

Deterministic stream-sampling for probabilistic programming: semantics and verification

Fredrik Dahlqvist
Queen Mary University of London
and University College London
London, United Kingdom
f.dahlqvist@qmul.ac.uk

Alexandra Silva
Department of Computer Science
Cornell University
Ithaca, USA
alexandra.silva@cornell.edu

William Smith
Department of Computer Science
University College London
London, United Kingdom
william.smith.19@ucl.ac.uk

Abstract—Probabilistic programming languages rely fundamentally on some notion of sampling, and this is doubly true for probabilistic programming languages which perform Bayesian inference using Monte Carlo techniques. Verifying samplers—proving that they generate samples from the correct distribution—is crucial to the use of probabilistic programming languages for statistical modelling and inference. However, the typical denotational semantics of probabilistic programs is incompatible with deterministic notions of sampling. This is problematic, considering that most statistical inference is performed using pseudorandom number generators.

We present a higher-order probabilistic programming language centred on the notion of *samplers* and *sampler operations*. We give this language an operational and denotational semantics in terms of continuous maps between topological spaces. Our language also supports discontinuous operations, such as comparisons between reals, by using the type system to track discontinuities. This feature might be of independent interest, for example in the context of differentiable programming.

Using this language, we develop tools for the formal verification of sampler correctness. We present an equational calculus to reason about equivalence of samplers, and a sound calculus to prove semantic correctness of samplers, i.e. that a sampler correctly targets a given measure by construction.

Index Terms—Probabilistic programming, operational and denotational semantics, verification

I. INTRODUCTION

Probabilistic programming languages without conditioning – that is to say, programming languages capable of drawing random samples – and the concepts of Monte Carlo methods and randomized algorithms have been around as long as true computers have¹; however, the introduction of languages with conditioning, higher-order features, continuous variables, recursion, and their application to statistical modelling and machine learning, is a product of the twenty-first century [36], [26], [15], [27], [40], [7], [3]. Since a probabilistic programming language with conditioning must come equipped with a range of inference algorithms and sampling methods, and since the rate of introduction of these has increased in

This work was supported by the Leverhulme Project Grant “Verification of Machine Learning Algorithms”.

¹See [4] for an overview of probabilistic programming languages, [6] for a historical overview of the Monte Carlo method, and [23] for a history of pseudorandom number generation.

recent years, new formal methods must be developed for the *verification* of these algorithms.

We aim to make the verification of inference algorithms straightforward by introducing a language endowed with a sampler type, featuring many of the sampler operations used in these inference algorithms, and a calculus for reasoning about correctness of these samplers relative to intended ‘target’ distributions. The semantics of our language is fully deterministic, in order to allow the use of deterministic – pseudorandom – samplers, in a simple manner and without paradox.

When assigning operational and denotational semantics to probabilistic programs, an interesting asymmetry emerges: the simplest reasonable denotational semantics of a first-order language with continuous datatypes is in terms of probability measures [21], while the simplest reasonable operational semantics is in terms of sampled *values* – the latter being much closer to the intuitions used by programmers of languages with the ability to draw samples.

The denotational semantics of probabilistic programming languages, in terms of measures (broadly construed) is well-understood [21], [16], [12], [38], [8]. The most common approaches to the operational semantics of such a language, as highlighted by [12], are *trace semantics* and *Markov chain semantics*. The latter is the chosen operational semantics for a number of probabilistic λ -calculi [11], [22], [5], [13], [12], [14] and probabilistic languages [34], [38], but does not speak of sampled values, only of distributions on execution paths. From the perspective of the asymmetry described above, it is thus closer to a denotational semantics.

Trace semantics, originally developed in [21] and later applied in [32], [5], [9], [2], assumes that for each distribution in the language, an infinite set of samples has been produced ahead-of-time. When a sample is requested, the head of this sequence is popped and used in the computation, and the tail of the sequence is kept available for further sampling. This perspective models samplers, such as `rand()`, as functions with hidden side effects on the state of the machine – in line with a programmer’s intuition on the nature of sequential calls to `rand()`. The natural notion of adequacy with respect to the denotational semantics is to show that subject to the assumption that each element of these sequences is sampled independently from its corresponding distribution, the resulting pushforward

through the program is identical to the program’s denotational semantics. We see two issues with this approach.

First, the supposition that all samples are pre-computed ahead-of-time is incompatible with pseudorandom generation of ‘random’ values, since computationally-generated samples and truly-random samples are in fact distinguishable. For example, if (x_0, x_1, \dots) is a sequence of samples targeting the distribution P which was produced via iteration of a computable map $x_{n+1} = T(x_n)$, then the program $T(x) - x$ will behave differently if x is a ‘truly random’ sample from P than if x is produced by the aforementioned iterative procedure, and so the operational and denotational semantics no longer cohere; similar counterexamples exist for any pseudorandom number generator. If (x_0, x_1, \dots) is a deterministic sequence, meaningful coherence between the operational and denotational sequence can only be assured if it is assumed that this sequence is *Martin-Löf random* (first defined in [29], later generalised to computable metric spaces in [17]); unfortunately, all Martin-Löf random sequences are uncomputable. Pseudorandom numbers are in fact used in simulation far more commonly than ‘true’ physical randomness, as they are typically faster to obtain, have the advantage of being reproducible given a particular seed, and, subject to certain assumptions, can even have better convergence properties [24]. We find the inability of trace semantics to describe pseudorandom number generation to be a significant weakness.

Second, from the trace semantics perspective, the notion of a sampler type is inextricably bound up in issues regarding side-effects, which makes verification challenging. In order to properly assign the syntax `rand`, without parentheses, a meaningful semantics as a function, we must give it a monadic interpretation as in [32], accounting for its hidden effect on the trace. The correctness of code which inputs and outputs samplers – sampler operations – is then subject to the state of the trace when computation is started, which is not contained within either the code of the sampler operation in question or the code of the samplers it inputs. This pattern, of using sampler operations to create composite samplers which make use of other ‘primitive’ samplers, is a central theme in computational statistics. In particular, inference algorithms within Bayesian statistics, to which probabilistic programming languages with conditioning compile, make heavy use of this technique. Commonly-used sampling methods, such as importance sampling, rejection sampling, and particle Markov chain Monte Carlo methods, are naturally understood as composite samplers of this type [1], [33]. We prefer a side-effect-free perspective from which the correctness of sampler operations can be demonstrated subject to assumptions about the samplers input to these programs, as opposed to a perspective in which their correctness depends also on the state of the machine on which these operations are run.

Contributions. We develop a language based around the idea of sampler types and sampler operations, which allows reasoning about deterministic and real-valued samplers, and which is designed to make verification of these samplers natural. A syntax, operational semantics, and denotational semantics for

our language are introduced in § III, and an adequacy result relating them is shown. § IV lays out a notion of equivalence of samplers, which will be applied to simplify programs. Finally, in § V, we discuss methods for proving that samplers target the desired probability measure (i.e. verifying samplers), and introduce a sound calculus for verifying the correctness of composite samplers which is capable of demonstrating the soundness of common Monte Carlo techniques such as importance sampling and rejection sampling.

II. EXAMPLES

The purpose of this section is twofold. First, we present examples of how samplers are transformed in order to create new samplers. We use this opportunity to informally introduce a language with a sampler type constructor Σ and operations for constructing and manipulating samplers. Second, we present techniques to reason about the correctness of sampling algorithms. These come in two flavours. We reason *equationally* about the equivalence between samplers (see § IV), and we reason *semantically* about whether a sampler does the job it is designed to do – namely, generate deviates from a target distribution (see § V).

A. Von Neumann extractor

We begin with a simple family of discrete samplers known as von Neumann extractors. This example will illustrate the concept of a sampler’s *self-product*, a central concept for a language featuring sampler types. The von Neumann extractor [37] is a simple procedure which, given flips from a *biased* coin on $\{\text{True}, \text{False}\}$ with probability $p \in (0, 1)$ of landing `True`, produces flips from an *unbiased* coin with probability $1/2$ of landing `True`. We view this as a sampler v of Boolean type – notation $v : \Sigma B$ – which, given another Boolean-valued sampler `flip` : ΣB representing our biased coin, constructs an unbiased Boolean-valued sampler. A simple implementation of the von Neumann extractor is given in Listing 1.

```
let choice = λb : B × B .
  if (fst(b) and snd(b)) or (not fst(b) and
    not snd(b)) then 0 else 1
in let proj = λb : B × B . fst(b)
  in map(proj, reweight(choice, flip2))
```

Listing 1: von Neumann extractor

The idea behind this algorithm is that if (b_1, b_2) are sampled independently from a Bernoulli distribution with parameter p , then the probabilities of the outcomes $(b_1, b_2) = (\text{False}, \text{True})$ and $(b_1, b_2) = (\text{True}, \text{False})$ are both $p(1-p)$, and so the first element b_1 of each pair is an unbiased flip. In Listing 1, samples (b_1, b_2) where $b_1 = b_2$ are removed by the `reweight` operation, which sets the *weight* of such pairs to zero; then, the command `map` applies the function `proj` to each pair (b_1, b_2) , returning the sampler whose outputs are only the first element b_1 .

Note the appearance of `flip2` in the von Neumann extractor; this is the ‘self-product’ of the sampler `flip`. Where `flip` : ΣB produces Boolean-valued samples, `flip2` : $\Sigma (B \times B)$ produces samples which are pairs of

$$\frac{\frac{\frac{\vdash \text{flip}^2 : \Sigma(\mathbb{B} \times \mathbb{B}) \rightsquigarrow \text{Ber}(p)^2 \quad \vdash \text{choice} : \mathbb{B} \times \mathbb{B} \rightarrow \mathbb{R}^+}{\vdash \text{reweight}(\text{choice}, \text{flip}^2) : \Sigma(\mathbb{B} \times \mathbb{B}) \rightsquigarrow \llbracket \text{choice} \rrbracket_* \cdot \text{Ber}(p)^2} \quad \vdash \text{proj} : \mathbb{B} \times \mathbb{B} \rightarrow \mathbb{B}}{\vdash \text{map}(\text{proj}, \text{reweight}(\text{choice}, \text{flip}^2)) : \Sigma \mathbb{B} \rightsquigarrow \llbracket \text{proj} \rrbracket_* (\llbracket \text{choice} \rrbracket_* \cdot \text{Ber}(p)^2) = \text{Ber}(1/2)}$$

Fig. 1: Validity of von Neumann extractor

$$\frac{\frac{\frac{\vdash \text{rand}^2 : \Sigma(\mathbb{R} \times \mathbb{R}) \rightsquigarrow U^2 \quad \vdash \text{plus} : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}}{\vdash \text{map}(\text{plus}, \text{rand}^2) : \Sigma \mathbb{R} \rightsquigarrow \llbracket \text{plus} \rrbracket_* U^2} \quad \vdash \text{phi} : \mathbb{R} \rightarrow \mathbb{R}^+}{\vdash \text{reweight}(\text{phi}, \text{map}(\text{plus}, \text{rand}^2)) : \Sigma \mathbb{R} \rightsquigarrow \llbracket \text{phi} \rrbracket_* (\llbracket \text{plus} \rrbracket_* U^2) = P}$$

Fig. 2: Validity of importance sampling in Listing 2

Booleans. Given a sampler $t : \Sigma T$ of type T , the self-product $t^2 : \Sigma(T \times T)$ is a sampler whose elements are adjacent samples from T ; the same construction, detailed in § III, is easily extended to arbitrary self-powers $t^K : \Sigma(T^K)$.

The main application of our language is to serve as a setting for the formal verification of its samplers. Let $v : \Sigma \mathbb{B}$ be the von Neumann extractor defined in Listing 1; the task of verifying v is the task of showing that v ‘targets’ the uniform distribution $\text{Ber}(1/2)$ – meaning, informally, that in the limit of increasing sample size, v generates unbiased flips. In § V-B, we define a relation \rightsquigarrow between samplers and distributions which reifies this notion: we read $\vdash v : \Sigma \mathbb{B} \rightsquigarrow \text{Ber}(1/2)$ as ‘the sampler v targets the measure $\text{Ber}(1/2)$ ’. The aforementioned self-product operation plays a crucial role in sampler verification, as it does *not* suffice, in order to conclude that v targets $\text{Ber}(1/2)$, to assume that flip targets $\text{Ber}(p)$ for some $p \in (0, 1)$. Instead, we are required to make the stronger assumption that flip^2 targets $\text{Ber}(p)^2$: in essence, that adjacent samples from flip act as if they are independent. Following the literature on pseudorandom number generation, we refer to this property as *K-equidistribution* (in this case, for $K = 2$): we will say that a sampler s is *K-equidistributed* with respect to the distribution P if s^K targets P^K . For example, the assertion that a pseudorandom number generator which takes values in $\{0, \dots, N - 1\}$ is *K-equidistributed* with respect to the uniform distribution is the assertion that all K -length words $w \in \{0, \dots, N - 1\}^K$ are produced in equal proportion. Several commonly-used discrete PRNGs, such as xorshift and the Mersenne twister, have well-known *K-equidistribution* guarantees; see [39], [30].

