



A Higher-Order Language for Markov Kernels and Linear Operators

Pedro H. Azevedo de Amorim^(✉)

Cornell University, Ithaca, NY, USA
pamorim@cs.cornell.edu

Abstract. Much work has been done to give semantics to probabilistic programming languages. In recent years, most of the semantics used to reason about probabilistic programs fall in two categories: semantics based on Markov kernels and semantics based on linear operators. Both styles of semantics have found numerous applications in reasoning about probabilistic programs, but they each have their strengths and weaknesses. Though it is believed that there is a connection between them there are no languages that can handle both styles of programming. In this work we address these questions by defining a two-level calculus and its categorical semantics which makes it possible to program with both kinds of semantics. From the logical side of things we see this language as an alternative resource interpretation of linear logic, where the resource being kept track of is sampling instead of variable use.

Keywords: Linear Logic, Probabilistic Programming, Categorical Semantics.

1 Introduction

Probabilistic primitives have been a standard feature of programming languages since the 70s. At first, randomness was mostly used to program so called random algorithms, i.e. algorithms that require access to a source of randomness. Recently, however, with the rise of computational statistics and machine learning, randomness is also used to program statistical models and inference algorithms.

Programming languages researchers have seen this rise in interest as an opportunity to further study the interaction of probability and programming languages, establishing it as an active subfield within the PL community.

One of the main goals of this subfield is giving semantics to programming languages that are both expressive in the regular PL sense as well as in its abilities to program with randomness. One particular difficulty is that the mathematical machinery used for probability theory, i.e. measure theory, does not interact well with higher-order functions [2].

Currently, there are two classes of models of probabilistic programming — in its broad sense — that have found numerous applications: models based on linear logic and models based on Markov kernels. Since each kind of semantics has peculiarities that make them more or less adequate to give semantics to expressive programming languages, it is an important theoretical question to understand how these classes of models are related.

Linear Logic for Probabilistic Semantics The models of linear logic that have been used to give semantics to probabilistic languages are usually based on categories of vector spaces where programs are denoted by linear operators. We highlight two of them:

- Ehrhard et. al [11,10,9] have defined models of linear logic with probabilistic primitives and have used the translation of intuitionistic logic into linear logic $A \rightarrow B = !A \multimap B$, where $!A$ is the exponential modality, to give semantics to a stochastic λ -calculus.
- Dahlqvist and Kozen [8] have defined an imperative, higher-order, linear probabilistic language and added a type constructor $!$ to accommodate non-linear programs.

The main advantage of models based on linear logic is that programs are denoted by linear operators between spaces of distributions, a formalism that has been extensively used to reason about stochastic processes, as illustrated by Dahlqvist and Kozen who have used results from ergodic theory to reason about a Gibbs sampling algorithm written in their language, and by Clerc et al. who have shown how Bayesian inference can be given semantics using adjoint of linear operators [7].

Unfortunately, these insights are hard to realize in practice, since languages based on linear logic enforce that variables must be used exactly once, making it hard to use it as a programming language. The usual way linear logic deals with this limitation is through the $!$ modality which allows variables to be reused.

The problem with the exponential modality, when it comes to probabilistic programming, is that they are usually difficult to construct, do not have any clear interpretation in terms of probability, making the linear operator formalism not applicable anymore and, more operationally, through its connections with call-by-name (CBN) semantics [18], makes it mathematically hard to reuse sampled values.

Ehrhard et al. have found a way around this problem by introducing a call-by-value (CBV) let operator that allows samples to be reused [11,24]. In the discrete case this operator is elegantly defined by a categorical argument which is unknown to scale to the continuous case, which they deal with by making use of an ad-hoc construction that is unclear if it can be generalized to other models of linear logic. Therefore, our current understanding of models of linear logic does not provide a uniform way of reusing samples.

The difference between CBV and CBN can be illustrated by the program let $x = \text{coin}$ in $x + x$, where coin is a primitive that outputs 0 or 1 with equal probability. In the CBN semantics each use of x corresponds to a new sample from coin , whereas in the CBV semantics the coin is only sampled once.

A subtler problem of probabilistic models based on linear logic is that they are ill-equipped to program with joint distributions. For instance, the language proposed by Ehrhard et. al can be easily extended with product types which, under their semantics, would make the type $\mathbb{R} \times \mathbb{R}$ be interpreted as $\mathcal{M}\mathbb{R} \times \mathcal{M}\mathbb{R}$, where $\mathcal{M}\mathbb{R}$ is the set of distributions over \mathbb{R} – which is isomorphic to the set of independent distributions over \mathbb{R}^2 . Dahlqvist and Kozen deal with this issue by

adding primitive types \mathbb{R}^n to their language which are interpreted as the set of joint distributions over \mathbb{R}^n . However, since they are not defined using the type constructors provided by the semantic domain, programs of type \mathbb{R}^n can only be manipulated by primitives defined outside the language.

Markov Kernel Semantics Markov kernels are a generalization of transition matrices, i.e. functions that map states to probability distributions over them. They are appealing from a programming languages perspective because their programming model is usually captured by monads and Kleisli arrows, a common abstraction in programming languages semantics, and have been extensively used to reason about probabilistic programs [1,22,3]. By being related to monadic programming they differ from their linear operator counterpart by being able to naturally capture a call-by-value semantics which, as we argued above, is the most natural one for probabilistic programming.

Unfortunately, even though these semantics can be generalized to continuous distributions, they are notoriously brittle when it comes to higher-order programming. Only recently, with the introduction of quasi Borel spaces [15] and its probability monad, it is possible to give a kernel-centric semantics to higher-order probabilistic programming with continuous distributions.

However, due to quasi Borel spaces being a different foundation to probability theory, it is unclear which theorems and theories can be generalized to higher-order. For instance, martingale theory has been used in Computer Science to reason about termination of probabilistic programs [6,20,16]. In order to generalize these ideas to higher-order functions it would be necessary to define a quasi Borel version of martingales and prove appropriate versions of the main theorems from martingale theory, a non-trivial task.

Our Work: Combining both Kinds of Semantics Though both styles of semantics provide insights into how to interpret probabilistic programming languages (PPL), it is still too early to claim that we have a “correct” semantics which subsumes all of the existing ones. Both approaches mentioned above have their advantages and drawbacks.

