrete random variables in probabilistic programming languages such as λ_o (Park et al. 2008) and Church (Goodman et al. 2008).

9.2 Continuous and Dominated Setting

The most common way to perform conditioning given a continuous random variable is via Bayes' rule, which requires the existence of a conditional density. Using the characterization of computable measures in terms of the computability of integration, we can characterize when Bayes' rule is computable.

PROPOSITION 9.4. *Let X and Y be $\mathbb{P}$ -almost computable random variables on computable Polish spaces S and T , and let $R \subseteq S$. If $p_{X|Y}(x|y)$ is a conditional density of X given Y that is positive, bounded, and computable on $R \times T$, then κ as defined in Equation (21) is computable on R . In particular, if R is a $\mathbb{P}_X$ -measure one subset, then κ is a $\mathbb{P}_X$ -almost computable version.*

PROOF. By Bayes' rule (Lemma 3.11), a version of $\mathbb{P}[Y|X = \cdot]$ is given by the ratio in Equation (21):

$$\kappa(x, B) = \frac{\int_B p_{X|Y}(x|y) \mathbb{P}_Y(dy)}{\int p_{X|Y}(x|y) \mathbb{P}_Y(dy)}, \quad B \in \mathcal{B}_T. \quad (80)$$

By Proposition 2.30, when B is a $\mathbb{P}_Y$ -almost decidable set, $x \mapsto \int_B p_{X|Y}(x|y) \mathbb{P}_Y(dy)$ is a computable function of x on R , uniformly in a witness for the $\mathbb{P}_Y$ -almost decidability of B . Further, we know that $\int p_{X|Y}(x|y) \mathbb{P}_Y(dy)$ is non-zero on a $\mathbb{P}_X$ -measure one set, and so the ratio $\kappa(\cdot, B)$ is a $\mathbb{P}_X$ -almost computable function, uniformly in a witness for the $\mathbb{P}_Y$ -almost decidability of B . Hence, for c.e. open sets E , the function $\kappa(\cdot, E)$ is lower semicomputable uniformly in a representation for E . Therefore, by Lemma 4.7, the function $\bar{\kappa}$ is computable, as desired. $\square$

Remark 9.5. One might likewise obtain an algorithm for conditioning in the context of Proposition 9.4 using the proof of Proposition 9.2 via computable rejection sampling. In particular, the following classical argument may be carried out computably: For $x \in R$, let E_x be a $\mathbb{P}_X$ -almost continuous random variable in $\{0, 1\}$ such that $\mathbb{P}[E_x = 1|Y = y] = \frac{1}{M} p_{X|Y}(x|y)$, where M satisfies $p_{X|Y}(x|y) < M$ for all $x \in R$ and $y \in T$. (Such an M exists because of the boundedness condition.)

Then, for every Borel $B \subseteq T$ and every $x \in R$, we have

$$\kappa(x, B) = \frac{\mathbb{P}\{Y \in B, E_x = 1\}}{\mathbb{P}\{E_x = 1\}} = \mathbb{P}\{Y \in B | E_x = 1\}, \quad (81)$$

providing the reduction to (classical) rejection sampling.

Because the computability result Proposition 9.4 was established in a way that obviously relativizes to any oracle, we now obtain a proof of its analogue for a continuous conditional density, as promised in Section 3.2.

PROOF OF PROPOSITION 3.12. Follows from the proof of Proposition 9.4 by relativizing with respect to an arbitrary oracle. $\square$

COROLLARY 9.6 (DENSITY AND INDEPENDENCE). *Let U, V , and X be $\mathbb{P}$ -almost computable random variables in computable Polish spaces, where X is independent of V given U . Assume that there exists a bounded and computable conditional density $p_{X|U}(x|u)$ of X given U . Then the kernel $\mathbb{P}[(U, V) | X = \cdot]$ is computable.*

PROOF. Let $Y = (U, V)$. Then $p_{X|Y}(x|(u, v)) = p_{X|U}(x|u)$ is the conditional density of X given Y (with respect to ν). The result then follows immediately from Proposition 9.4. $\square$

9.3 Conditioning on Noisy Observations

As an immediate consequence of Corollary 9.6, 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 9.7 (INDEPENDENT NOISE). *Let U and E be $\mathbb{P}$ -almost computable random variables in $\mathbb{R}$, and let V be a $\mathbb{P}$ -almost computable random variable in a computable Polish space. Define $X = U + E$. If $\mathbb{P}_E$ is absolutely continuous (with respect to Lebesgue measure) with a bounded computable density p_E , and E is independent of U and V , then the conditional distribution map $\mathbb{P}[(U, V) | X = \cdot]$ is computable.*

PROOF. We have that

$$p_{X|U}(x|u) = p_E(x - u) \quad (82)$$

is the conditional density of X given U (with respect to Lebesgue measure). The result then follows from Corollary 9.6. $\square$

Pour-El and Richards (1989, Ch. 1, Thm. 2) show that a twice continuously differentiable computable function has a computable derivative (despite the fact that Myhill (1971) exhibits a computable function from $[0, 1]$ to $\mathbb{R}$ whose derivative is continuous, but not computable). Therefore, noise with a sufficiently smooth computable distribution has a computable density, and by Corollary 9.7, if such noise has a bounded density, then an almost computable random variable corrupted by such noise still admits a computable conditional distribution map.