In our calculus for asymptotic targeting, the validity of the von Neumann extractor is shown in Fig. 1. Before we begin this derivation, it must be shown that the von Neumann extractor v is equivalent to the simplified sampler $\text{map}(\text{proj}, \text{reweight}(\text{choice}, \text{flip}^2))$ in a context in which access to a Boolean-typed sampler flip is assumed, where proj and choice are defined as in Listing 1. We write this equivalence as $\text{flip} : \Sigma \mathbb{B} \vdash v \approx \text{map}(\text{proj}, \text{reweight}(\text{choice}, \text{flip}^2)) : \Sigma \mathbb{B}$; this equivalence relation is discussed in § IV, and is shown in this particular case using the let-binding rule of Table IV.

Having rewritten v in this way, and using the hypothesis of 2-equidistribution $\vdash \text{flip}^2 : \Sigma \mathbb{B} \rightsquigarrow \text{Ber}(p)^2$, we derive our conclusion in Fig. 1 by applying the rules from § V-B corre-

sponding to the sampler operations reweight and map . These rules show us that v targets the measure $\llbracket \text{proj} \rrbracket_* (\llbracket \text{choice} \rrbracket_* \cdot \text{Ber}(p)^2)$. (Here, as we will discuss in § V-B, $\llbracket f \rrbracket_* \mu$ denotes the pushforward of the measure μ through the function f , and the notation $f \cdot \mu$ denotes the measure μ reweighted by the density f .) To complete the proof, we show that this measure is identical to $\text{Ber}(1/2)$, the uniform measure on \mathbb{B} ; this is straightforward. It is easily seen first that $\llbracket \text{choice} \rrbracket_* \cdot \text{Ber}(p)^2$ assigns probability $1/2$ to the samples $(\text{True}, \text{False})$ and $(\text{False}, \text{True})$ and zero probability to all other samples; the desired result then follows by observing that the function proj simply drops the second sample.

B. Importance sampling

A central application of the reweighting operation is its role in *importance sampling*. This is a commonly used technique [33], [6] in Bayesian learning and statistical inference, which transforms samples from a ‘proposal’ distribution Q on the latent space X into approximate samples from a ‘target’ distribution P , where P is absolutely continuous with respect to Q with Radon-Nikodym derivative $\frac{dP}{dQ}(x)$. Its operation is straightforward: for each sample $x_n \sim Q$, compute the sample’s weight $w_n = \frac{dP}{dQ}(x)$, and then the *weighted sample* (x_n, w_n) is informally understood as an approximate sample from the target P . Formally, the normalised empirical measure $\frac{\sum_{n=1}^N w_n \delta_{x_n}}{\sum_{i=1}^N w_i}$, where δ_x is the Dirac measure at $x \in X$, converges weakly as $N \rightarrow \infty$ to the target measure P .

For example, consider the Bayesian inference problem in which the prior P_0 is the triangular distribution on $[0, 2]$ and the likelihood of the observation $y = 3$ given the latent value x is a standard Gaussian $L(x) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{(3-x)^2}{2}\right)$. Let P represent the corresponding posterior distribution, whose density is proportional to the pointwise product of the triangular and Gaussian densities; a simple importance-sampling procedure for sampling from P in our language is shown in Listing 2. Here, we assume access to a sampler rand which targets the uniform distribution on $[0, 1]$; recalling that a triangular random variable is the sum of two independent uniform random variables, and assuming rand has the necessary independence property of 2-equidistribution, we sum two draws from rand to produce a triangular random variable. Finally, we *reweight* the result according to the likelihood $L(x)$, yielding a sampler which targets the posterior distribution P corresponding to the observed datum $y = 3$.

$$\frac{\frac{\vdash \text{tri} \otimes \text{rand} : \Sigma T \rightsquigarrow \text{Tri} \otimes U \quad \vdash \text{accept} : T \rightarrow \mathbb{R}^+}{\vdash \text{reweight}(\text{accept}, \text{tri} \otimes \text{rand}) : \Sigma T \rightsquigarrow \llbracket \text{accept} \rrbracket \cdot (\text{Tri} \otimes U)} \quad \vdash \text{proj} : T \rightarrow \mathbb{R}}{\vdash \text{map}(\text{proj}, \text{reweight}(\text{accept}, \text{tri} \otimes \text{rand})) : \Sigma \mathbb{R} \rightsquigarrow \llbracket \text{proj} \rrbracket_* (\llbracket \text{accept} \rrbracket \cdot (\text{Tri} \otimes U)) = P}$$

Fig. 3: Validity of rejection sampling in Listing 3

```

let phi = λx : ℝ . 1/sqrt(2*pi) * exp(-1/2*(3-x)
  * (3-x))
in let plus = λu : ℝ × ℝ . fst(u) + snd(u)
  in reweight(phi, map(plus, rand2))

```

Listing 2: Importance sampling

The validity of this sampler – i.e. the fact that it targets the correct posterior distribution – follows easily in our targeting calculus. Under the hypothesis that `rand` produces 2-equidistributed samples with respect to the uniform distribution U , the derivation Fig. 2 proves that the sampler defined in Listing 2 targets the measure $\llbracket \text{phi} \rrbracket \cdot (\llbracket \text{plus} \rrbracket_* U^2)$. It remains to show that this measure is the desired P ; once one shows that the sum of two independent uniform variates is triangular, this follows by definition of the reweighting operation \cdot . The same argument suffices for any observation y .

C. Rejection sampling

Our final example of sampler verification is an instance of the technique known as *rejection sampling*. Listing 3 applies rejection sampling from the prior to the same Bayesian inference problem discussed in § II-B to yield a sampler which targets the same posterior distribution P . Of particular importance is the *discontinuity* of the accept-reject step, which significantly complicates the argument of sampler verification in the presence of pseudorandom number generation.

```

let phi = λx : ℝ . 1/sqrt(2*pi) * exp(-1/2*(3-x)
  * (3-x))
in let accept = λ(u,v) : T .
  if v ≤ phi(u)*sqrt(2*pi) then 1 else 0
  in let proj = λz : T. fst(cast(ℝ×ℝ)(z)) in
  map(proj, reweight(accept, tri ⊗ rand))

```

Listing 3: Rejection sampling

In order to show the validity of rejection sampling from the prior, we must assume access to a sampler on the prior distribution (here `tri`), an independent standard uniform random sampler (here `rand`), and an upper bound $\sup_{x \in \mathbb{R}} L(x) = \frac{1}{2\pi} \sup_{x \in \mathbb{R}} \exp(-\frac{(3-x)^2}{2}) = \frac{1}{2\pi}$ on the likelihood, which is used in the acceptance condition. Fig. 3 proves that, subject to the natural independence assumption for the samplers `tri` and `rand`, the rejection sampler defined by Listing 3 targets the measure $\llbracket \text{proj} \rrbracket_* (\llbracket \text{accept} \rrbracket \cdot (\text{Tri} \otimes U))$. We can then show, using standard methods, that this measure is identical to P , the posterior distribution also targeted by Listing 2.

We have omitted, for the moment, one crucial part of the proof. Note that, in both Listing 3 and Fig. 3, the function `accept`, which one might expect to have type $\mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}^+$, instead has type $T \rightarrow \mathbb{R}^+$. Correspondingly, the product `tri` \otimes `rand` must be assumed to produce samples of

type T , rather than $\mathbb{R} \times \mathbb{R}$, and `proj` must accept inputs of type T rather than pairs $\mathbb{R} \times \mathbb{R}$. The nature of this type T , a *subtype* of $\mathbb{R} \times \mathbb{R}$, will be explained in § III, but it encodes the fact that `accept` is discontinuous when viewed as a function on the standard topologies, as well as where those discontinuities are allowed to lie. The type-inference of Listing 3, detailing the structure of T , is given in [10, Fig. 6 and Fig. 7].

III. LANGUAGE

A. Syntax

We use a λ -calculus with a notion of subtype and a type constructor Σ for samplers.

1) *Types*: Types are generated by the mostly standard grammar in Fig. 4a, where the set `Ground` of ground types is

$$\{\mathbb{N}, \mathbb{R}, \mathbb{R}^+\} \cup \{f^{-1}(i) \mid f \in \{\leq, <, \geq, >, =, \neq\}, i = 0, 1\}.$$

Our ground types include the natural, real and nonnegative real numbers, as well as important sets of pairs of reals: for example, $<^{-1}(1)$ will be denoted, as the notation suggests, by the pairs of reals whose first component is strictly smaller than the second. The boolean type $\mathbb{B} \triangleq 1 + 1$ will be treated as a ground type.

The only unusual type constructors are the *pullback types* ${}_s T_t$ which – as the name suggests – will be interpreted as pullbacks (in fact inverse images), and the *sampler types* ΣT which will be defined as the coinductive (stream) types defined by the (syntactic) functors $T \times \mathbb{R}^+ \times -$. In other words, we assume that samplers can be weighted; this covers the special case of unweighted samplers, in which every weight is set to 1. As these are the only coinductive types we need, and to highlight the central role played by samplers, we choose not to add generic coinductive types to the language.

The *subtyping relation* \triangleleft on types is the reflexive transitive closure of the relation generated by the rules of Fig. 4b.

2) *Terms*: Fig. 4c presents the grammar generating the set `Expr` of terms in our language. We assume the existence of a set `Func` of built-in functions which come equipped with typing information $f : T \rightarrow G$, where G is a ground type. Some built-in functions will be continuous w.r.t. to the usual topologies, such as the addition operation $+$: $\mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$, but others will be discontinuous, such as the comparison operators $\{\leq, <, \geq, >, =, \neq\} : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{B}$. Dealing with such functions is the main reason for adding coproducts to the grammar, as we will discuss in § III-C. We also employ the syntactic sugar

$$\text{if } b \text{ then } s_{\text{True}} \text{ else } s_{\text{False}} \triangleq \text{case } (b, _) \text{ of } \{(i, _) \Rightarrow s_i\}_{i \in \mathbb{B}}.$$

Most of our language constructs are standard for a typed functional language without recursion, but we endow our language with several nonstandard (sampler) operations:

$S, T ::= G \in \text{Ground} \mid 1 \mid S \times T \mid S + T \mid {}_s T_t \mid S \rightarrow T \mid \Sigma T \quad s, t : T$
(a) Type grammar

$$\frac{}{f^{-1}(0) + f^{-1}(1) \triangleleft \mathbb{R} \times \mathbb{R}} \quad f \in \{<, \leq, >, \geq, =, \neq\} \quad \frac{S_1 \triangleleft S_2 \quad T_1 \triangleleft T_2}{S_1 \times T_1 \triangleleft S_2 \times T_2} \quad \frac{S_1 \triangleleft S_2 \quad T_1 \triangleleft T_2}{S_1 + T_1 \triangleleft S + T_2}$$

$$\frac{S \triangleleft T \quad \sum_{i \in n} S_i \triangleleft S \quad \sum_{j \in m} S'_j \triangleleft S}{\Sigma S \triangleleft \Sigma T} \quad \frac{\sum_{i \in n, j \in m} S_i \cap S'_j \triangleleft \sum_{i \in n} S_i}{\sum_{i \in n, j \in m} S_i \cap S'_j \triangleleft \sum_{j \in m} S_j}$$

(b) Subtyping rules

$t ::= x \in \text{Var} \mid b \in \{\text{True}, \text{False}\} \mid n \in \mathbb{N} \mid r \in \mathbb{R} \mid$ *Variables and constants*
 $f(t, \dots, t), f \in \text{Func} \mid \text{cast}(T)t \mid$ *Built-in functions*
 $\text{case } t \text{ of } \{(i, x_i) \Rightarrow s_i\}_{i \in n} \mid \text{in}_i(t) \mid \lambda x: T. t \mid t(t) \mid \text{let } x = t \text{ in } t \mid$ *Programming constructs*
 $(t, t) \mid \text{fst}(t) \mid \text{snd}(t) \mid$ *Products*
 $\text{prng}(t, t) \mid t \otimes t \mid \text{map}(t, t) \mid \text{reweight}(t, t) \mid \text{hd}(t) \mid \text{wt}(t) \mid \text{tl}(t) \mid \text{thin}(t, t)$ *Sampler operations*

(c) Term grammar

Fig. 4: Grammars and subtyping rules

- The operation $\text{prng}(f, t)$ is used to construct a sampler as a pseudo-random number generator, using an initial value t and a deterministic endomap f .
- $s \otimes t$ represents the product of samplers s, t .
- The syntax $\text{map}(f, t)$ maps the function f over the elements produced by the sampler t to produce a new sampler, in analogy to the pushforward of a measure.
- The operation $\text{reweight}(f, t)$ applies the reweighting scheme f to the sampler t to form a new sampler.
- Given a sampler t , the operation $\text{hd}(t)$ returns the first sample produced by t , $\text{wt}(t)$ the weight of the first sample produced by t , and $\text{tl}(t)$ returns the sampler t but with its first sample-weight pair dropped.
- The operation $\text{thin}(n, t)$, given a natural number n and a sampler t , returns the sampler which includes only those elements of t whose index is a multiple of n .