In this work we shed some light into how both semantics relate to one another by showing that it is possible to use both styles of semantics to interpret a linear calculus that has higher-order functions, looser linearity restrictions, a uniform way of dealing with sample reuse and better syntax for programming joint distributions while still being close to their kernel and linear operator counterparts. Interestingly, we identify the joint distribution problem described above to be a consequence of linear logic requiring the non-linear product to be cartesian. In order to tackle this problem we build on categorical semantics of linear logic and on recent work on Markov categories, a suitable categorical generalization of Markov kernels defined using semicartesian products.

We bridge the gap between these semantics by noting that the regular resource interpretation of linear logic, i.e. $A \multimap B$ being equivalent to “by using one copy of A I get one copy of B ” is too restrictive an interpretation for probabilistic

programming. Instead, we should think of usage as being equivalent to sampling. Therefore the linear arrow $A \multimap B$ should be thought of as “by sampling from A once I get B ”, which is the computational interpretation of Markov kernels.

We realize this interpretation through a multilanguage approach: we have one language that programs Markov kernels, a second language that programs linear operators and add syntax that transports programs from the former language into the latter one. To justify the viability of our categorical framework we show how existing probabilistic semantics are models to our language and show how, under mild conditions, this semantics can be generalized to commutative effects.

Our contributions are:

- We define a multi-language syntax that can program both Markov kernels as well as linear operators. (§3)
- We define its categorical semantics and prove certain interesting equations satisfied by it. (§4)
- We show that our semantics is already present in existing models for discrete and continuous probabilistic programming. (§5)
- We show how our semantics can be generalized to commutative effects. (§6)

2 Mathematical Preliminaries

We are assuming that the reader is familiar with basic notions from category theory such as categories, functors and monads.

Probability Theory

Transition matrices are one of the simplest abstractions used to model stochastic processes. Given two countable sets A and B , the entry (a, b) of a transition matrix is the probability of ending up in state $b \in B$ whenever you start from the initial state $a \in A$ and every row adds up to 1.

Definition 1. *The category **CountStoch** has countable sets as objects and transition matrices as morphisms. The identity morphism is the identity matrix and composition is given by matrix multiplication.*

Though transition matrices are conceptually simple, they can only model discrete probabilistic processes and, in order to generalize them to continuous probability we must use measurable sets and Markov kernels.

Definition 2. *A measurable set is a pair (A, Σ_A) , where A is a set and $\Sigma_A \subseteq \mathcal{P}(A)$ is a σ -algebra, i.e. it contains the empty set and it is closed under complements and countable unions.*

Definition 3. *A function $f : (A, \Sigma_A) \rightarrow (B, \Sigma_B)$ is called measurable if for every $\mathcal{B} \in \Sigma_B$, $f^{-1}(\mathcal{B}) \in \Sigma_A$.*

Definition 4. Let (A, Σ_A) be a measurable space. A probability distribution (A, Σ_A) is a function $\mu : \Sigma_A \rightarrow [0, 1]$ such that $\mu(\emptyset) = 0$, $\mu(A) = 1$ and $\mu(\bigsqcup_{i \in \mathbb{N}} A_i) = \sum_{i \in \mathbb{N}} \mu(A_i)$.

Given two measurable sets (A, Σ_A) and (B, Σ_B) it is possible to define a σ -algebra over $A \times B$ generated by the sets $X \times Y$ which we denote by $\Sigma_A \otimes \Sigma_B$, where $X \in \Sigma_A$ and $Y \in \Sigma_B$. Furthermore, every pair of distributions μ_A and μ_B over A and B respectively, can be lifted to a distribution $\mu_A \otimes \mu_B$ over $A \times B$ such that $(\mu_A \otimes \mu_B)(X \times Y) = \mu_A(X)\mu_B(Y)$, for $X \in \Sigma_A$ and $Y \in \Sigma_B$.

Definition 5. Let (A, Σ_A) and (B, Σ_B) be two measurable spaces. A Markov kernel is a function $f : A \times \Sigma_B \rightarrow [0, 1]$ such that

- For every $a \in A$, $f(a, -)$ is a probability distribution.
- For every $\mathcal{B} \in \Sigma_B$, $f(-, \mathcal{B})$ is a measurable function.

Definition 6. The category **Kern** has measurable sets as objects and Markov kernels as morphisms. The identity arrow is the function $id_A(a, \mathcal{A}) = 1$ if $a \in \mathcal{A}$ and 0 otherwise and Composition is given by $(f \circ g)(a, \mathcal{C}) = \int f(-, \mathcal{C})d(g(a, -))$.

Markov Categories

The field of categorical probability was developed in order to get a more conceptual understanding of Markov kernels. One of its cornerstone definitions is that of a Markov category which are categories where objects are abstract sample spaces, morphisms are abstract Markov kernels and every object has “contraction” and “weakening” morphisms which correspond to duplicating and discarding a sample, respectively, without adding any new randomness.

Definition 7 (Markov category [12]). A Markov category is a semicartesian symmetric monoidal category $(\mathbf{C}, \otimes, 1)$ in which every object A comes equipped with a commutative comonoid structure, denoted by $\text{copy}_X : X \rightarrow X \otimes X$ and $\text{delete}_X : X \rightarrow 1$, where copy satisfies

$$\text{copy}_{X \otimes Y} = (id_X \otimes b_{Y, X} \otimes id_Y) \circ (\text{copy}_X \otimes \text{copy}_Y),$$

where $b_{Y, X}$ is the natural isomorphism $Y \otimes X \cong X \otimes Y$. The category being semicartesian means that the monoidal product comes equipped with projection morphisms $\pi_1 : A \otimes B \rightarrow A$ and $\pi_2 : A \otimes B \rightarrow B$, but it is not Cartesian because the equation $(\pi_1 \circ f, \pi_2 \circ f) = f$ does not hold in general which, intuitively, corresponds to the fact that joint distributions might be correlated.

Theorem 1 ([12]). **CountStoch** is a Markov category.

The monoidal product is given by the Cartesian product and the monoidal unit is the singleton set. The copy_X morphism is the matrix $X \times X \times X \rightarrow [0, 1]$ which is 1 in the positions (x, x, x) and 0 elsewhere, and the delete_X morphism is the constant 1 matrix indexed by X .

Theorem 2 ([12]). *Kern is a Markov category.*

This category is the continuous generalization of **CountStoch** and the monoidal product is the Cartesian product with the *product* σ -algebra and the monoidal unit is the singleton set $\{*\}$. The copy_X morphism is the Markov kernel $\text{copy}_X : X \times \Sigma_X \otimes \Sigma_X \rightarrow [0, 1]$ such that $\text{copy}_X(x, S \times T) = 1$ if $x \in S \cap T$ and 0 otherwise. Its delete morphism is simply the function that given any element in X , returns the function which is 1 on the measurable set $\{*\}$ and 0 on the empty measurable set.