Furthermore, Corollary 9.7 implies that the conditional distribution of a random variable given a continuous observation cannot always be uniformly approximated using noise in the sense that one cannot computably tell how little noise must be present to obtain a given accuracy. For example, consider our main construction (Theorem 7.6) corrupted with noise $E = \sigma Z$, where Z is a standard Gaussian noise and $\sigma > 0$ determines the standard deviation. Even though, as $\sigma \rightarrow 0$, the conditional distribution given the corrupted observation converges weakly to the uncorrupted conditional distribution with $\sigma = 0$, the noncomputability of the uncorrupted conditional distribution implies that one cannot uniformly compute a value of σ from a desired bound on the error introduced to the conditional distribution corrupted by the noise σZ .

9.4 Exchangeable Setting

Freer and Roy (2010) show how to compute conditional distributions in the setting of *exchangeable sequences*, i.e., sequences of random variables whose joint distribution is invariant to permutations of the indices. 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 (Freer and Roy 2009, 2012) and covers a wide range of common scenarios in nonparametric Bayesian statistics (often where no conditional density exists).

ACKNOWLEDGMENTS

The authors would like to thank Jeremy Avigad, Henry Cohn, Leslie Kaelbling, Vikash Mansinghka, Arno Pauly, Hartley Rogers, Michael Sipser, and Joshua Tenenbaum for helpful discussions, Bill Thurston for a useful comment regarding Lemma 8.1, Martín Escardó for a helpful discussion regarding computable random variables, and Quinn Culver, Noah Goodman, Jonathan Huggins, Leslie Kaelbling, Bjørn Kjos-Hanssen, Vikash Mansinghka, Timothy O'Donnell, Geoff Patterson, and the anonymous referees for comments on earlier drafts.

REFERENCES

- Nathanael L. Ackerman, Cameron E. Freer, and Daniel M. Roy. 2010. On the computability of conditional probability. Retrieved from <http://arxiv.org/abs/1005.3014v1>.
- Nathanael L. Ackerman, Cameron E. Freer, and Daniel M. Roy. 2011. Noncomputable conditional distributions. In *Proceedings of the 26th IEEE Symposium on Logic in Computer Science (LICS'11)*. IEEE Computer Society, 107–116.
- Shai Ben-David, Benny Chor, Oded Goldreich, and Michael Luby. 1992. On the theory of average case complexity. *J. Comput. System Sci.* 44, 2 (1992), 193–219.
- Jens Blanck. 1997. Domain representability of metric spaces. *Ann. Pure Appl. Logic* 83, 3 (1997), 225–247.
- Volker Bosserhoff. 2008. Notions of probabilistic computability on represented spaces. *J. Univ. Comput. Sci.* 14, 6 (2008), 956–995.
- Mark Braverman. 2005. On the complexity of real functions. In *Proceedings of the 46th Symposium on Foundations of Computer Science*. IEEE Computer Society, 155–164.
- Mark Braverman and Stephen Cook. 2006. Computing over the reals: Foundations for scientific computing. *Notices Amer. Math. Soc.* 53, 3 (2006), 318–329.
- Gregory F. Cooper. 1990. The computational complexity of probabilistic inference using Bayesian belief networks. *Artific. Intell.* 42, 2–3 (1990), 393–405.
- Paul Dagum and Michael Luby. 1993. Approximating probabilistic inference in Bayesian belief networks is NP-hard. *Artific. Intell.* 60, 1 (1993), 141–153.
- Paul Dagum and Michael Luby. 1997. An optimal approximation algorithm for Bayesian inference. *Artific. Intell.* 93, 1–2 (1997), 1–27.
- François G. Dorais, Gerald Edgar, and Jason Rute. 2013. Continuity on a measure one set versus measure one set of points of continuity. *MathOverflow*. Retrieved from <http://mathoverflow.net/q/146063> (version: 2013-10-27).
- Abbas Edalat. 1996. The Scott topology induces the weak topology. In *Proceedings of the 11th IEEE Symposium on Logic in Computer Science*. IEEE Computer Society Press, Los Alamitos, CA, 372–381.
- Abbas Edalat and Reinhold Heckmann. 1998. A computational model for metric spaces. *Theoret. Comput. Sci.* 193, 1–2 (1998), 53–73.
- Cameron E. Freer and Daniel M. Roy. 2009. Computable exchangeable sequences have computable de Finetti measures. In *Mathematical Theory and Computational Practice: Proceedings of the 5th Conference on Computability in Europe (CiE 2009)*. (Lecture Notes in Computer Science), Klaus Ambos-Spies, Benedikt Löwe, and Wolfgang Merkle (Eds.), Vol. 5635. Springer, 218–231.
- Cameron E. Freer and Daniel M. Roy. 2010. Posterior distributions are computable from predictive distributions. In *Proceedings of the 13th International Conference on Artificial Intelligence and Statistics (AISTATS'10)*. Y. W. Teh and M. Titterton (Eds.). JMLR: W&CP 9 (2010), 233–240.
- Cameron E. Freer and Daniel M. Roy. 2012. Computable de Finetti measures. *Ann. Pure Appl. Logic* 163, 5 (2012), 530–546.