The intuition and purposes of most of these language constructs was explained in § II, and their precise meaning will be made clear when we introduce their semantics.

3) *Well-formed terms*: Our typing system is mostly standard and presented in Table I. The only non-standard rules are the *context-restriction rule* on the second line of Table I, and the typing rules for the sampler operations, which should be straightforward given their descriptions above. The purpose of the context-restriction rule is, in a nutshell, to be able to pass the result of a computation of type T which is continuous w.r.t. a topology τ on the denotation of T , to a computation using a variable of type T but which is continuous w.r.t. to a finer topology $\tau' \supset \tau$ on the denotation of T . After application of this rule, it is no longer possible to λ -abstract on the individual variables of the context. There are good semantic reasons for this feature, which we discuss in § III-C. For readability and intuition's sake, the rule is written using the syntactic sugar

$$t^{-1}(T_i) \triangleq_{\text{cast}(T)\text{in}_i(x)} T_t \quad \text{where } x : T_i \quad (1)$$

For the subtyping rules Fig. 4b, we use the syntactic sugar

$$S_i \cap S'_j \triangleq_{\text{cast}(S)\text{in}_i(x_i)} S_{\text{cast}(S)\text{in}_j(x'_j)} \quad \text{where } x_i : S_i, x'_j : S'_j$$

Our typed lambda calculus does not feature recursion for two reasons. First, it is not necessary: as any computable probability measure can be obtained as a computable pushforward of the uniform measure on the unit interval [18], [17], any sampler language which features the sampler operation map can, given a uniform sampler, target any computable probability measure. In particular, many rejection samplers, which are commonly implemented recursively, can alternatively be implemented using the operation reweight , as shown in Listing 3. Second, the categorical semantics of a typed, probabilistic, higher-order lambda calculus with recursion are a very recent area of investigation [38]; we consider the inclusion of recursive samplers to be further work.

B. Operational semantics

In practice, in order to evaluate a program containing a sampler, one must specify a finite number of samples $N \in \mathbb{N}$ which are to be produced. Our (big-step) operational semantics correspondingly takes the form of a reduction relation $(t, N) \rightarrow v$, where the left side consists of a well-typed closed term $t \in \text{Expr}$ and a number of samples $N \in \mathbb{N}$, and the right side is a *value* $v \in \text{Value}$, i.e. a term generated by the grammar

$$v ::= x \in \text{Var} \mid g \in G \mid (v, v) \mid \text{in}_i(v) \mid \lambda x: T. v \quad (2)$$

The rules of this big-step operational semantics, shown in full in [10, Table VI], are the usual rules for the standard language constructs, together with additional rules for our implemented sampler operations; these are given in Table II. For notational simplicity, these operations make use of lists (a, b, c, d) , which are in fact interpreted within our language as nested pairs $(a, (b, (c, d)))$. In order to keep the rules readable, we also introduce the shorthand $(t, N) \rightarrow ((v_1, w_1), \dots, (v_N, w_N))$ to denote the N reductions

$$\begin{aligned} (\text{hd}(t), \text{wt}(t)) &\rightarrow (v_1, w_1), \\ (\text{hd}(\text{tl}(t)), \text{wt}(\text{tl}(t))) &\rightarrow (v_2, w_2), \dots, \\ (\text{hd}(\text{tl}^{N-1}(t)), \text{wt}(\text{tl}^{N-1}(t))) &\rightarrow (v_N, w_N). \end{aligned}$$

| | | | |
|---|--|--|---|
| $\frac{}{\Gamma \vdash g : \mathbb{G}} \quad g \in \llbracket \mathbb{G} \rrbracket$ | $\frac{}{\Gamma, x : \mathbb{T}, \Delta \vdash x : \mathbb{T}}$ | $\frac{\Gamma \vdash t : \mathbb{T}}{\Gamma \vdash f(t) : \mathbb{G}} \quad \text{Func} \ni f : \mathbb{T} \rightarrow \mathbb{G}$ | $\frac{\Delta \vdash t : \mathbb{S}}{\Gamma \vdash \text{cast}(\mathbb{T})t : \mathbb{T}} \quad \mathbb{S} \triangleleft \mathbb{T}, \Gamma \triangleleft \Delta$ |
| $\frac{\Gamma \vdash t : \mathbb{T},}{(x_1, \dots, x_n) : \sum_{i \in m} t^{-1}(\mathbb{T}_i) \vdash t : \sum_{i \in m} \mathbb{T}_i} \quad \sum_{i \in m} \mathbb{T}_i \triangleleft \mathbb{T}, \Gamma = x_1 : \mathbb{S}_1, \dots, x_n : \mathbb{S}_n$ | | | |
| $\frac{\Gamma \vdash s : \mathbb{S} \quad \Gamma \vdash t : \mathbb{T}}{\Gamma \vdash (s, t) : \mathbb{S} \times \mathbb{T}}$ | $\frac{\Gamma \vdash t : \mathbb{S} \times \mathbb{T}}{\Gamma \vdash \text{fst}(t) : \mathbb{S}}$ | $\frac{\Gamma \vdash t : \mathbb{S} \times \mathbb{T}}{\Gamma \vdash \text{snd}(t) : \mathbb{T}}$ | $\frac{\Gamma, x : \mathbb{S} \vdash t : \mathbb{T} \quad \Gamma \vdash s : \mathbb{S}}{\Gamma \vdash \text{let } x = s \text{ in } t : \mathbb{T}}$ |
| $\frac{\Gamma, x : \mathbb{S} \vdash t : \mathbb{T}}{\Gamma \vdash \lambda x : \mathbb{S}. t : \mathbb{S} \rightarrow \mathbb{T}}$ | $\frac{\Gamma \vdash s : \mathbb{S} \quad \Gamma \vdash t : \mathbb{S} \rightarrow \mathbb{T}}{\Gamma \vdash t(s) : \mathbb{T}}$ | $\frac{\Gamma \vdash t : \mathbb{T}_j}{\Gamma \vdash \text{inj}_j(t) : \sum_{i \in n} \mathbb{T}_i} \quad j \in n$ | $\frac{\Gamma \vdash t : \sum_{i \in I} \mathbb{T}_i \quad \Gamma, x_i : \mathbb{T}_i \vdash s_i : \mathbb{T}}{\Gamma \vdash \text{case } t \text{ of } \{(i, x_i) \Rightarrow s_i\}_{i \in I} : \mathbb{T}}$ |
| $\frac{\Gamma \vdash t : \Sigma \mathbb{T}}{\Gamma \vdash \text{hd}(t) : \mathbb{T}}$ | $\frac{\Gamma \vdash t : \Sigma \mathbb{T}}{\Gamma \vdash \text{wt}(t) : \mathbb{R}^+}$ | $\frac{\Gamma \vdash t : \Sigma \mathbb{T}}{\Gamma \vdash \text{tl}(t) : \Sigma \mathbb{T}}$ | $\frac{\Gamma \vdash s : \Sigma \mathbb{S} \quad \Gamma \vdash t : \Sigma \mathbb{T}}{\Gamma \vdash s \otimes t : \Sigma(\mathbb{S} \times \mathbb{T})}$ |
| $\frac{\Gamma \vdash s : \mathbb{T} \rightarrow \mathbb{T} \quad \Gamma \vdash t : \mathbb{T}}{\Gamma \vdash \text{prng}(s, t) : \Sigma \mathbb{T}}$ | $\frac{\Gamma \vdash t : \Sigma \mathbb{T} \quad \Gamma \vdash n : \mathbb{N}}{\Gamma \vdash \text{thin}(t, n) : \Sigma \mathbb{T}}$ | $\frac{\Gamma \vdash s : \Sigma \mathbb{S} \quad \Gamma \vdash t : \mathbb{S} \rightarrow \mathbb{T}}{\Gamma \vdash \text{map}(t, s) : \Sigma \mathbb{T}}$ | $\frac{\Gamma \vdash s : \mathbb{T} \rightarrow \mathbb{R}^+ \quad \Gamma \vdash t : \Sigma \mathbb{T}}{\Gamma \vdash \text{reweight}(s, t) : \Sigma \mathbb{T}}$ |

TABLE I: Typing rules

Note that the product of two weighted samplers has as its weights the product of its factors' weights. The product and the operation `reweight` are the only operations modifying the weights of samplers.

The following proposition shows that the operational semantics is well-formed in that for any $N \in \mathbb{N}$, samplers can only reduce to weighted lists of length N .

Proposition III.1. [10, Appendix A] *If $\vdash s : \Sigma \mathbb{S}$ is a closed sampler, then for any $N \in \mathbb{N}$, if $(s, N) \rightarrow v$, then v has the form $((v_1, w_1), \dots, (v_N, w_N))$, where v_n are values and $w_n \in \mathbb{R}_{\geq 0}$ are weights. If \mathbb{S} is not a sampler type, then $v_n : \mathbb{S}$; more generally, each v_n might be a weighted list itself.*

The self-product operation: Having clarified the meaning of the product and of the `thin` operation, we are now in a position to formally justify the operation which we referred to, in § II, as the ‘self-product’ of a sampler. To motivate it, consider a sampler $t : \Sigma \mathbb{T}$ which evaluates as $(t, 2N) \rightarrow (x_1, \dots, x_{2N})$, where for notational clarity we have omitted the weights. From the above operational semantics, the lagged sampler `thin`(2, $t \otimes \text{tl}(t)$) : $\Sigma(\mathbb{T} \times \mathbb{T})$ evaluates to

$$(\text{thin}(2, t \otimes \text{tl}(t)), N) \rightarrow ((x_1, x_2), (x_3, x_4), \dots, (x_{2N-1}, x_{2N})).$$

This is the ‘self-product’ which was denoted t^2 in § II. This notion is important because it is the construction which allows us to generate pairs of independent samples from a given sampler. Note that simply taking $t \otimes t$ will produce pairs of perfectly correlated samples: the operational semantics gives $(t \otimes t, N) \rightarrow ((x_1, x_1), \dots, (x_N, x_N))$. More generally, for any $K \in \mathbb{N}$, we define the K -fold self-product of a sampler as

$$t^K \triangleq \text{thin}(K, t \otimes \text{tl}(t) \otimes \dots \otimes \text{tl}^{K-1}(t)). \quad (3)$$

Sampling from t^K is intended to allow the sampling of K -tuples of independent deviates generated by the sampler K . Ultimately, it is only to define this self-product operation that the sampler operation `thin` is included at all, it being somewhat of an unnatural construct.

C. Denotational semantics

1) *Denotational universe:* We will see in § V that continuous maps play a special role in the verification of sampler

properties. We therefore need a denotational domain in which continuity is a meaningful concept. We also need a Cartesian closed model, as we want to interpret the lambda-abstraction operation of our calculus. A standard solution is to consider the category of *compactly generated topological spaces* [35], [31], [25] (henceforth *CG-spaces*). A topological space X is compactly generated if it is Hausdorff and has the property that $C \subseteq X$ is closed iff $C \cap K$ is closed in K for every compact K in X [35, §1]. We need not worry about the theory of these spaces, but the following facts are essential in what follows.

Proposition III.2 ([35], [25]). 1) *The category \mathbf{CG} of CG-spaces and continuous functions is Cartesian closed.*
2) *The category \mathbf{CG} is complete and cocomplete.*
3) *Every metrizable topological space is CG.*
4) *Locally closed subsets (i.e. intersections of an open and a closed subset) of CG-spaces are compactly generated.*

It is worth briefly describing the Cartesian closed structure of \mathbf{CG} . The product is in general different from the product in \mathbf{Top} , the category of topological spaces: if the usual product topology is not already compactly generated, then it needs to be modified to enforce compact generation [35, §4]. However, in most practical instances the usual product topology is already compactly generated – for example, any countable product of metrizable spaces is metrizable, and thus compactly generated by Prop. III.2. The internal hom $[X, Y]$ between CG-spaces X, Y is given by the set of continuous maps $X \rightarrow Y$ together with the topology of uniform convergence on compact sets, also known as the compact-open topology [35, §5].