Linear Logic and Monoidal Categories

We recall the categorical semantics of the multiplicative fragment of linear logic (MLL):

Definition 8 ([21]). *A category \mathbf{C} is an MLL model if it is symmetric monoidal closed (SMCC), i.e. the functors $A \otimes -$ have a right adjoint $A \multimap -$.*

We denote the monoidal product as \otimes and the space of linear maps between objects X and Y as $X \multimap Y$, $\text{ev} : ((X \multimap Y) \otimes X) \rightarrow Y$ is the counit of the monoidal closed adjunction and $\text{cur} : \mathbf{C}(X \otimes Y, Z) \rightarrow \mathbf{C}(X, Y \multimap Z)$ is the linear curryfication map. We use the triple $(\mathcal{C}, \otimes, \multimap)$ to denote such models.

Definition 9. *Let $(\mathbf{C}, \otimes_{\mathbf{C}}, 1_{\mathbf{C}})$ and $(\mathbf{D}, \otimes_{\mathbf{D}}, 1_{\mathbf{D}})$ be two monoidal categories. We say that a functor $F : \mathbf{C} \rightarrow \mathbf{D}$ is lax monoidal if there is a morphism $\epsilon : 1_{\mathbf{D}} \rightarrow F(1_{\mathbf{C}})$ and a natural transformation $\mu_{X,Y} : F(X) \otimes_{\mathbf{D}} F(Y) \rightarrow F(X \otimes_{\mathbf{C}} Y)$ making the diagrams in Figure 8 (in Appendix B) commute.*

If ϵ and $\mu_{X,Y}$ are isomorphisms we say that F is *strong* monoidal.

One key observation of this paper is that there are many lax monoidal functors between Markov categories and models of linear logic that can interpret probabilistic processes.

3 Syntax

In this section we will design a syntax that reflects the fact that linearity corresponds to sampling, not variable usage. We achieve this by making use of a multi-language semantics that enables the programmer to transport programs defined in a Markov kernel-centric language (MK) to a linear, higher-order, language (LL).

Our thesis is that in the context of probabilistic programming, linear logic, through its connection with linear algebra, departs from its usual Computer Science applications of enforcing syntactic invariants and, instead, provides a natural mathematical formalism to express ideas from probability theory, as shown by Dahlqvist and Kozen [8].

Therefore, since many probabilistic programming constructs, such as Bayesian inference and Markov kernels, can be naturally interpreted in linear logic terms,

$$\begin{aligned}
& \tau := 1 \mid \tau \times \tau \\
M, N & := x \mid \text{unit} \mid \text{let } x = M \text{ in } N \mid (M, N) \mid \pi_1 M \mid \pi_2 N \mid f(M) \\
\Gamma & := \cdot \mid x : \tau, \Gamma
\end{aligned}$$

Fig. 1: Syntax MK

we believe that our calculus allows the user to benefit from the insights linearity provides to PPL while unburdening them from worrying about syntactic restrictions by making it possible to also program using kernels.

We use standard notation from the literature: $\Gamma \vdash t : \tau$ means that the program t has type τ under context Γ , $t\{x/u\}$ means substitution of u for x in t and $t\{\vec{x}/\vec{u}\}$ is the simultaneous substitution of the term list \vec{u} for a variable list \vec{x} in t .

Both languages will be defined in this section and, for presentation's sake, we are going to use orange to represent MK programs and purple to represent LL programs.

3.1 A Markov Kernel Language

We need a language to program Markov kernels. Since we are aiming at generality, we are assuming the least amount of structure possible. As such we will be working with the internal language of Markov categories, as presented in Figure 1 and Figure 4¹. Note that we are implicitly assuming a set of primitives for the functions f .

By construction, every Markov category can interpret this language, as we show in Figure 6, with types being interpreted as

$$\begin{aligned}
\llbracket 1 \rrbracket &= 1 \\
\llbracket \tau_1 \times \tau_2 \rrbracket &= \llbracket \tau_1 \rrbracket \times \llbracket \tau_2 \rrbracket
\end{aligned}$$

and the contexts are interpreted using \times over the interpretation of the types. However, as it stands, it is not very expressive, since it does not have any probabilistic primitives nor does it have any interesting types since $1 \times 1 \cong 1$.

When working with concrete models (c.f. Section 5) we can extend the language with more expressive types as well as with concrete probabilistic primitives. For instance, in the context of continuous probabilities we could add a \mathbb{R} datatype and a $\cdot \vdash \text{uniform} : \mathbb{R}$ uniform distribution primitive.

Note that even though this language does not have any explicit sampling operators, this is implicitly achieved by the `let` operator. For instance, the program

¹ c.f. Appendix A.

$$\begin{aligned}
\tau &:= 1 \mid \tau \multimap \tau \mid \tau \otimes \tau \\
t, u &:= x \mid \text{unit} \mid \lambda x. t \mid t u \mid t \otimes u \mid \text{let } x \otimes y = t \text{ in } u \\
\Gamma &:= \cdot \mid x : \tau, \Gamma
\end{aligned}$$

Fig. 2: Syntax LL

`let $x = \text{uniform in } x + x$` samples from a uniform distribution, binds the result to the variable x and adds the sample to itself (Fig. 2).

3.2 A Linear Language

Our second language is a linear simply-typed λ -calculus, with the usual typing rules shown in Figure 5 in Appendix A, which can be interpreted in every symmetric monoidal closed category as shown in Figure 7, also in Appendix A, with types interpreted by

$$\begin{aligned}
\llbracket 1 \rrbracket &= 1 \\
\llbracket \tau_1 \otimes \tau_2 \rrbracket &= \llbracket \tau_1 \rrbracket \otimes \llbracket \tau_2 \rrbracket \\
\llbracket \tau_1 \multimap \tau_2 \rrbracket &= \llbracket \tau_1 \rrbracket \multimap \llbracket \tau_2 \rrbracket
\end{aligned}$$

and the contexts are interpreted using \otimes over the interpretation of the types. Once again, we are aiming at generality instead of expressivity. In a concrete setting it would be fairly easy to extend the calculus with a datatype \mathbb{N} for natural numbers and probabilistic primitives such as $\cdot \vdash \text{coin} : \mathbb{N}$ that flips a fair coin.