- Peter Gács. 2005. Uniform test of algorithmic randomness over a general space. *Theoret. Comput. Sci.* 341, 1–3 (2005), 91–137.
- Stefano Galatolo, Mathieu Hoyrup, and Cristóbal Rojas. 2010. Effective symbolic dynamics, random points, statistical behavior, complexity and entropy. *Inform. and Comput.* 208, 1 (2010), 23–41.
- E. Mark Gold. 1967. Language identification in the limit. *Inform. and Control* 10, 5 (1967), 447–474.
- Noah D. Goodman, Vikash K. Mansinghka, Daniel M. Roy, Keith Bonawitz, and Joshua B. Tenenbaum. 2008. Church: A language for generative models. In *Proceedings of the 24th Conference on Uncertainty in Artificial Intelligence*.
- Armin Hemmerling. 2002. Effective metric spaces and representations of the reals. *Theoret. Comput. Sci.* 284, 2 (2002), 347–372.
- Mathieu Hoyrup. 2008. *Computability, Randomness and Ergodic Theory on Metric Spaces*. Ph.D. Dissertation. Université Paris Diderot (Paris VII), Paris, France.
- Mathieu Hoyrup and Cristóbal Rojas. 2009a. An application of Martin-Löf randomness to effective probability theory. In *Mathematical Theory and Computational Practice*, Klaus Ambos-Spies, Benedikt Löwe, and Wolfgang Merkle (Eds.). Lecture Notes in Computer Science, Vol. 5635. Springer, 260–269.
- Mathieu Hoyrup and Cristóbal Rojas. 2009b. Computability of probability measures and Martin-Löf randomness over metric spaces. *Inform. and Comput.* 207, 7 (2009), 830–847.
- Mathieu Hoyrup and Cristóbal Rojas. 2011. Absolute continuity of measures and preservation of randomness. (2011). Retrieved from <https://members.loria.fr/MHoyrup/abscont.pdf>.
- Mathieu Hoyrup, Cristóbal Rojas, and Klaus Weihrauch. 2011. Computability of the Radon-Nikodym derivative. In *Models of Computation in Context, Proceedings of the 7th Conference on Computability in Europe (CiE 2011)*, Benedikt Löwe, Dag Normann, Ivan N. Soskov, and Alexandra A. Soskova (Eds.). Lecture Notes in Computer Science, Vol. 6735. Springer, 132–141.
- Marcus Hutter. 2007. On universal prediction and Bayesian confirmation. *Theoret. Comput. Sci.* 384, 1 (2007), 33–48.
- Olav Kallenberg. 2002. *Foundations of Modern Probability* (2nd ed.). Springer, New York.
- Alexander S. Kechris. 1995. *Classical Descriptive Set Theory*. Graduate Texts in Mathematics, Vol. 156. Springer-Verlag, New York.
- Oleg Kiselyov and Chung-chieh Shan. 2009. Embedded probabilistic programming. In *Domain-Specific Languages*, Walid Mohamed Taha (Ed.). Lecture Notes in Computer Science, Vol. 5658. Springer, 360–384.
- Stephen Cole Kleene. 1938. On notation for ordinal numbers. *J. Symbolic Logic* 3, 4 (1938), 150–155.
- Donald E. Knuth and Andrew C. Yao. 1976. The complexity of nonuniform random number generation. In *Proceedings of the Symposium on Algorithms and Complexity*. Academic Press, New York, 357–428.
- A. N. Kolmogorov. 1933. *Grundbegriffe der Wahrscheinlichkeitsrechnung*. Springer.
- Leonid A. Levin. 1986. Average case complete problems. *SIAM J. Comput.* 15, 1 (1986), 285–286.
- Irwin Mann. 1973. Probabilistic recursive functions. *Trans. Amer. Math. Soc.* 177 (1973), 447–467.
- T. Minka, J. M. Winn, J. P. Guiver, and D. A. Knowles. 2010. *Infer.NET 2.4*. Microsoft Research Cambridge. Cambridge, UK. Retrieved from <http://research.microsoft.com/infernet>.
- Kenshi Miyabe. 2013. L^1 -computability, layerwise computability and Solovay reducibility. *Computability* 2, 1 (2013), 15–29.
- Yiannis N. Moschovakis. 2009. *Descriptive Set Theory* (2nd ed.). Mathematical Surveys and Monographs, Vol. 155. American Mathematical Society, Providence, RI.
- Norbert Th. Müller. 1999. Computability on random variables. *Theoret. Comput. Sci.* 219, 1–2 (1999), 287–299.
- J. Myhill. 1971. A recursive function, defined on a compact interval and having a continuous derivative that is not recursive. *Michigan Math. J.* 18 (1971), 97–98.
- André Nies. 2009. *Computability and Randomness*. Oxford Logic Guides, Vol. 51. Oxford University Press, Oxford.
- Daniel N. Osherson, Michael Stob, and Scott Weinstein. 1988. Mechanical learners pay a price for Bayesianism. *J. Symbolic Logic* 53, 4 (1988), 1245–1251.
- Sungwoo Park, Frank Pfenning, and Sebastian Thrun. 2008. A probabilistic language based on sampling functions. *ACM Trans. Program. Lang. Syst.* 31, 1 (2008), 1–46.
- J. Pfanzagl. 1979. Conditional distributions as derivatives. *Ann. Probab.* 7, 6 (1979), 1046–1050.
- Avi Pfeffer. 2001. IBAL: A probabilistic rational programming language. In *Proceedings of the 17th International Joint Conference on Artificial Intelligence*. Morgan Kaufmann Publishers, 733–740.
- David Poole. 1991. Representing Bayesian networks within probabilistic Horn abduction. In *Proceedings of the 7th Conference on Uncertainty in Artificial Intelligence*. 271–278.
- Marian B. Pour-El and J. Ian Richards. 1989. *Computability in Analysis and Physics*. Springer-Verlag, Berlin.
- Hilary Putnam. 1965. Trial and error predicates and the solution to a problem of Mostowski. *J. Symbolic Logic* 30 (1965), 49–57.
- M. M. Rao. 1988. Paradoxes in conditional probability. *J. Multivariate Anal.* 27, 2 (1988), 434–446.
- M. M. Rao. 2005. *Conditional Measures and Applications* (2nd ed.). Pure and Applied Mathematics, Vol. 271. Chapman & Hall/CRC.