2) *Semantics of types:* With this categorical model in place we define the semantics of types. The semantics of ground types is as expected: $\llbracket \mathbb{N} \rrbracket = \mathbb{N}$, equipped with the discrete topology, and $\llbracket \mathbb{R} \rrbracket = \mathbb{R}$, $\llbracket \mathbb{R}^+ \rrbracket = [0, \infty)$ with the usual topology. The spaces $f^{-1}(i)$, $f \in \{\leq, <, \geq, >, =, \neq\}$, $i \in 2$ are interpreted precisely as the notation suggests, e.g.

$$\begin{aligned} \llbracket <^{-1}(0) \rrbracket &= \{(x, y) \mid x, y \in \mathbb{R} \wedge x \geq y\}, \\ \llbracket =^{-1}(1) \rrbracket &= \{(x, x) \mid x \in \mathbb{R}\} \end{aligned}$$

together with the subspace topology inherited from $\mathbb{R} \times \mathbb{R}$. Since all these spaces are metrizable, our ground types are

$$\begin{array}{c}
\frac{((s(\text{hd}(t)), \text{wt}(t)), N) \rightarrow (v_1, w_1) \quad \dots \quad ((s(\text{hd}(\text{tl}^{N-1}(t)), \text{wt}(\text{tl}^{N-1}(t))), N) \rightarrow (v_N, w_N))}{(\text{map}(s, t), N) \rightarrow ((v_1, w_1), \dots, (v_N, w_N))} \\
\frac{((\text{hd}(t), s(\text{hd}(t)) \cdot \text{wt}(t)), N) \rightarrow (v_1, w_1) \quad \dots \quad ((\text{hd}(\text{tl}^{N-1}(t)), s(\text{hd}(\text{tl}^{N-1}(t))) \cdot \text{wt}(\text{tl}^{N-1}(t))), N) \rightarrow (v_N, w_N)}{(\text{reweight}(s, t), N) \rightarrow ((v_1, w_1), \dots, (v_N, w_N))} \\
\frac{(s, N) \rightarrow ((v_1, w_1), \dots, (v_N, w_N)) \quad (t, N) \rightarrow ((v'_1, w'_1), \dots, (v'_N, w'_N))}{(s \otimes t, N) \rightarrow (((v_1, v'_1), w_1 \cdot w'_1), \dots, ((v_N, v'_N), w_N \cdot w'_N))} \\
\frac{(t, N) \rightarrow ((v_1, w_1), \dots, (v_N, w_N))}{(\text{hd}(t), N) \rightarrow v_1} \quad \frac{(t, N) \rightarrow ((v_1, w_1), \dots, (v_N, w_N))}{(\text{tl}(t), N-1) \rightarrow ((v_2, w_2), \dots, (v_N, w_N))} \quad \frac{(t, N) \rightarrow ((v_1, w_1), \dots, (v_N, w_N))}{(\text{wt}(t), N) \rightarrow w_1} \\
\frac{(s, N) \rightarrow i \quad (t, Ni) \rightarrow ((v_1, w_1), \dots, (v_{Ni}, w_{Ni}))}{(\text{thin}(s, t), N) \rightarrow ((v_1, w_1), (v_{i+1}, w_{i+1}), (v_{2i+1}, w_{2i+1}), \dots, (v_{(N-1)i+1}, w_{(N-1)i+1}))} \\
\frac{(t, N) \rightarrow v_1 \quad (s(t), N) \rightarrow v_2 \quad \dots \quad (s^{N-1}(t), N) \rightarrow v_N}{(\text{prng}(s, t), N) \rightarrow ((v_1, 1), \dots, (v_N, 1))}
\end{array}$$

TABLE II: Big-step operational semantics of sampler operations

interpreted in **CG** by Prop. III.2.

Products (including the unit type) and function types are interpreted in the obvious way using the Cartesian closed structure of **CG**. Coproduct types are interpreted by coproducts in **CG**, and given two terms $s, t : \mathbb{T}$ interpreted as **CG**-morphisms $\llbracket s \rrbracket : A \rightarrow \llbracket \mathbb{T} \rrbracket, \llbracket t \rrbracket : B \rightarrow \llbracket \mathbb{T} \rrbracket$, the pullback type ${}_s\mathbb{T}_t$ is interpreted as the pullback $A \times_{\llbracket \mathbb{T} \rrbracket} B$ of $\llbracket s \rrbracket$ along $\llbracket t \rrbracket$. All these spaces live in **CG** by Prop. III.2.

Since sampler types are coinductive types, their semantics will hinge on the existence of terminal coalgebras.

Theorem III.1 (Adámek). *Let \mathcal{C} be a category with terminal object 1, and $F : \mathcal{C} \rightarrow \mathcal{C}$ be a functor. If \mathcal{C} has and F preserves ω^{op} -indexed limits, then the limit νF of $1 \xleftarrow{!} F1 \xleftarrow{F!} FF1 \xleftarrow{FF!} \dots$ is the terminal coalgebra of F .*

Since **CG** is complete, it has ω^{op} -indexed limits. Recall that we want to interpret $\Sigma\mathbb{T}$ as the coinductive type defined by the ‘functor’ $\mathbb{T} \times \mathbb{R}^+ \times -$. Formally, given a type \mathbb{T} we want

$$\llbracket \Sigma\mathbb{T} \rrbracket \triangleq \nu(\llbracket \mathbb{T} \rrbracket \times \mathbb{R}^+ \times \text{Id}). \quad (4)$$

Since products are limits, and limits commute with limits, it is clear that the functor $\llbracket \mathbb{T} \rrbracket \times \mathbb{R}^+ \times \text{Id}$ preserves limits, and in particular ω^{op} -indexed ones. Adámek’s theorem thus guarantees the existence of an object satisfying (4). More concretely, since the terminal object 1 is trivially metrizable, and since \mathbb{R}^+ is metrizable, each object in the terminal sequence will be metrizable provided $\llbracket \mathbb{T} \rrbracket$ is, and thus $\prod_n (\llbracket \mathbb{T} \rrbracket \times \mathbb{R}^+)^n$ will be metrizable whenever $\llbracket \mathbb{T} \rrbracket$ is, and will therefore be equipped with the usual product topology. The limit defining (4) is a closed subspace of this product, which means that the limit in **CG** defining $\llbracket \Sigma\mathbb{T} \rrbracket$ is the same as in **Top** when $\llbracket \mathbb{T} \rrbracket$ is metrizable (for example, if \mathbb{T} is a ground type or a product of ground types). However, by defining $\llbracket \Sigma\mathbb{T} \rrbracket$ coinductively rather than simply as $(\llbracket \mathbb{T} \rrbracket \times \mathbb{R}^+)^{\omega}$, we obtain a terminal coalgebra structure on $\llbracket \Sigma\mathbb{T} \rrbracket$, and therefore the ability to define sampler operations coinductively.

3) *Semantics of the subtyping relation:* Our language contains the predicates $f \in \{\leq, <, \geq, >, =, \neq\}$ (essential for rejection sampling § II-C) and yet is meant to be interpreted in

a universe of topological spaces and continuous maps. These predicates are not continuous maps $\mathbb{R} \times \mathbb{R} \rightarrow 2$ for the usual topology on $\mathbb{R} \times \mathbb{R}$. However, for each such predicate f , the sets $\llbracket f^{-1}(0) \rrbracket$ and $\llbracket f^{-1}(1) \rrbracket$ are *locally closed sets*, that is to say the intersection of an open set and a closed set (for the usual topology on $\mathbb{R} \times \mathbb{R}$), and therefore **CG**-spaces by Prop. III.2, e.g. $\llbracket <^{-1}(0) \rrbracket$ is closed and $\llbracket <^{-1}(1) \rrbracket$ open.

Our central idea for dealing with discontinuities is that since **CG** is cocomplete, the space $\llbracket f^{-1}(0) \rrbracket + \llbracket f^{-1}(1) \rrbracket$ is a **CG**-space. This space has the nice property that f is *continuous* as a map $f : \llbracket f^{-1}(0) + f^{-1}(1) \rrbracket \rightarrow 2$. Since each $f^{-1}(i)$ is a type, we can enforce this semantics by simply *typing* these built-in functions in **Func** as $f : f^{-1}(0) + f^{-1}(1) \rightarrow \mathbb{B}$.

The topology on $\llbracket f^{-1}(0) + f^{-1}(1) \rrbracket$ is finer than the usual topology on $\mathbb{R} \times \mathbb{R}$, which means that the identity map $\text{Id} : \llbracket f^{-1}(0) + f^{-1}(1) \rrbracket \rightarrow \mathbb{R} \times \mathbb{R}$ is continuous. This is the semantic basis for the axiom in Fig. 4b. From the other rules it is easy to see by induction that the subtyping relation is always between spaces *sharing the same carrier set* and is semantically given by coarsening the topology. In other words, if $\mathbb{S} \triangleleft \mathbb{T}$, then $\llbracket \mathbb{S} \rrbracket$ and $\llbracket \mathbb{T} \rrbracket$ share the same carrier and the corresponding identity map $\text{Id} : \llbracket \mathbb{S} \rrbracket \rightarrow \llbracket \mathbb{T} \rrbracket$ is continuous.

Example III.1. *Let $\text{p} \triangleq \text{if } x = 0 \text{ then } 1 \text{ else } -1$; we will first show how the context-restriction rule allows us to type-check this program. For readability’s sake, let $\text{Eq} \triangleq =^{-1}(1)$ and $\text{Neq} \triangleq =^{-1}(0)$. We now derive, using $= : \text{Neq} + \text{Eq} \rightarrow \mathbb{R}$,*

$$\begin{array}{c}
\frac{x : \mathbb{R} \vdash x : \mathbb{R} \quad \vdash 0 : \mathbb{R}}{x : \mathbb{R} \vdash (x, 0) : \mathbb{R} \times \mathbb{R}} \\
\frac{x : (x, 0)^{-1}\text{Neq} + (x, 0)^{-1}\text{Eq} \vdash (x, 0) : \text{Neq} + \text{Eq} \quad \text{Neq} + \text{Eq} \triangleleft \mathbb{R} \times \mathbb{R}}{x : (x, 0)^{-1}\text{Neq} + (x, 0)^{-1}\text{Eq} \vdash x = 0 : \mathbb{B} \quad \vdash 1 : \mathbb{R} \quad \vdash -1 : \mathbb{R}} \\
\frac{x : (x, 0)^{-1}\text{Neq} + (x, 0)^{-1}\text{Eq} \vdash \text{if } x = 0 \text{ then } 1 \text{ else } -1 : \mathbb{R}}{}
\end{array}$$

Anticipating the semantics on terms discussed shortly, it can easily be shown that

$$\llbracket (x, 0)^{-1}\text{Neq} + (x, 0)^{-1}\text{Eq} \rrbracket = ((-\infty, 0) \cup (0, \infty)) + \{0\}$$

and thus $\llbracket \mathbf{p} \rrbracket$ is the continuous map

$$\llbracket \mathbf{p} \rrbracket : ((-\infty, 0) \cup (0, \infty)) + \{0\} \rightarrow \mathbb{R}, x \mapsto \begin{cases} 1 & \text{if } x = 0 \\ -1 & \text{else} \end{cases}$$

4) *Semantics of well-formed terms:* Axioms, weakening, subtyping, product, projections, let -binding, λ -abstraction, function application, injections and pattern matching are interpreted in the expected way (given that \mathbf{CG} is a Cartesian closed category with coproducts).

Continuous built-in functions, for example $+$: $\mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ or exp : $\mathbb{R} \rightarrow \mathbb{R}$, are interpreted in the obvious way. As explained above, discontinuous built-in functions $\{\leq, <, \geq, >, =, \neq\}$ are typed in such a way that their natural interpretations are tautologically continuous.

We can now describe the semantics of the context-restriction rule. From the premise, our observations in § III-C3, and the side-conditions, we have morphisms

$$\llbracket t \rrbracket : \prod_{j \in n} \llbracket \mathbf{S}_j \rrbracket \rightarrow \llbracket \mathbf{T} \rrbracket, \quad \text{and} \quad \text{Id} : \prod_{i \in m} \llbracket \mathbf{T}_i \rrbracket \rightarrow \llbracket \mathbf{T} \rrbracket.$$

By Eq. (1) we interpret each ‘inverse image type’ $t^{-1}(\mathbf{T}_i)$ as the pullback (inverse image) of $\llbracket t \rrbracket$ along the inclusion $\llbracket \mathbf{T} \rrbracket_i \hookrightarrow \prod_{i \in m} \llbracket \mathbf{T}_i \rrbracket$ which is, as the notation implies, simply given by $\llbracket t \rrbracket^{-1}(\prod_{i \in m} \llbracket \mathbf{T}_i \rrbracket)$. Since $\prod_{i \in m} \llbracket \mathbf{T}_i \rrbracket$ and $\llbracket \mathbf{T} \rrbracket$ share the same carrier, it is clear that this defines a partition of $\llbracket \mathbf{T} \rrbracket$, and we can thus retype t as a continuous map $\prod_{i \in m} \llbracket t^{-1}(\mathbf{T}_i) \rrbracket \rightarrow \prod_{i \in m} \llbracket \mathbf{T}_i \rrbracket$, interpreting the rule.

As mentioned earlier in this section, context-restriction prevents λ -abstraction; the following example illustrates why this must be the case.