The idea behind the particular linear logic models that we are interested in is that, by integration, Markov kernels can be seen as linear operators between vector spaces of probability distributions. As such, an LL program $x : \mathbb{N} \vdash_{LL} t : \mathbb{N}$ will be denoted by a linear function between distributions over the natural numbers. Therefore, from a programming point of view, variables are placeholders for probability distributions, i.e. computations, not values, and sampling occurs when variables are used.

3.3 Combining Languages

The main drawback of the linear calculus above is that the syntactic linearity restriction makes it hard to program with it, while the main drawback of the Markov language is that it does not have higher-order functions. In this section we will show how we can combine both language so that we get a calculus with looser linearity restrictions while still being higher-order.

$$\begin{aligned}
 \tau &:= 1 \mid \tau \times \tau \\
 \underline{\tau} &:= 1 \mid \mathcal{M}\tau \mid \underline{\tau} \multimap \underline{\tau} \mid \underline{\tau} \otimes \underline{\tau} \\
 M, N &:= x \mid \text{unit} \mid \text{let } x = M \text{ in } N \mid f(M) \\
 &\mid (M, N) \mid \pi_1 M \mid \pi_2 M \\
 t, u &:= x \mid \text{unit} \mid \lambda x. t \mid t u \mid t \otimes u \mid \text{let } x \otimes y = t \text{ in } u \\
 &\mid \text{sample } t_i \text{ as } x_i \text{ in } M
 \end{aligned}$$

Fig. 3: Syntax LL+MK

As we will show in Section 5, when looking at concrete models for these languages we can see that the semantic interpretations of variables in both languages are completely different: in the MK language variables should be thought of as values, i.e. the values that were sampled from a distribution, whereas in the LL language, variables of ground type are distributions. In order to bridge these languages we must use the observation that Markov kernels — i.e. open MK terms — have a natural resource-aware interpretation of being “sample-once” stochastic processes and, by integration, can be seen as linear maps between measure spaces — i.e. open LL terms. The combined syntax for the language is depicted in Figure 3.

We now have a language design problem: we want to capture the fact that every open MK program is, semantically, also an open LL term. The naive typing rule is:

$$\frac{x_1 : \tau_1, \dots, x_n : \tau_n \vdash_{MK} M : \tau}{x_1 : \mathcal{M}\tau_1, \dots, x_n : \mathcal{M}\tau_n \vdash_{LL} \text{MK}(M) : \mathcal{M}\tau}$$

The problem with this rule is that it breaks substitution: the variables in the premise are MK variables whereas the ones in the conclusion are LL variables.

We solve this problem by making the syntax reflect a common idiom of PPLs: compute distributions (elements of $\mathcal{M}\tau$), sample from it and then use the result in a non-linear continuation. This is captured by the following syntax:

$$\text{sample } t_1, \dots, t_n \text{ as } x_1, \dots, x_n \text{ in } M$$

Note that we are sampling from LL programs t_i (possibly an empty list), outputting the results to MK variables x_i and binding them to an MK program M . When clear from the context we simply use $\text{sample } t_i \text{ as } x_i \text{ in } M$. Its corresponding typing rule is:

$$\frac{\text{SAMPLE} \quad x_1 : \tau_1 \cdots x_n : \tau_n \vdash_{MK} M : \tau \quad \Gamma_i \vdash_{LL} t_i : \mathcal{M}\tau_i \quad 0 \leq i < n}{\Gamma_1, \dots, \Gamma_n \vdash_{LL} \text{sample } t_i \text{ as } x_i \text{ in } M : \mathcal{M}\tau}$$

As the typing rule suggests, its semantics should be some sort of composition. However, since we are composing programs that are interpreted in different categories, we must have a way of translating MK programs into LL programs — as we will see in Section 4 this translation will be functorial. The operational interpretation of this rule is that we have a set of distributions $\{t_i\}$ defined using the linear language — possibly using higher-order programs — we sample from them, bind the samples to the variables $\{x_i\}$ in the MK program M where there are no linearity restrictions. Note that the rule above looks very similar to a monadic composition, though they are semantically different (cf. Section 4).

With this new syntax we can finally program in accordance with our new resource interpretation of linear logic, allowing us to write the program

$$\text{sample coin as } x \text{ in } (x = x),$$

which flips a coin once and tests the result for equality with itself, making it equivalent to true.

This combined calculus enjoys the expected syntactic properties².

Theorem 3. *Let $\Gamma, x : \underline{\tau}_1 \vdash_{LL} t : \underline{\tau}$ and $\Delta \vdash_{LL} u : \underline{\tau}_1$ be well-typed terms, then $\Gamma, \Delta \vdash_{LL} t\{x/u\} : \underline{\tau}$*

Proof. The proof can be found in Appendix D.

The following example illustrates how we can use the MK language to duplicate and discard linear variables.

Example 1. The program which samples from a distribution t and then returns a perfectly correlated pair is given by:

$$\cdot \vdash_{LL} \text{sample } t \text{ as } x \text{ in } (x, x) : \mathcal{M}(\tau \times \tau)$$

Similarly, the program that samples from a distribution t and does not use its sampled value is represented by the term

$$\cdot \vdash_{LL} \text{sample } t \text{ as } x \text{ in unit} : \mathcal{M}1$$

Example 2. Suppose that we have a Markov kernel given by an open MK term $x : \mathbb{N} \vdash M : \mathbb{N}$. If we want to encapsulate it as a linear program of type $\mathcal{M}\mathbb{N} \multimap \mathcal{M}\mathbb{N}$ we can write:

$$\cdot \vdash_{LL} \lambda \text{meas.}(\text{sample } \text{meas} \text{ as } x \text{ in } M) : \mathcal{M}\mathbb{N} \multimap \mathcal{M}\mathbb{N}$$

Example 3. As we explain in the introduction, Dahlqvist and Kozen must add many primitives to their language to work around their linearity restrictions. For instance, in order to write projection functions $\mathbb{R}^n \rightarrow \mathbb{R}^m$, $n > m$ they must add projection primitives to the language.

² To avoid visually polluting the proofs we will drop the color code in Theorem 3 and Theorem 7

By having compositional type constructors that can represent joint distributions, i.e. $\mathcal{M}(\tau \times \tau)$, it is possible to write the program `sample t as x in ($\pi_1 x, \pi_3 x$)` which samples from a distribution over triples and returns only the first and third components by only using the syntax of products in MK.