- Matthew Richardson and Pedro Domingos. 2006. Markov logic networks. *Machine Learn.* 62, 1–2 (2006), 107–136.
- Hartley Rogers, Jr. 1987. *Theory of Recursive Functions and Effective Computability* (2nd ed.). MIT Press, Cambridge, MA.
- Daniel M. Roy. 2011. *Computability, Inference and Modeling in Probabilistic Programming*. Ph.D. Dissertation. Massachusetts Institute of Technology, Cambridge, MA.
- Mark J. Schervish. 1995. *Theory of Statistics*. Springer-Verlag, New York.
- Matthias Schröder. 2007. Admissible representations for probability measures. *Math. Log. Q.* 53, 4–5 (2007), 431–445.
- R. J. Solomonoff. 1964. A formal theory of inductive inference II. *Inform. and Control* 7 (1964), 224–254.
- Hayato Takahashi. 2008. On a definition of random sequences with respect to conditional probability. *Inform. and Comput.* 206, 12 (2008), 1375–1382.
- Tue Tjur. 1974. *Conditional Probability Distributions*. Institute of Mathematical Statistics, University of Copenhagen, Copenhagen, Denmark.
- Tue Tjur. 1975. *A Constructive Definition of Conditional Distributions*. Institute of Mathematical Statistics, University of Copenhagen, Copenhagen, Denmark.
- Tue Tjur. 1980. *Probability Based on Radon Measures*. John Wiley & Sons Ltd., Chichester, UK.
- Klaus Weihrauch. 1993. Computability on computable metric spaces. *Theoret. Comput. Sci.* 113, 2 (1993), 191–210.
- Klaus Weihrauch. 1999. Computability on the probability measures on the Borel sets of the unit interval. *Theoret. Comput. Sci.* 219, 1–2 (1999), 421–437.
- Klaus Weihrauch. 2000. *Computable Analysis: An Introduction*. Springer-Verlag, Berlin.
- Tomoyuki Yamakami. 1999. Polynomial time samplable distributions. *J. Complexity* 15, 4 (1999), 557–574.
- A. K. Zvonkin and L. A. Levin. 1970. The complexity of finite objects and the development of the concepts of information and randomness by means of the theory of algorithms. *Uspekhi Mat. Nauk* 25, 6 (156) (1970), 85–127.

Received December 2011; revised July 2017; accepted March 2019