Example III.2. Consider the program $x < y$ derived by:

$$\frac{x : \mathbb{R}, y : \mathbb{R} \vdash (x, y) : \mathbb{R} \times \mathbb{R} \quad (x, y) : (x, y)^{-1}(<^{-1}(0)) + (x, y)^{-1}(<^{-1}(1)) \vdash (x, y) : <^{-1}(0) + <^{-1}(1)}{(x, y) : (x, y)^{-1}(<^{-1}(0)) + (x, y)^{-1}(<^{-1}(1)) \vdash x < y : \mathbb{B}}$$

The interpretation of $x < y$ is given by the continuous function

$$\llbracket < \rrbracket : \{(x, y) \mid x < y\} + \{(x, y) \mid x \geq y\} \rightarrow 2.$$

Although it has the same carrier $\mathbb{R} \times \mathbb{R}$, the domain of this map is no longer of product of topological spaces; it is now a coproduct of topological spaces. This means that it is no longer possible to λ -abstract over one of the variables of this function using the Cartesian closed structure of \mathbf{CG} .

In order to be able to λ -abstract the map $<$, we would need a topology on $\mathbb{R} \times \mathbb{R}$ with the property that for any given $x_0 \in \mathbb{R}$ the function $x_0 < - : \mathbb{R} \rightarrow 2$ is continuous. This would introduce the open sets $[x_0, \infty)$ to the topology of \mathbb{R} for each $x_0 \in \mathbb{R}$, meaning that we must equip \mathbb{R} with the notoriously problematic lower limit topology (a.k.a. the Sorgenfrey line). Whether or not this is a CG-space seems to be a thorny question, possibly independent of ZF [20].

Finally, we define the denotational semantics of sampler operations using the coinductive nature of sampler types.

Recall that for a type \mathbf{T} , $\llbracket \Sigma \mathbf{T} \rrbracket \triangleq \nu(\llbracket \mathbf{T} \rrbracket \times \mathbb{R}^+ \times \text{Id})$. In particular, $\llbracket \Sigma \mathbf{T} \rrbracket$ comes equipped with a coalgebra structure map

$$\text{unfold}_{\mathbf{T}} : \llbracket \Sigma \mathbf{T} \rrbracket \rightarrow \llbracket \mathbf{T} \rrbracket \times \mathbb{R}^+ \times \llbracket \Sigma \mathbf{T} \rrbracket.$$

Moreover, for any other (continuous) coalgebra structure map $\gamma : X \rightarrow \llbracket \mathbf{T} \rrbracket \times \mathbb{R}^+ \times X$, the terminal nature of $\llbracket \Sigma \mathbf{T} \rrbracket$ provides a unique $\llbracket \mathbf{T} \rrbracket \times \mathbb{R}^+ \times \text{Id}$ -coalgebra morphism

$$\text{beh}(\gamma) : X \rightarrow \llbracket \Sigma \mathbf{T} \rrbracket.$$

Since $\llbracket \Sigma \mathbf{T} \rrbracket$ is interpreted in \mathbf{CG} , it follows automatically that both $\text{unfold}_{\mathbf{T}}$ and $\text{beh}(\gamma)$ are continuous. However, what is not immediately clear is that beh is in fact continuous in γ .

Proposition III.3. [10, Appendix A] Let $F : \mathbf{CG} \rightarrow \mathbf{CG}$ satisfy the condition of Thm. III.1 as well as the condition that $\text{int}(\nu F) \neq \emptyset$ in $\prod_i F^i 1$, and let $\text{beh}_X : [X, FX] \rightarrow [X, \nu F]$ be the (behaviour) map associating to any F -coalgebra structure on X the unique coalgebra morphism into the terminal coalgebra. The map beh_X is continuous, i.e. is a \mathbf{CG} -morphism.

Using unfold and beh we define the denotational semantics of all the sampler operations in Table III. These definitions are precisely the infinite (coinductive) versions of the finitary transformations defined in the operational semantics of Table II. All the maps involved in these definitions are continuous; this follows from Prop. III.3 and the fact that evaluation and function composition are continuous operations on the internal hom sets of \mathbf{CG} ([35, 5.2,5.9]).

D. Adequacy

This language features an interesting asymmetry in that its denotational semantics is written in terms of the coinductive sampler type $\llbracket \Sigma \mathbf{T} \rrbracket$, while its operational semantics is written in terms of finitary operations on finite sequences of samples. Moreover, the operational semantics is given in terms of reductions to *values*, i.e. terms whose types are constructed without the type constructor Σ , whereas the denotational semantics does not make this distinction. To establish a connection, we start by defining a generic way to convert terms of arbitrary types into values, following the idea behind the operational semantics. Given a type \mathbf{T} and an integer N we inductively define its associated value type $\text{val}^N(\mathbf{T}) \in \text{Value}$ by:

$$\begin{aligned} \text{val}^N(\mathbf{G}) &= \mathbf{G} & \text{val}^N(\Sigma \mathbf{T}) &= (\text{val}^N \mathbf{T})^N \\ \text{val}^N(\mathbf{S} * \mathbf{T}) &= \text{val}^N(\mathbf{S}) * \text{val}^N(\mathbf{T}), & * &\in \{\times, +, \rightarrow\} \end{aligned}$$

where $\mathbf{G} \in \text{Ground}$.² We now define the *generalized projection* maps $p_{\mathbf{T}}^N : \llbracket \mathbf{T} \rrbracket \rightarrow \llbracket \text{val}^N(\mathbf{T}) \rrbracket$ recursively via

$$\begin{aligned} p_{\mathbf{G}}^N &= \text{id}_{\llbracket \mathbf{G} \rrbracket}, & p_{\mathbf{S} * \mathbf{T}}^N &= p_{\mathbf{S}}^N * p_{\mathbf{T}}^N, * \in \{\times, +\} \\ p_{\Sigma \mathbf{T}}^N &= \text{id}_{\llbracket \Sigma \mathbf{T} \rrbracket} & p_{\mathbf{T} \times \mathbf{R}^+}^N &= \pi_{1:N} \circ (p_{\mathbf{T} \times \mathbf{R}^+}^N)^\omega \end{aligned}$$

The reader will have noticed that we have defined $p_{\mathbf{S} \rightarrow \mathbf{T}}^N$ trivially. The reason is that, as a quick examination of the rules of Table II will reveal, there is no conclusion and no premise

²Since we’re only interested in closed samplers here, and since pullback types can only occur in a context, we need not define val^N on pullback types.

$$\begin{array}{c}
\frac{[\Gamma \vdash t : \Sigma \mathbf{T}] = f}{[\Gamma \vdash \text{hd}(t) : \mathbf{T}] = \pi_1 \circ \text{unfold}_{\mathbf{T}} \circ f} \quad \frac{[\Gamma \vdash t : \Sigma \mathbf{T}] = f}{[\Gamma \vdash \text{wt}(t) : \mathbf{R}^+] = \pi_2 \circ \text{unfold}_{\mathbf{T}} \circ f} \quad \frac{[\Gamma \vdash t : \Sigma \mathbf{T}] = f}{[\Gamma \vdash \text{tl}(t) : \Sigma \mathbf{T}] = \pi_3 \circ \text{unfold}_{\mathbf{T}} \circ f} \\
\frac{[\Gamma \vdash s : \mathbf{N}] = f \quad [\Gamma \vdash t : \Sigma \mathbf{T}] = g}{[\Gamma \vdash \text{thin}(s, t) : \Sigma(\mathbf{T})] = \text{ev}_{\Sigma \mathbf{T}, \Sigma \mathbf{T}} \circ (\text{id}_{\Sigma \mathbf{T}} \times \text{beh}_{\Sigma \mathbf{T}}) \circ (\text{id}_{\Sigma \mathbf{T}} \times (\text{unfold}_{\mathbf{T}} \circ (\pi_3 \circ \text{unfold}_{\mathbf{T}})^{(-1)})) \circ \langle f, g \rangle} \\
\frac{[\Gamma \vdash s : \Sigma \mathbf{S}] = f \quad [\Gamma \vdash t : \Sigma \mathbf{T}] = g}{[\Gamma \vdash s \otimes t : \Sigma(\mathbf{S} \times \mathbf{T})] = \text{beh}_{\Sigma \mathbf{S}, \Sigma \mathbf{T}} (\pi_1 \times \pi_4 \times (\pi_2 \cdot \pi_5) \times \pi_3 \times \pi_6 \circ \text{unfold}_{\mathbf{S}} \times \text{unfold}_{\mathbf{T}}) \circ \langle f, g \rangle} \\
\frac{[\Gamma \vdash s : \Sigma \mathbf{S}] = f \quad [\Gamma \vdash t : \mathbf{S} \rightarrow \mathbf{T}] = g}{[\Gamma \vdash \text{map}(t, s) : \Sigma \mathbf{T}] = \text{ev}_{\Sigma \mathbf{S}, \Sigma \mathbf{T}} \circ (\text{id}_{\Sigma \mathbf{S}} \times \text{beh}_{\Sigma \mathbf{S}}) \circ (\text{id}_{\Sigma \mathbf{S}} \times ((-\times \text{id}_{\mathbf{R}^+} \times \text{id}_{\Sigma \mathbf{S}}) \circ \text{unfold}_{\mathbf{S}})) \circ \langle f, g \rangle} \\
\frac{[\Gamma \vdash s : \Sigma \mathbf{T}] = f \quad [\Gamma \vdash t : \mathbf{T} \rightarrow \mathbf{R}^+] = g}{[\Gamma \vdash \text{reweight}(t, s)] = \text{ev}_{\Sigma \mathbf{T}, \Sigma \mathbf{T}} \circ (\text{id}_{\Sigma \mathbf{T}} \times \text{beh}_{\Sigma \mathbf{T}}) \circ (\text{id}_{\Sigma \mathbf{T}} \times ((\text{id}_{\mathbf{T}} \times - \times \text{id}_{\Sigma \mathbf{T}}) \circ \text{unfold}_{\mathbf{T}})) \circ \langle f, g \rangle} \\
\frac{[\Gamma \vdash t : \mathbf{T}] = f \quad [\Gamma \vdash s : \mathbf{T} \rightarrow \mathbf{T}] = g}{[\Gamma \vdash \text{prng}(s, t) : \Sigma \mathbf{T}] = \text{ev}_{\mathbf{T}, \Sigma \mathbf{T}} \circ (\text{id}_{\mathbf{T}} \times \text{beh}_{\mathbf{T}}) \circ (\text{id}_{\mathbf{T}} \times (\text{id}_{\mathbf{T}} \times 1 \times -)) \circ \langle f, g \rangle}
\end{array}$$

TABLE III: Denotational semantics of sampler operations

of the type $(t, N) \rightarrow v$ where t is of function type. The only occurrence of terms of function types are within an evaluation, or are *values*, i.e. terms trivially reducing to themselves.

Theorem III.2. [10, Appendix A] For any program $\vdash t : \mathbf{T}$, we have

$$(t, N) \rightarrow v \Leftrightarrow p_{\mathbf{T}}^N([\![t]\!]]) = [\![v]\!].$$

IV. EQUIVALENCE OF SAMPLERS

In order to implement a system for reasoning about whether a deterministic sampler targets a particular probability distribution, it is necessary to first define a notion of equivalence between samplers. Having such a system gives a natural path towards verifying a sampler: first rewrite a given sampler s in an equivalent but simpler form, and then show that this simplified form targets the correct distribution. This is the approach taken in the derivations in § II, which implicitly used several equivalence results – in particular, `let`-reduction and the equivalence of the nested self-product $(s^m)^n$ to the self-product $s^{m \cdot n}$ for any sampler s . In this section, we introduce a relation \approx on programs which justifies this type of reasoning.

Definition IV.1. We say that two programs $\Gamma \vdash s : \mathbf{T}$ and $\Gamma \vdash t : \mathbf{T}$ are equivalent, notation $\Gamma \vdash s \approx t : \mathbf{T}$, if they are related by the smallest congruence relation on well-typed terms containing the rules of Table IV.³

The rules of Table IV employ a number of shorthand conventions for a more concise presentation. We introduce identity functions $\text{id}_{\mathbf{S}} \triangleq \lambda x : \mathbf{S}. x : \mathbf{S} \rightarrow \mathbf{S}$, constant functions $1_{\mathbf{S}} \triangleq \lambda x : \mathbf{S}. 1 : \mathbf{S} \rightarrow \mathbf{R}^+$, function composition $t \circ s \triangleq \lambda x : \mathbf{S}. t(s(x)) : \mathbf{S} \rightarrow \mathbf{U}$ where $s : \mathbf{S} \rightarrow \mathbf{T}, t : \mathbf{T} \rightarrow \mathbf{U}$, compositions $f^0 \triangleq \text{id}_{\mathbf{S}} : \mathbf{S} \rightarrow \mathbf{S}, f^n \triangleq f \circ f^{n-1}$ for any $n \in \mathbb{N}$, pointwise products $s \cdot t \triangleq \lambda x : \mathbf{S}, y : \mathbf{T}. s(x) * t(y) : \mathbf{S} \times \mathbf{T} \rightarrow \mathbf{R}^+$ of real-valued functions $s : \mathbf{S} \rightarrow \mathbf{R}^+, t : \mathbf{T} \rightarrow \mathbf{R}^+$, and finally Cartesian products $s \times t \triangleq \lambda x : \mathbf{S}, y : \mathbf{T}. (s(x), t(y)) : \mathbf{S} \times \mathbf{T} \rightarrow \mathbf{S}' \times \mathbf{T}'$ of functions $s : \mathbf{S} \rightarrow \mathbf{S}', t : \mathbf{T} \rightarrow \mathbf{T}'$.