Unfortunately there are some aspects of this language that still are restrictive. For instance, imagine that we want to write an LL program that receives two “Markov kernels” $\mathcal{M}\mathbb{N} \multimap \mathcal{M}\mathbb{N}$ and a distribution over \mathbb{N} as inputs, samples from the input distribution, feeds the result to the Markov kernels, samples from them and adds the results. Its type would be

$$(\mathcal{M}\mathbb{N} \multimap \mathcal{M}\mathbb{N}) \multimap (\mathcal{M}\mathbb{N} \multimap \mathcal{M}\mathbb{N}) \multimap \mathcal{M}\mathbb{N} \multimap \mathcal{M}\mathbb{N}$$

Even though the program only requires you to sample once from each distribution, it is still not possible to write it in the linear language.

We will show in Section 4 how the type constructor \mathcal{M} actually corresponds to an applicative functor [19], and the limitation above is actually a particular case of a fundamental difference between programming with applicative functors compared to programming with monads.

Remark 1. We now have two languages that can interpret probabilistic primitives such as `coin`. However, every primitive M in the MK language can be easily transported to an LL program by using an empty list of LL programs: `sample _ as _ in M`. Therefore it makes sense to only add these primitives to the MK language.

4 Categorical Semantics

As it is the case with categorical interpretations of languages/logics, types and contexts are interpreted as objects in a category and every well-typed program/proof gives rise to a morphism.

In our case, MK types τ are interpreted as objects $\llbracket \tau \rrbracket$ in a Markov category (\mathbf{M}, \times) and well-typed programs $\Gamma \vdash_{MK} M : \tau$ are interpreted as an \mathbf{M} morphism $\llbracket \Gamma \rrbracket \rightarrow \llbracket \tau \rrbracket$, as shown in Figure 6. Similarly, LL types $\underline{\tau}$ are interpreted as objects $\llbracket \underline{\tau} \rrbracket$ in a model of linear logic $(\mathbf{C}, \otimes, \multimap)$ and well-typed programs $\Gamma \vdash_{LL} t : \underline{\tau}$ are interpreted as a \mathbf{C} morphism $\llbracket \Gamma \rrbracket \rightarrow \llbracket \underline{\tau} \rrbracket$, as shown in Figure 7.

To give semantics to the combined language is not as straightforward. The `sample` rule allows the programmer to run LL programs, bind the results to MK variables and use said variables in an MK continuation. The implication of this rule in our formalism is that our semantics should provide a way of translating MK programs into LL programs. In category theory this is usually achieved by a functor $\mathcal{M} : \mathbf{M} \rightarrow \mathbf{C}$.

However, we can easily see that functors are not enough to interpret the `sample` rule. Consider what happens when you apply \mathcal{M} to an MK program $x : \tau_1, y : \tau_2 \vdash_{MK} N : \tau$:

$$\mathcal{M} \llbracket N \rrbracket : \mathcal{M}(\tau_1 \otimes \tau_2) \rightarrow \mathcal{M}\tau$$

To precompose it with two LL programs outputting $\mathcal{M}\tau_1$ and $\mathcal{M}\tau_2$ we need a mediating morphism $\mu_{\tau_1, \tau_2} : \mathcal{M}\tau_1 \otimes \mathcal{M}\tau_2 \rightarrow \mathcal{M}(\tau_1 \times \tau_2)$. Furthermore, if N has three or more free variables, there would be several ways of applying μ . Since from a programming standpoint it should not matter how the LL programs are associated, we require that μ_{τ_1, τ_2} makes the lax monoidality diagrams to commute. Therefore, assuming lax monoidality of μ we can interpret the sample rule:

$$\text{SAMPLE} \quad \frac{\tau_1 \times \cdots \times \tau_n \xrightarrow{N} \tau \quad \Gamma_i \xrightarrow{t_i} \mathcal{M}\tau_i}{\Gamma \xrightarrow{t_1 \otimes \cdots \otimes t_n} \mathcal{M}\tau_1 \otimes \cdots \otimes \mathcal{M}\tau_n \xrightarrow{\mu} \mathcal{M}(\tau_1 \times \cdots \times \tau_n) \xrightarrow{\mathcal{M}N} \mathcal{M}\tau}$$

In case it only has one MK variable, the semantics is given by $\llbracket t \rrbracket ; \mathcal{M} \llbracket N \rrbracket$ and in case it does not have any free variables the semantics is $\epsilon ; \mathcal{M} \llbracket N \rrbracket$.

The equational theory of the LL languages is the well-known theory of the simply-typed λ -calculus and the MK equational theory has been described, in graphical notation, by Fritz [12]. Something which is not obvious is understanding how they interact at their boundary. This is where \mathcal{M} being a functor becomes relevant, since from functoriality it follows the two program equivalences:

Theorem 4. *Let t , M and N be well-typed programs,*

$$\begin{aligned} & \llbracket (\lambda y. \text{sample } y \text{ as } z \text{ in } N) (\text{sample } t \text{ as } x \text{ in } M) \rrbracket = \\ & \llbracket \text{sample } t \text{ as } x \text{ in } (\text{let } y = M \text{ in } N) \rrbracket \end{aligned}$$

Proof.

$$\begin{aligned} & \llbracket (\lambda y. \text{sample } y \text{ as } z \text{ in } N) (\text{sample } t \text{ as } x \text{ in } M) \rrbracket = \\ & \llbracket t \rrbracket ; \mathcal{M} \llbracket M \rrbracket ; \mathcal{M} \llbracket N \rrbracket = \llbracket t \rrbracket ; \mathcal{M}(\llbracket M \rrbracket ; \llbracket N \rrbracket) = \\ & \llbracket \text{sample } t \text{ as } x \text{ in } (\text{let } y = M \text{ in } N) \rrbracket \end{aligned}$$

Theorem 5. *Let t be a well-typed program,*

$$\llbracket \text{sample } t \text{ as } x \text{ in } x \rrbracket = \llbracket t \rrbracket$$

Proof. $\llbracket \text{sample } t \text{ as } x \text{ in } x \rrbracket = \llbracket t \rrbracket ; \mathcal{M}(\llbracket x \rrbracket) = \llbracket t \rrbracket ; \mathcal{M}(id) = \llbracket t \rrbracket ; id = \llbracket t \rrbracket$

Furthermore, we also have a modularity property that can be easily proven:

Theorem 6. *Let t , M and N be well-typed programs. If $\llbracket M \rrbracket = \llbracket N \rrbracket$ then*

$$\llbracket \text{sample } t \text{ as } x \text{ in } M \rrbracket = \llbracket \text{sample } t \text{ as } x \text{ in } N \rrbracket$$

The expected compositionality of the semantics also holds:

Theorem 7. *Let $x_1 : \tau_1, \dots, x_n : \tau_n \vdash t : \tau$ and $\Gamma_i \vdash t_i : \tau_i$ be well-typed terms.*
 $\llbracket \Gamma_1, \dots, \Gamma_n \vdash t\{\vec{x}_i/\vec{t}_i\} : \tau \rrbracket = (\llbracket \Gamma_1 \vdash t_1 : \tau_1 \rrbracket \otimes \cdots \otimes \llbracket \Gamma_n \vdash t_n : \tau_n \rrbracket); \llbracket \Gamma_1, \dots, \Gamma_n \rrbracket \vdash t : \tau.$

Proof. The proof can be found in Appendix D.

$$\frac{\text{SUBST} \quad \Gamma \vdash u_1 : \tau' \quad \Gamma \vdash u_2 : \tau' \quad \Gamma, x : \tau' \vdash t : \tau \quad \Gamma \vdash u_1 \equiv u_2 : \tau'}{\Gamma \vdash t\{x/u_1\} \equiv t\{x/u_2\} : \tau}$$

From this theorem we can conclude:

Corollary 1. *The Subst rule shown above is sound with respect to the categorical semantics.*

Lax monoidal functors, under the name *applicative functors*, are widely used in programming languages research[19]. They are often used to define embedded domain-specific languages (eDSL) within a host language. This suggests that from a design perspective the Markov kernel language can be thought of as an eDSL inside a linear language.

We have just shown that \mathcal{M} being lax monoidal is sufficient to give semantics to our combined language, but what would happen if it had even more structure? If it were also full it would be possible to add a reification command³:

$$\frac{\mathcal{M}\Gamma \vdash_{LL} t : \mathcal{M}\tau}{\Gamma \vdash_{MK} \text{reify}(t) : \tau}$$

where $\mathcal{M}\Gamma$ is notation for every variable in Γ being of the form $\mathcal{M}\tau'$, for some τ' . The semantics for the rule would be taking the inverse image of \mathcal{M} . As we will show in the next section, there are some concrete models where \mathcal{M} is full and some other models where it is not. Computationally, fullness of \mathcal{M} can be interpreted as every program of type $\mathcal{M}\tau \multimap \mathcal{M}\tau'$ being equal to a Markov kernel.

A property which is easier to satisfy is faithfulness, which is verified by both models in the next section. In this case the translation of the MK language into the LL language would be fully-abstract in the following sense:

Theorem 8. *Let $x : \tau_1 \vdash M : \tau_2$ and $x : \tau_1 \vdash N : \tau_2$ be two well-typed MK programs. If \mathcal{M} is faithful then $\llbracket \text{sample } y \text{ as } x \text{ in } M \rrbracket = \llbracket \text{sample } y \text{ as } x \text{ in } N \rrbracket$ implies $\llbracket M \rrbracket = \llbracket N \rrbracket$.*

Proof. $\llbracket \text{sample } y \text{ as } x \text{ in } M \rrbracket = \llbracket \text{sample } y \text{ as } x \text{ in } N \rrbracket \implies id_{\mathcal{M}\tau_1}; \mathcal{M} \llbracket M \rrbracket = id_{\mathcal{M}\tau_1}; \mathcal{M} \llbracket N \rrbracket \implies \llbracket M \rrbracket = \llbracket N \rrbracket$.

5 Concrete Models

In this section we show how existing models for both discrete as well as continuous probabilities fit within our formalism.

³ The proposed rule breaks the substitution theorem, but it is possible to define a variant for it where this is not the case.

5.1 Discrete Probability

For the sake of simplicity we will denote the monoidal product of **CountStoch** as \times .

The probabilistic coherence space model of linear logic has been extensively studied in the context of semantics of discrete probabilistic languages[9].

Definition 10 (Probabilistic Coherence Spaces [9]). *A probabilistic coherence space (PCS) is a pair $(|X|, \mathcal{P}(X))$ where $|X|$ is a countable set and $\mathcal{P}(X) \subseteq |X| \rightarrow \mathbb{R}^+$ is a set, called the web, such that:*

- $\forall a \in X \exists \varepsilon_a > 0 \ \varepsilon_a \cdot \delta_a \in \mathcal{P}(X)$, where $\delta_a(a') = 1$ iff $a = a'$ and 0 otherwise, and we use the notation $\varepsilon_a = \varepsilon(a)$;
- $\forall a \in X \exists \lambda_a \forall x \in \mathcal{P}(X) \ x_a \leq \lambda_a$;
- $\mathcal{P}(X)^{\perp\perp} = \mathcal{P}(X)$, where $\mathcal{P}(X)^\perp = \{x \in X \rightarrow \mathbb{R}^+ \mid \forall v \in \mathcal{P}(X) \ \sum_{a \in X} x_a v_a \leq 1\}$.

We can define a category **PCoh** where objects are probabilistic coherence spaces and morphisms $X \multimap Y$ are matrices $f : |X| \times |Y| \rightarrow \mathbb{R}^+$ such that for every $v \in \mathcal{P}(X)$, $(fv) \in \mathcal{P}(Y)$, where $(fv)_b = \sum_{a \in |A|} f_{(a,b)} v_a$.

Definition 11. *Let $(|X|, \mathcal{P}(X))$ and $(|Y|, \mathcal{P}(Y))$ be PCS, we define $X \otimes Y = (|X| \times |Y|, \{x \otimes y \mid x \in \mathcal{P}(X), y \in \mathcal{P}(Y)\}^{\perp\perp})$, where $(x \otimes y)(a, b) = x(a)y(b)$*

Lemma 1. *Let X be a countable set, the pair $(X, \{\mu : X \rightarrow \mathbb{R}^+ \mid \sum_{x \in X} \mu(x) \leq 1\})$ is a PCS.*

Proof. The first two points are obvious, as the Dirac measure is a subprobability measure and every subprobability measure is bounded above by the constant function $\mu_1(x) = 1$.