³By *congruence relation*, we mean that \approx is an equivalence relation preserved by all operations in the language. For example, if $\Gamma \vdash s \approx t : \Sigma \mathbf{T}$ holds, then $\Gamma \vdash \text{tl}(s) \approx \text{tl}(t) : \Sigma \mathbf{T}$ must hold as well, and the same for all operations in the language.

Theorem IV.1. [10, Appendix A] The rules of Table IV are sound: if $\Gamma \vdash s \approx t : \mathbf{T}$, then $[\![\Gamma \vdash s : \mathbf{T}]\!] = [\![\Gamma \vdash t : \mathbf{T}]\!]$.

The proof is a straightforward exercise in coinductive reasoning and can be found in the Appendix, along with the full list of equivalence rules. The soundness of these rules with respect to operational equivalence then follows from abstraction, though it is also straightforward to show directly.

Recall that an important application of our sampler operations is to provide a formal definition of the *self-product* of samplers, given in (3). It is crucial that our equivalence rules should show that this self-product is well-defined.

Proposition IV.1. [10, Appendix A] For any $\Gamma \vdash s : \Sigma \mathbf{S}$, $m, n \in \mathbb{N}$, the self-product satisfies $\Gamma \vdash (s^m)^n \approx s^{m \cdot n} : \Sigma(\mathbf{S}^{m \cdot n})$.

The equivalence rules in Table IV suggest a procedure for simplifying samplers. Consider samplers which have no occurrences of the operation `prng`. Of our remaining sampler operations, we identify two groups: $\{\text{tl}, \text{hd}, \text{wt}, \text{thin}, \otimes\}$ and $\{\text{map}, \text{reweight}\}$. Applying the rules of Table IV, we see that for each combination of operations in the first and second group, there is a rule which enables us to pull the first operation into the body of the second. Therefore, any sampler with no instances of `prng` can be written so that the sampler operations `tl`, `hd`, `wt`, `thin`, \otimes are pulled all the way inwards.

Proposition IV.2. Let $\Gamma \vdash t : \Sigma \mathbf{T}$ be a sampler which contains no instances of `prng`. Applying the rules of Table IV, it follows that we can equivalently rewrite such a sampler in either the form $\Gamma \vdash \text{map}(f, \text{reweight}(g, \text{map}(f', \dots, s))) : \Sigma \mathbf{T}$ or $\Gamma \vdash \text{reweight}(g, \text{map}(f, \text{reweight}(g', \dots, s))) : \Sigma \mathbf{T}$, i.e. a composition of invocations of `map` and `reweight` (including the trivial case of zero occurrences of either), where crucially the sampler s does not contain the sampler operations `map` or `reweight`.

We can also show that the self-product distributes over the operations `map` self- and `reweight`; such operations are useful for representing the self-product of a composite sampler in a

| | | |
|---|---|---|
| $\Gamma, s : \mathbf{S} \vdash (\lambda x : \mathbf{S}. t)(s) \approx t[x \leftarrow s] : \mathbf{T}$ | $\Gamma \vdash \lambda x : \mathbf{S}. t(x) \approx t : \mathbf{S} \rightarrow \mathbf{T}$ | |
| $\Gamma, s : \mathbf{S} \vdash \text{let } x = s \text{ in } t \approx (\lambda x : \mathbf{S}. t)(s) : \mathbf{T}$ | | |
| $\Gamma \vdash \text{if True then } s \text{ else } t \approx s : \mathbf{T}$ | $\Gamma \vdash \text{if False then } s \text{ else } t \approx t : \mathbf{T}$ | |
| $\Gamma \vdash \text{fst}((s, t)) \approx s : \mathbf{S}$ | $\Gamma \vdash \text{snd}((s, t)) \approx t : \mathbf{T}$ | |
| $\Gamma \vdash \text{hd}(\text{map}(s, t)) \approx s(\text{hd}(t)) : \mathbf{T}$ | $\Gamma \vdash \text{wt}(\text{map}(s, t)) \approx \text{wt}(t) : \mathbf{R}^+$ | $\Gamma \vdash \text{tl}(\text{map}(s, t)) \approx \text{map}(s, \text{tl}(t)) : \Sigma \mathbf{T}$ |
| $\Gamma \vdash (\text{hd}(s), \text{hd}(t)) \approx \text{hd}(s \otimes t) : \mathbf{S} \times \mathbf{T}$ | $\Gamma \vdash \text{wt}(s) * \text{wt}(t) \approx \text{wt}(s \otimes t) : \mathbf{R}^+$ | $\Gamma \vdash \text{tl}(s) \otimes \text{tl}(t) \approx \text{tl}(s \otimes t) : \Sigma(\mathbf{S} \times \mathbf{T})$ |
| $\Gamma \vdash \text{hd}(\text{thin}(n, t)) \approx \text{hd}(t) : \mathbf{T}$ | $\Gamma \vdash \text{wt}(\text{thin}(n, t)) \approx \text{wt}(t) : \mathbf{R}^+$ | $\{\Gamma \vdash \text{tl}(\text{thin}(n, t)) \approx \text{thin}(n, \text{tl}^n(t)) : \Sigma \mathbf{T} \mid n \in \mathbb{N}\}$ |
| | $\Gamma \vdash \text{thin}(1, t) \approx t : \Sigma \mathbf{T}$ | |
| $\Gamma \vdash \text{hd}(\text{prng}(s, t)) \approx t : \mathbf{T}$ | $\Gamma \vdash \text{wt}(\text{prng}(s, t)) \approx 1 : \mathbf{R}^+$ | $\Gamma \vdash \text{tl}(\text{prng}(s, t)) \approx \text{prng}(s, \text{tl}(t)) : \Sigma \mathbf{T}$ |
| $\Gamma \vdash \text{hd}(\text{reweight}(s, t)) \approx \text{hd}(t) : \mathbf{T}$ | $\Gamma \vdash \text{wt}(\text{reweight}(s, t)) \approx s(\text{hd}(t)) * \text{wt}(t) : \mathbf{R}^+$ | $\Gamma \vdash \text{tl}(\text{reweight}(s, t)) \approx \text{reweight}(s, \text{tl}(t)) : \Sigma \mathbf{T}$ |
| $\Gamma \vdash \text{thin}(n, \text{thin}(m, t)) \approx \text{thin}(n * m, t) \approx \Sigma \mathbf{T}$ | $\Gamma \vdash \text{map}(g, \text{map}(f, t)) \approx \text{map}(g \circ f, t) : \Sigma \mathbf{T}$ | |
| $\Gamma \vdash \text{reweight}(g, \text{reweight}(f, t)) \approx \text{reweight}(f \cdot g, t) : \Sigma \mathbf{T}$ | | |
| $\{\Gamma \vdash \text{thin}(n, \text{prng}(s, t)) \approx \text{prng}(s^n, t) : \Sigma \mathbf{T} \mid n \in \mathbb{N}\}$ | $\Gamma \vdash \text{thin}(n, \text{map}(s, t)) \approx \text{map}(s, \text{thin}(n, t)) : \Sigma \mathbf{T}$ | |
| $\Gamma \vdash s \otimes \text{map}(f, t) \approx \text{map}(\text{id}_{\mathbf{S}} \times f, s \otimes t) : \Sigma(\mathbf{S} \times \mathbf{T})$ | $\Gamma \vdash \text{map}(f, s) \otimes t \approx \text{map}(f \times \text{id}_{\mathbf{T}}, s \otimes t) : \Sigma(\mathbf{S} \times \mathbf{T})$ | |
| $\Gamma \vdash s \otimes \text{reweight}(g, t) \approx \text{reweight}(1_{\mathbf{S}} \cdot g, s \otimes t) : \Sigma(\mathbf{S} \times \mathbf{T})$ | $\Gamma \vdash \text{reweight}(f, s) \otimes t \approx \text{reweight}(f \cdot 1_{\mathbf{T}}, s \otimes t) : \Sigma(\mathbf{S} \times \mathbf{T})$ | |
| $\Gamma \vdash \text{prng}(f, a) \otimes \text{prng}(g, b) \approx \text{prng}(f \times g, (a, b)) : \Sigma(\mathbf{S} \times \mathbf{T})$ | $\Gamma \vdash \text{thin}(n, s) \otimes \text{thin}(n, t) \approx \text{thin}(n, s \otimes t) : \Sigma(\mathbf{S} \times \mathbf{T})$ | |

TABLE IV: Rules for sampler equivalence

simpler form whose correctness can then be verified.

Proposition IV.3. [10, Appendix A] *For any mapped sampler $\Gamma \vdash \text{map}(f, s) : \Sigma \mathbf{T}$ and any $n \in \mathbb{N}$, it follows that $\Gamma \vdash \text{map}(f, s)^n \approx \text{map}(f \times \dots \times f, s^n) : \Sigma(\mathbf{T}^n)$; for a reweighted sampler $\Gamma \vdash \text{reweight}(f, s) : \Sigma \mathbf{S}$, it follows that $\Gamma \vdash \text{reweight}(f, s)^n \approx \text{reweight}(f \cdot \dots \cdot f, s^n) : \Sigma(\mathbf{S}^n)$.*

V. SEMANTIC CORRECTNESS OF SAMPLERS

The fundamental correctness criterion for a sampler is that it should produce samples which are distributed according to the desired target distribution. This section aims to sketch a simple ‘targeting calculus’ to compositionally verify this property. We frame this correctness in terms of weak convergence of measures; while other notions of convergence could be used, weak convergence is standard and will suffice for our purposes.

A. The empirical transformation

First, we need to formalise what we mean when we say that a sampler $s : \Sigma \mathbf{T}$ *targets* a probability distribution on $\llbracket \mathbf{T} \rrbracket$. Given a topological space X , let us write $\mathcal{P}X$ for the space of probability measures on (the Borel σ -algebra generated by) X , equipped with the topology of weak convergence, i.e. $\lim_{n \rightarrow \infty} \mu_n = \mu$ in $\mathcal{P}X$ if for any *bounded continuous* map $f : X \rightarrow \mathbb{R}$, $\lim_{n \rightarrow \infty} \int f d\mu_n = \int f d\mu$. In fact, \mathcal{P} defines a functor $\mathbf{Top} \rightarrow \mathbf{Top}$: if $f : X \rightarrow Y$ is a continuous map, then $\mathcal{P}(f) \triangleq f_* : \mathcal{P}X \rightarrow \mathcal{P}Y$ is the pushforward map, which is easily shown to be continuous. We do not know if $\mathcal{P}X$ is a CG-space when X is, and in particular we do not know if \mathcal{P} can be given a monad structure on \mathbf{CG} . These questions are, however, orthogonal to this work since \mathcal{P} plays no role in the semantics of § III-C.

For any stream $\sigma : \mathbb{N} \rightarrow X \times \mathbb{R}^+$ we define $\hat{\sigma}_n \in \mathcal{P}X$ as

$$\hat{\sigma}_n \triangleq \frac{1}{n} \sum_{i=1}^n \frac{\pi_2(\sigma(i))}{\sum_{j=1}^n \pi_2(\sigma(j))} \delta_{\pi_1(\sigma(i))},$$

the *empirical distribution* based on the first n (weighted) samples of σ . We also define $\mathcal{P}_{\perp} X \triangleq \mathcal{P}X + 1$, where $1 = \{\perp\}$ is the terminal object and $+$ the coproduct in \mathbf{Top} .

Definition V.1. *The empirical measure transformation is the $\mathbf{Top}^{\text{obj}}$ -collection of maps*

$$\varepsilon_X : (X \times \mathbb{R}^+)^{\mathbb{N}} \rightarrow \mathcal{P}_{\perp} X, \sigma \mapsto \begin{cases} \lim_{n \rightarrow \infty} \hat{\sigma}_n & \text{if it exists} \\ \perp \in 1 & \text{else} \end{cases}$$

The empirical measure transformation cannot be natural, as the following example shows.

Example V.1. *Let $\sigma : \mathbb{N} \rightarrow X$ be a diverging unweighted sampler on X , i.e. $\varepsilon_X(\sigma) = \perp$, and consider the map to the terminal object $! : X \rightarrow 1$. Then $\varepsilon_1(! \circ \sigma) = \varepsilon_1(\perp, \perp, \dots) = \delta_{\perp}$, but $\mathcal{P}_{\perp}(!)(\varepsilon_X(\sigma)) = \mathcal{P}_{\perp}(!)(\perp) = \perp$.*

However, if a sampler does define a probability measure via ε , this is preserved by continuous maps.

Proposition V.1. [10, Appendix A] *Let $\sigma : \mathbb{N} \rightarrow X \times \mathbb{R}^+$ and $f : X \times \mathbb{R}^+ \rightarrow Y \times \mathbb{R}^+$ be continuous. If $\varepsilon_X(\sigma) = \mu$, then $\varepsilon_Y(f \circ \sigma) = f_*(\mu)$.*

It is tempting to try to generalise this nice property of continuous maps to more general maps – for example, measurable maps. The following example shows that this is not possible.

Example V.2. *Let $X = [0, 1]$ and $\sigma : \mathbb{N} \rightarrow [0, 1]$ denote any unweighted sampler such that $\varepsilon(\sigma)$ is the Lebesgue measure on $[0, 1]$. Now consider the map $f : [0, 1] \rightarrow \{0, 1\}$ defined by $f(x) = 1$ if $x = \sigma(i)$ for some i and 0 else. This function is the indicator function of a countable, therefore closed, set, and so is Borel-measurable. On the one hand we have that $\varepsilon(f \circ \sigma) = \delta_1$ since $f \circ \sigma$ is the constant stream on ones, but on the other we have $f_*(\varepsilon(\sigma))(1) = \varepsilon(\sigma)(f^{-1}(1)) = 0$ since only a countable set is mapped onto 1 by f .*

Even for functions with finitely many discontinuities, it is impossible to extend the class of functions for which Prop. V.1

holds. However, the semantic framework adopted in § III-C allows us to bypass this problem altogether. We illustrate these two points by revisiting Ex. III.1.

Example V.3. Consider the sampler $\mathfrak{s} \triangleq \text{prng}(\lambda x : \mathbb{R} . x/2, 1)$ and the term $\mathfrak{p} \triangleq \text{if } x = 0 \text{ then } 1 \text{ else } -1$ of Ex. III.1. Assume first that \mathbb{R} is equipped with its standard topology, i.e. that $\llbracket \mathfrak{p} \rrbracket$ is not continuous at 0. Since \mathbb{R} is a metric space we can use the Portmanteau Lemma and rephrase weak convergence by limiting ourselves to bounded Lipschitz functions. It is then easy to show that $\varepsilon(\llbracket \mathfrak{s} \rrbracket) = \delta_0$: letting $f : \mathbb{R} \rightarrow \mathbb{R}$ be bounded Lipschitz, we have

$$\begin{aligned} \lim_{n \rightarrow \infty} \left| \int f d\widehat{\llbracket \mathfrak{s} \rrbracket}_n - \int f d\delta_0 \right| &= \lim_{n \rightarrow \infty} \left| \frac{1}{n} \sum_{i=1}^n f\left(\frac{1}{2^i}\right) - f(0) \right| \\ &\leq \lim_{n \rightarrow \infty} \left| \frac{1}{n} \sum_{i=1}^n \frac{1}{2^i} \right| \leq \lim_{n \rightarrow \infty} \frac{2}{n} = 0 \end{aligned}$$

Prop. V.1 now fails on $\llbracket \mathfrak{p} \rrbracket$, since $\varepsilon(\llbracket \mathfrak{p} \rrbracket \circ \llbracket \mathfrak{s} \rrbracket) = \varepsilon(-1, -1, \dots) = \delta_{-1} \neq \mathcal{P}(\llbracket \mathfrak{p} \rrbracket)(\varepsilon(\llbracket \mathfrak{s} \rrbracket)) = \llbracket \mathfrak{p} \rrbracket_*(\delta_0) = \delta_1$.

Let us now equip \mathbb{R} with the topology given by type-checking \mathfrak{p} as described in Ex. III.1. This makes $\llbracket \mathfrak{p} \rrbracket$ bounded and continuous, and we therefore no longer have $\varepsilon(\llbracket \mathfrak{s} \rrbracket) = \delta_0$; indeed $\lim_n \int \llbracket \mathfrak{p} \rrbracket d\widehat{\llbracket \mathfrak{s} \rrbracket}_n = -1 \neq \llbracket \mathfrak{p} \rrbracket(0) = 1$. In fact we now have $\varepsilon(\llbracket \mathfrak{s} \rrbracket) = \perp$, i.e. \mathfrak{s} is no longer a sampler targeting anything for this topology, which prevents the failure of Prop. V.1 on $\llbracket \mathfrak{p} \rrbracket$.

This example also shows that our semantics has provided us with many more morphisms satisfying Prop. V.1 than would have been the case had we only considered programs which are continuous w.r.t. the usual topology on the denotation of types. Our semantics allows us to push forward a sampler s through any piecewise continuous function, except in the narrow case where this function has a point of discontinuity which is asymptotically assigned positive mass by s . We illustrate this further in the next example.

Example V.4. Consider the sampler of Ex. V.3, but now let $\mathfrak{p} \triangleq \text{if } x = 2 \text{ then } 1 \text{ else } -1$ instead. To make this function continuous, our semantics adds the open set $\{2\}$ to the usual topology of \mathbb{R} . This does not interfere with the derivation that $\varepsilon(\llbracket \mathfrak{s} \rrbracket) = \delta_0$, since we can write $\int f d\widehat{\llbracket \mathfrak{s} \rrbracket}_n = \int_{\{2\}^c} f d\widehat{\llbracket \mathfrak{s} \rrbracket}_n + \int_{\{2\}} f d\widehat{\llbracket \mathfrak{s} \rrbracket}_n = \int_{\{2\}^c} f d\widehat{\llbracket \mathfrak{s} \rrbracket}_n$, and similarly for δ_0 . Because the discontinuity of $\llbracket \mathfrak{p} \rrbracket$ is not assigned any mass by δ_0 , the topology on \mathbb{R} making $\llbracket \mathfrak{p} \rrbracket$ continuous no longer prevents $\varepsilon(\llbracket \mathfrak{s} \rrbracket)$ from converging, and we can therefore safely push \mathfrak{s} forward through \mathfrak{p} using map.

B. Calculus for asymptotic targeting

Definition V.2. We will say that the sampler $\Gamma \vdash s : \Sigma S$ asymptotically targets, or simply targets, the continuous map $\mu : \llbracket \Gamma \rrbracket \rightarrow \mathcal{P} \llbracket S \rrbracket$ if for every $\gamma \in \llbracket \Gamma \rrbracket$,

$$\varepsilon_{\llbracket S \rrbracket} \circ \llbracket s \rrbracket (\gamma) = \mu(\gamma).$$

In particular, $\widehat{\llbracket s \rrbracket}(\gamma)_n$ always converges as $n \rightarrow \infty$; diverging samplers do not target anything.

We will say that $\Gamma \vdash s : \Sigma S$ is K -equidistributed with respect to the morphism $\mu : \llbracket \Gamma \rrbracket \rightarrow \mathcal{P} \llbracket S \rrbracket$ if for every $\gamma \in \llbracket \Gamma \rrbracket$, $\varepsilon_{\llbracket S \rrbracket} \circ \llbracket s^K \rrbracket (\gamma) = \mu^K(\gamma)$, where the self-product s^K is defined in (3) and $\mu^K(\gamma) \in \mathcal{P}(\llbracket S \rrbracket^K)$ is the K -fold product of the measure $\mu(\gamma)$ with itself.

We introduce in Table V a relation \rightsquigarrow which is sound with respect to asymptotic targeting; this is the relation which was used in the proofs of § II. That is, if $\Gamma \vdash s : \Sigma S \rightsquigarrow \mu$, then s is a parametrised sampler on S which asymptotically targets a parametrised distribution μ on $\llbracket S \rrbracket$. Here, we use Greek lower case letters μ, ν to represent (parametrised) distributions in order to emphasise their role as meta-variables, used only in the context of the targeting calculus, and not within the language itself. In the rule for `reweight`, we abbreviate the operation of reweighting a measure μ on X by $f : X \rightarrow \mathbb{R}_{\geq 0}$ as the product $(f \cdot \mu)(A) = \frac{\int_A f(x) d\mu(x)}{\int_X f(x) d\mu(x)}$, assuming that the integral in question is finite and nonzero.

Table V incorporates a rule for building samplers from scratch as pseudo-random number generators defined by a deterministic endomap $t : T \rightarrow T$ and an initial value $x : T$ via $\text{prng}(t, x) : \Sigma T$.⁴ However, Table V also incorporates a set of ‘axioms’ for built-in samplers rand_i , each targeting distributions $\mu_i \in \mathcal{P} \llbracket T_i \rrbracket$; in some settings, it may be defensible to assume access to ‘truly random’ samplers which generate samples using a physical process.

The reader might wonder why Table V does not have a rule for the `thin` operation: after all, if σ is a sampler targeting a distribution μ , then only keeping every n samples should produce a good sampler as well. Whilst this is true of the sequences produced by ‘true’ i.i.d. samplers (for example physical samplers) with probability 1, this rule is in general not sound, as the following simple example shows.

Example V.5. Consider the sampler on $\{0, 1\}$ defined by the program $\text{prng}(\lambda x : \mathbb{R} . 1 - x, 0)$. This sampler, which generates the unweighted samples $(0, 1, 0, 1, \dots)$, targets the uniform Bernoulli distribution; however, applying `thin`(2, `-`) to it yields a sampler which targets the Dirac measure δ_0 .

This example highlights the fact that samplers can be manifestly non-random, and yet from the perspective of inference – that is to say, from the perspective of the topology of weak convergence – target bona fide probability distributions.

Theorem V.1. [10, Appendix A] Targeting \rightsquigarrow is sound: if $\Gamma \vdash s : \Sigma S \rightsquigarrow \mu$, then $\varepsilon_{\llbracket S \rrbracket} \circ \llbracket s \rrbracket = \mu$.

We saw in § V-A how our (sub-)typing system can be used to safely pushforward samplers through maps which are only piecewise continuous. Our typing system also allows us to add additional constraints to samplers. Specifically, we can ensure that a sampler visits certain subsets infinitely often.

⁴Applying this rule requires showing that the initial point of the sampler is *typical*; a point $x \in \llbracket T \rrbracket$ is called *typical* if it belongs to the μ -mass 1 subset $X \subseteq \llbracket T \rrbracket$ in which the ergodic theorem holds [19, Theorem 9.6].

$$\begin{array}{c}
\frac{}{\Gamma_i \vdash \mathbf{rand}_i : \Sigma T_i \rightsquigarrow \mu_i} \quad i \in I \qquad \frac{\Gamma \vdash s \approx t : \Sigma T \quad \Gamma \vdash s : \Sigma T \rightsquigarrow \mu}{\Gamma \vdash t : \Sigma T \rightsquigarrow \mu} \qquad \frac{\Gamma \vdash s : \Sigma S \rightsquigarrow \mu}{\Gamma \vdash \mathbf{tl}(s) : \Sigma S \rightsquigarrow \mu} \\
\frac{\Gamma \vdash s : \Sigma S \rightsquigarrow \mu \quad \Gamma \vdash f : S \rightarrow T}{\Gamma \vdash \mathbf{map}(f, s) : \Sigma T \rightsquigarrow \gamma \mapsto ([f](\gamma))_* \mu(\gamma)} \qquad \frac{\Gamma \vdash s : \Sigma S \rightsquigarrow \mu \quad \Gamma \vdash f : S \rightarrow \mathbb{R}^+}{\Gamma \vdash \mathbf{reweight}(f, s) : \Sigma S \rightsquigarrow \gamma \mapsto [f](\gamma) \cdot \mu(\gamma)} \quad \int_{[S]} [f](\gamma) d\mu(\gamma) \in (0, \infty) \\
\frac{}{\Gamma \vdash \mathbf{prng}(f, x) \rightsquigarrow \mu} \quad [f] : [T] \rightarrow [T] \text{ ergodic w.r.t. } \mu, x \text{ typical}
\end{array}$$

TABLE V: Rules for asymptotic targeting

$$\begin{array}{c}
\frac{}{\vdash \mathbf{rand}^3 : \Sigma S \rightsquigarrow U^3} \quad \vdash \mathbf{id}_R \times \mathbf{box} : S \rightarrow T \\
\frac{}{\vdash \mathbf{map}(\mathbf{id}_R \times \mathbf{box}, \mathbf{rand}^3) : \Sigma T \rightsquigarrow [\mathbf{id}_R \times \mathbf{box}]_* (U^3) = U \otimes N(0, 1)} \quad \vdash \mathbf{accept} : T \rightarrow \mathbb{R}^+ \\
\frac{}{\vdash \mathbf{reweight}(\mathbf{accept}, \mathbf{joint}) : \Sigma T \rightsquigarrow [\mathbf{accept}] \cdot (U \otimes N(0, 1))} \quad \vdash \mathbf{produce} : T \rightarrow \mathbb{R} \\
\frac{}{\vdash \mathbf{map}(\mathbf{produce}, \mathbf{reweight}(\mathbf{accept}, \mathbf{joint})) : \Sigma \mathbb{R}^+ \rightsquigarrow [\mathbf{produce}]_* ([\mathbf{accept}] \cdot (U \otimes N(0, 1))) = \Gamma(\alpha, 1)}
\end{array}$$

Fig. 5: Validity of Marsaglia sampler in Listing 4

Proposition V.2. [10, Appendix A] Assume $\Gamma \vdash s : \Sigma S \rightsquigarrow \mu$, $S \triangleleft T$ and $[T]$ second-countable; then $\Gamma \vdash \mathbf{map}(\lambda x : S.\mathbf{cast}(T)x, s) : \Sigma T$ targets the same measure μ on T . Moreover, if $[T]$ is metrizable, if U is in the topology of $[S]$ but not $[T]$ and $\mu(\partial_T U) > 0$ (where ∂_T denotes the boundary in $[T]$) then s must visit $\partial_T U$ i.o. (infinitely often).