To prove the last point we use the — easy to prove — fact that $\mathcal{P}X \subseteq \mathcal{P}X^{\perp\perp}$. Therefore we must only prove the other direction. First, observe that, if $\mu \in \{\mu : X \rightarrow \mathbb{R}^+ \mid \sum_{x \in X} \mu(x) \leq 1\}$, then we have $\sum \mu(x) \mu_1(x) = \sum 1\mu(x) = \sum \mu(x) \leq 1$, $\mu_1 \in \{\mu : X \rightarrow \mathbb{R}^+ \mid \sum_{x \in X} \mu(x) \leq 1\}^\perp$.

Let $\tilde{\mu} \in \{\mu : X \rightarrow \mathbb{R}^+ \mid \sum_{x \in X} \mu(x) \leq 1\}^{\perp\perp}$. By definition, $\sum \tilde{\mu}(x) = \sum \tilde{\mu}(x) \mu_1(x) \leq 1$ and, therefore, the third point holds.

This lemma can be used to give semantics to probabilistic primitives. For instance, a fair coin is interpreted as a function $\text{coin} : \mathbb{N} \rightarrow [0, 1]$ which is .5 at 0 and 1 and 0 elsewhere and is an element of $\mathcal{P}(\mathbb{N})$.

Lemma 2. *Let $X \rightarrow Y$ be a **CountStoch** morphism. It is also a **PCoh** morphism.*

Theorem 9. *There is a lax monoidal functor $\mathcal{M} : \mathbf{CountStoch} \rightarrow \mathbf{PCoh}$.*

Proof. The functor is defined using the lemmas above. Functoriality holds due to the functor being the identity on arrows. The lax monoidal structure is given by $\epsilon = id_1$ and $\mu_{X,Y} = id_{X \times Y}$

Lemma 3. *If $\mu \in \{x \otimes y \mid x \in \mathcal{M}(X), y \in \mathcal{M}(Y)\}^\perp$ then for every $x \in X$ and $y \in Y$, $\mu(x, y) \leq 1$.*

Proof. If there were such indices such that $\mu(x_1, y_1) > 1$ then $\sum \sum \mu(x, y)(\delta_{x_1} \otimes \delta_{y_1})(x, y) > \mu(x_1, y_1)(\delta_{x_1} \otimes \delta_{y_1})(x_1, y_1) = \mu(x_1, y_1) > 1$, which is a contradiction.

Lemma 4. *Let X and Y be two countable sets, then*

$$\mathcal{M}X \otimes \mathcal{M}Y = \left(X \times Y, \{ \mu : X \times Y \rightarrow \mathbb{R}^+ \mid \sum_{x \in X} \sum_{y \in Y} \mu(x, y) \leq 1 \} \right) = \mathcal{M}(X \times Y).$$

Proof. By the lemma above it follows that if we have a joint probability distribution $\tilde{\mu}$ over $X \times Y$ and an element $\mu \in \{x \otimes y \mid x \in \mathcal{M}(X), y \in \mathcal{M}(Y)\}^\perp$ then $\sum \sum \mu(x, y)\tilde{\mu}(x, y) \leq \sum \sum \tilde{\mu}(x, y) \leq 1$.

Theorem 10. *Both ϵ and $\mu_{X,Y}$ are isomorphisms.*

Proof. Since ϵ is the identity morphism, it is trivially an isomorphism. The morphisms $\mu_{X,Y}$ being an isomorphism is a direct consequence of the lemmas above.

Theorem 11. *The functor \mathcal{M} is full.*

Both results above can be directly used to enhance the syntax of the combined language. From Theorem 10 we can conclude that elements of type $\mathcal{M}(\tau_1 \times \tau_2)$, by projecting their marginal distributions, can be manipulated as if they had type $\mathcal{M}\tau_1 \otimes \mathcal{M}\tau_2$. Something to note is that when we do this marginalization process we lose potential correlations between the elements of the pair.

5.2 Continuous Probability

In order to accommodate continuous distributions we can use regularly ordered Banach spaces, whose detailed definition goes beyond the scope of this paper.

Definition 12 ([8]). *The category **RoBan** has regularly ordered Banach spaces as objects and regular linear functions as morphisms.*

Theorem 12. *There is a lax monoidal functor $\mathcal{M} : \mathbf{Kern} \rightarrow \mathbf{RoBan}$.*

Proof. The functor acts on objects by sending a measurable space to the set of signed measures over it, which can be equipped with a **RoBan** structure. On morphisms it sends a Markov kernel f to the linear function $\mathcal{M}(f)(\mu) = \int f d\mu$.

The monoidal structure of **RoBan** satisfies the universal property of tensor products and, therefore, we can define the natural transformation $\mu_{X,Y} : \mathcal{M}(X) \otimes \mathcal{M}(Y) \rightarrow \mathcal{M}(X \times Y)$ as the function generated by the bilinear function $\mathcal{M}(X); \mathcal{M}(Y) \multimap \mathcal{M}(X \times Y)$ which maps a pair of distributions to its product measure. The map ϵ is, once again, equal to the identity function.

The commutativity of the lax monoidality diagrams follows from the universal property of the tensor product: it suffices to verify it for elements $\mu_A \otimes \mu_B \otimes \mu_C$.

In **RoBan** the uniform distribution over the interval $[0, 1]$ is an element of $\mathcal{M}\mathbb{R}$, meaning that it can soundly interpret a $\cdot \vdash_{LL} \text{uniform} : \mathcal{M}\mathbb{R}$ primitive.

Even though \mathcal{M} looks very similar to the discrete case, it follows from a well-known theorem from functional analysis that the functor is *not* strong monoidal, meaning that there are joint probability distributions (elements of $\mathcal{M}(A \times B)$) that cannot be represented as an element of the tensor product $\mathcal{M}(A) \otimes \mathcal{M}(B)$ and, as such, programs of type $M(A \times B)$ must be manipulated in MK language, as shown in Example 3.