Example V.6. Suppose we want $s : \Sigma \mathbb{R} \rightsquigarrow \text{Bern}(1/2)$. A sampler alternating between the sampler $z \triangleq \mathbf{prng}(\lambda x : \mathbb{R} . x/2, 1)$ of Ex. V.3 and its shifted version $\mathbf{map}(\lambda x : \mathbb{R} . 1 + x, z)$ will satisfy the condition, but will never visit 0 or 1! We can use the previous result to enforce that a sampler s targeting $\text{Bern}(1/2)$ should visit 0 i.o. by constructing s in such a way that it has type $\Sigma((x, 0)^{-1}\text{Neq} + (x, 0)^{-1}\text{Eq})$ (see Ex. III.1). We can, in the same manner, enforce that a sampler s' targeting $\text{Bern}(1/2)$ visits 1 i.o. Finally, using the last two rules of Fig. 4b which build the coarsest common refinement of two topologies, we can combine s and s' to create a sampler targeting $\text{Bern}(1/2)$ and guaranteed to visit 0, 1 i.o.

Example V.7. We conclude with an example highlighting sampler compositionality by chaining two well-known sampling algorithms. Consider the following program:

```

let box = λu : ℝ+ × ℝ+ . sqrt(-2*log(fst(u))) *
  cos(2*pi*snd(u)) in
let joint = map(idℝ × box, rand3) in
let d = α - 1/3 in
let c = 1/(3*sqrt(d)) in
let accept = λ(u, x) : T .
  let v = (1+c*x)^3 in
  if v > 0 and log(u) < x^2 + d - d*v + d*
    log(v) then 1 else 0 in
let produce = λz : T . d*(1+c*fst(cast(ℝ × ℝ)z))
  ^3 in
map(proj, reweight(accept, joint))

```

Listing 4: Marsaglia sampler for gamma random variables

First, the Box-Muller technique is a well-known technique for generating standard normal random variates using two independent uniform samples; its verification using the \mathbf{map} rule of Table V is straightforward. We then form a joint sampler consisting of independent uniform and Gaussian samples, and then consume both of these samples to generate gamma-

distributed random variables $z \sim \Gamma(\alpha, 1)$, for shape $\alpha \geq 1$, using a well-known rejection sampling technique; see [28] for a proof. The validity of this sampler is sketched in Fig. 5; we omit some types for brevity.

VI. DISCUSSION

We have presented a ‘probabilistic’ language designed to compositionally construct *samplers*. We have given this language an intuitive operational semantics and a denotational semantics in the category of CG-spaces, and shown that the two are equivalent for closed samplers. This denotational universe is sufficiently rich to interpret sampler types coinductively, and to interpret functions which are only piecewise-continuous on the standard topologies given by the type system (§ III-C).

With the support of this language, we have shown how to compositionally reason about the validity of sampler constructions, an essential aspect in the practice of probabilistic programming. Our approach draws on a sound equational system to reason about equivalent ways of constructing the same sampler (§ IV) and a sound system for reasoning about semantic correctness (§ V).

What distinguishes our approach is that we are in effect providing a purely deterministic semantics for probabilistic programs. This approach is much closer to the practice of probabilistic programming, in which samples and samplers are the most important concrete entities; this distinction between samplers and the measures they target is necessary in order to support pseudo-random number generation. Measure-theoretic entities, which have typically been a part of the denotational semantics of probabilistic languages, e.g. [21], [16], [12], [38], [8], instead take a meta-theoretic role as verification criteria.

Two commonly-used schemes for producing samplers are missing from our calculus: Markov chain Monte Carlo methods and resampling techniques, as applied in e.g. particle filters. We consider adding these constructions to our language, together with the corresponding correctness proofs in our targeting calculus, to be future work.

REFERENCES

- [1] ANDRIEU, C., DOUCET, A., AND HOLENSTEIN, R. Particle markov chain monte carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 72, 3 (2010), 269–342.
- [2] BAGNALL, A., STEWART, G., AND BANERJEE, A. Formally verified samplers from probabilistic programs with loops and conditioning, 2023.
- [3] BINGHAM, E., CHEN, J. P., JANKOWIAK, M., OBERMEYER, F., PRADHAN, N., KARALETSOS, T., SINGH, R., SZERLIP, P. A., HORSFALL, P., AND GOODMAN, N. D. Pyro: Deep universal probabilistic programming. *J. Mach. Learn. Res.* 20 (2019), 28:1–28:6.
- [4] BLACKWELL, A., KOHN, T., ERWIG, M., BAYDIN, A. G., CHURCH, L., GEDDES, J., GORDON, A., GORINOVA, M., GRAM-HANSEN, B., LAWRENCE, N., MANSINGHA, V., PAIGE, B., PETRICEK, T., ROBINSON, D., SARKAR, A., AND STRICKSON, O. Usability of probabilistic programming languages. In *Psychology of Programming Interest Group Annual Workshop (PPIG 2019), Newcastle, UK, 28–30 August 2019* (2019).
- [5] BORGSTRÖM, J., DAL LAGO, U., GORDON, A. D., AND SZYMCAK, M. A lambda-calculus foundation for universal probabilistic programming. *ACM SIGPLAN Notices* 51, 9 (2016), 33–46.
- [6] BROOKS, S., GELMAN, A., JONES, G., AND MENG, X.-L. *Handbook of Markov Chain Monte Carlo*. CRC press, 2011.
- [7] CARPENTER, B., GELMAN, A., HOFFMAN, M. D., LEE, D., GOODRICH, B., BETANCOURT, M., BRUBAKER, M., GUO, J., LI, P., AND RIDDELL, A. Stan: A probabilistic programming language. *Journal of statistical software* 76, 1 (2017).
- [8] DAHLQVIST, F., AND KOZEN, D. Semantics of higher-order probabilistic programs with conditioning. *Proceedings of the ACM on Programming Languages* 4, POPL (2019), 1–29.
- [9] DAHLQVIST, F., KOZEN, D., AND SILVA, A. *Semantics of Probabilistic Programming: A Gentle Introduction*. Cambridge University Press, 2020, pp. 1–42.
- [10] DAHLQVIST, F., SILVA, A., AND SMITH, W. Deterministic stream-sampling for probabilistic programming: semantics and verification. *arXiv preprint* (2023).
- [11] EHRHARD, T., PAGANI, M., AND TASSON, C. The computational meaning of probabilistic coherence spaces. In *2011 IEEE 26th Annual Symposium on Logic in Computer Science* (2011), IEEE, pp. 87–96.
- [12] EHRHARD, T., PAGANI, M., AND TASSON, C. Measurable cones and stable, measurable functions: a model for probabilistic higher-order programming. *Proceedings of the ACM on Programming Languages* 2, POPL (2017), 1–28.
- [13] EHRHARD, T., PAGANI, M., AND TASSON, C. Full abstraction for probabilistic pcf. *Journal of the ACM (JACM)* 65, 4 (2018), 1–44.
- [14] FAGGIAN, C., AND DELLA ROCCA, S. R. Lambda calculus and probabilistic computation. In *2019 34th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)* (2019), IEEE, pp. 1–13.
- [15] GOODMAN, N., MANSINGHA, V., ROY, D. M., BONAWITZ, K., AND TENENBAUM, J. B. Church: a language for generative models. *arXiv preprint arXiv:1206.3255* (2012).
- [16] HEUNEN, C., KAMMAR, O., STATON, S., AND YANG, H. A convenient category for higher-order probability theory. In *2017 32nd Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)* (2017), IEEE, pp. 1–12.
- [17] HOYRUP, M., AND ROJAS, C. Computability of probability measures and martin-löf randomness over metric spaces. *Information and Computation* 207, 7 (2009), 830–847.
- [18] HUANG, D., MORRISSETT, G., AND SPITTERS, B. *Application of Computable Distributions to the Semantics of Probabilistic Programs*. Cambridge University Press, 2020, p. 75–120.
- [19] KALLENBERG, O. *Foundations of modern probability*. Springer, 1997.
- [20] KEREMEDIS, K., ÖZEL, C., PIEKOSZ, A., SHUMRANI, M. A., AND WAJCH, E. Compact complement topologies and k-spaces. *arXiv preprint arXiv:1806.10177* (2018).
- [21] KOZEN, D. Semantics of probabilistic programs. *J. Comput. Syst. Sci.* 22, 3 (June 1981), 328–350.
- [22] LAGO, U. D., AND ZORZI, M. Probabilistic operational semantics for the lambda calculus. *RAIRO-Theoretical Informatics and Applications-Informatique Théorique et Applications* 46, 3 (2012), 413–450.
- [23] L’ECUYER, P. History of uniform random number generation. In *2017 Winter Simulation Conference (WSC)* (2017), pp. 202–230.
- [24] LEOBACHER, G., AND PILLICHSHAMMER, F. *Introduction to quasi-Monte Carlo integration and applications*. Springer, 2014.
- [25] LEWIS, L. G. *The stable category and generalized Thom spectra. Appendix A*. PhD thesis, University of Chicago, Department of Mathematics, 1978.
- [26] LUNN, D., THOMAS, A., BEST, N., AND SPIEGELHALTER, D. Winbugs - a bayesian modeling framework: Concepts, structure and extensibility. *Statistics and Computing* 10 (10 2000), 325–337.
- [27] MANSINGHA, V., SELSAM, D., AND PEROV, Y. Venture: a higher-order probabilistic programming platform with programmable inference. *arXiv e-prints* (Mar. 2014), arXiv:1404.0099.
- [28] MARSAGLIA, G., AND TSANG, W. W. A simple method for generating gamma variables. *ACM Trans. Math. Softw.* 26, 3 (sep 2000), 363–372.
- [29] MARTIN-LÖF, P. The definition of random sequences. *Information and control* 9, 6 (1966), 602–619.
- [30] MATSUMOTO, M., AND NISHIMURA, T. Mersenne twister: A 623-dimensionally equidistributed uniform pseudo-random number generator. *ACM Trans. Model. Comput. Simul.* 8, 1 (1998), 3–30.
- [31] MCCORD, M. C. Classifying spaces and infinite symmetric products. *Transactions of the American Mathematical Society* 146 (1969), 273–298.
- [32] PARK, S., PFENNING, F., AND THRUN, S. A probabilistic language based upon sampling functions. *ACM SIGPLAN Notices* 40, 1 (2005), 171–182.
- [33] ROBERT, C., AND CASELLA, G. *Monte Carlo statistical methods*. Springer Science & Business Media, 2013.
- [34] STATON, S., WOOD, F., YANG, H., HEUNEN, C., AND KAMMAR, O. Semantics for probabilistic programming: higher-order functions, continuous distributions, and soft constraints. In *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)* (2016), IEEE, pp. 1–10.
- [35] STEENROD, N. E. A convenient category of topological spaces. *Michigan Mathematical Journal* 14, 2 (1967), 133–152.
- [36] THOMAS, A. Bugs: a statistical modelling package. *RTA/BCS Modular Languages Newsletter* 2 (1994), 36–38.
- [37] TREVISAN, L., AND VADHAN, S. Extracting randomness from samplable distributions. In *Proceedings of the 41st Annual Symposium on Foundations of Computer Science (USA, 2000), FOCS ’00*, IEEE Computer Society, p. 32.
- [38] VÁKÁR, M., KAMMAR, O., AND STATON, S. A domain theory for statistical probabilistic programming. *Proceedings of the ACM on Programming Languages* 3, POPL (2019), 1–29.
- [39] VIGNA, S. Further scramblings of marsaglia’s xorshift generators. *Journal of Computational and Applied Mathematics* 315 (2017), 175–181.
- [40] WOOD, F., VAN DE MEENT, J. W., AND MANSINGHA, V. A new approach to probabilistic programming inference. In *Proceedings of the 17th International conference on Artificial Intelligence and Statistics* (2014), pp. 1024–1032.