6 Beyond Probability

We have seen that this new resource interpretation is present in different models of linear logic models for probabilistic programming. In this section we show that this model can be generalized to commutative effects, i.e. effects where the program equation Commutativity below holds. Categorically, these effects are captured by monoidal monads⁴. Due to length issues, we will not fully detail the definition of monoidal monads, but we suggest the interested reader to read Seal [23].

$$\begin{array}{c} \text{COMMUTATIVITY} \\ \frac{\Gamma \vdash t_1 : \tau_1 \quad \Gamma \vdash t_2 : \tau_2 \quad \Gamma, x : \tau_1, y : \tau_2 \vdash u : \tau}{\text{let } x_1 = t_1 \text{ in (let } x_2 = t_2 \text{ in } u) \equiv \text{let } x_2 = t_2 \text{ in (let } x_1 = t_1 \text{ in } u) : \tau} \end{array}$$

Definition 13 ([23]). *Let (\mathbf{C}, \otimes, I) be a monoidal category and (T, η, μ) a monad over it. The monad T is called monoidal if it comes equipped with a natural transformation $\kappa_{X,Y} : TX \otimes TY \rightarrow T(X \otimes Y)$ making certain diagrams commute*

For probability monads the transformation κ corresponds to forming the product probability distribution and, more generally, this can be thought of a program that runs both of its (effectful) inputs and pairs the outputs.

Every monad give rise to the interesting categories \mathbf{C}_T and \mathbf{C}^T which are, respectively, the Kleisli category and Eilenberg-Moore category. The objects of \mathbf{C}_T are the same as \mathbf{C} and morphisms between A and B are \mathbf{C} morphisms $A \rightarrow TB$, with the identity morphism being equal to the unit η of the monad and composition is given by $f; g = f; Tg; \mu$.

The objects of the category \mathbf{C}^T are pairs (X, x) , where X is a \mathbf{C} object and $x : TX \rightarrow X$ is a \mathbf{C} morphism such that $\mu; x = Tx; x$ and $\eta; x = id_X$, and morphisms between objects (X, x) and (Y, y) are \mathbf{C} morphisms $f : X \rightarrow Y$ such that $x; f = Tf; y$.

For every monad T there is a canonical inclusion functor $\iota : \mathbf{C}_T \rightarrow \mathbf{C}^T$ which maps X to (TX, μ) and $f : X \rightarrow Y$ to $Tf; \mu_Y$.

Theorem 13 ([5]). *The functor ι is full and faithful.*

⁴ Monoidal monads are equivalent to commutative monads, which is the nomenclature usually used in the context of programming languages semantics.

As we explain in Appendix C, assuming enough structure on the category \mathbf{C} we can show that the triple $(\mathbf{C}_T, \mathbf{C}^T, \iota)$ is a model to the MK+LL language and we can bring our new resource interpretation of linear logic to other commutative effects.

An illustrative example is the powerset monad $\mathcal{P} : \mathbf{Set} \rightarrow \mathbf{Set}$ which is monoidal and since \mathbf{Set} has the necessary structure, the triple $(\mathbf{C}_{\mathcal{P}}, \mathbf{C}^{\mathcal{P}}, \mathcal{P})$ is a model to our language and can be used to give semantics to non-deterministic computation.

In the context of commutative effects other than randomness, the syntax sample t as x in M does not make as much sense, in which case we can use the syntax observe t_i as x_i in M instead. Once again, operationally, the programs t_i are fully executed, the values are bound to x_i in M which is then executed.

Furthermore, other effects have other relevant effectful operations and, therefore, we can assume that there is a set of operations in the MK language that are interpreted in the Kleisli category and can be transported to LL using observe, similar to how it was done in the probabilistic case.

For the non-deterministic case we can assume the existence of typing rules for non-deterministic choice and failure:

$$\begin{array}{c} \text{CHOICE} \\ \frac{\Gamma \vdash_{MK} t_1 : \tau \quad \Gamma \vdash_{MK} t_2 : \tau}{\Gamma \vdash_{MK} t_1 \oplus t_2 : \tau} \end{array} \qquad \begin{array}{c} \text{NULL} \\ \frac{}{\Gamma \vdash_{MK} 0_{\tau} : \tau} \end{array}$$

satisfying the expected equations and interpreted using set-theoretic union and the empty set, respectively.

A similar connection between linear logic and monoidal monads has been made by Benton and Wadler[4], where they want to relate Moggi's monadic λ -calculus with linear logic by showing that if a monad is monoidal and the category has equalizers and coequalizers, then the Eilenberg-Moore category is a model of linear logic.

7 Related Work

Semantics of Probabilistic Programming Ehrhard et al. [11,10] have defined a model of linear logic \mathbf{CLin} which can be used to interpret a higher-order probabilistic programming language. They have used the call-by-name translation of intuitionistic logic into linear logic $A \rightarrow B = !A \multimap B$ to give semantics to their language. The authors extend their language with a call-by-value let syntax which makes it possible to reuse sampled values. In order to give semantics to this new language they introduce a new category \mathbf{CLin}_m which can interpret this new operator, at the cost of complicating their model.

Because there is an analogous proof of Theorem 12 with the category \mathbf{CLin} replacing \mathbf{RoBan} , we can use their original, simpler, model to interpret our language, while not needing to use the linear logic exponential to interpret non-linear programs.

Dahlqvist and Kozen [8] have defined a category of partially ordered Banach spaces and shown that it is a model of intuitionistic linear logic. An important difference from their approach and the one mentioned above is that they embrace variable linearity as part of their syntax. As we argued in this paper, we believe that the syntactic restriction of linearity they have used is not adequate for the purposes of probabilistic programming. They deal with this limitation by adding primitives to their languages which, by using the results of Section 5, could be programmed using the MK language.

Quasi Borel spaces [15] are a conservative extension of **Meas** that are Cartesian closed and have a commutative probability monad. The drawback of this model is that it is still not as well understood as its measure-theoretic counterpart, and there are theorems from probability theory used to reason about programs that may not hold in the category of quasi Borel spaces **QBS**.

Recently, Geoffroy [13] has made progress in connecting linear logic and quasi Borel Spaces by showing that a certain subcategory of the Eilenberg-Moore category for the probability monad in **QBS** is a model of classical linear logic, which we see as an instance of our model where the MK language can have higher-order functions as well.

Call-by-Push-Value The idea of having two distinct type systems that are connected by a functorial layer is reminiscent of Call-by-Push-Value (CBPV) [17], which has a type system for values and a type system for computations that are connected by an adjunction. In recent work, Ehrhard and Tasson [24] use the Eilenberg-Moore adjunction of the linear logic exponential ! to give semantics to a calculus that can interpret lazy and eager probabilistic computation, allowing for the interpretation of an eager let operator which is operationally similar to our **sample** construct. However, the existence of the let operator depends on properties of the ! that are unknown to hold for continuous distributions, while our semantics can naturally deal with continuous distributions as we have shown in Section 5.

Furthermore, the exponential which lies at the center of their approach is, semantically, hard to work with and does not have any clear connections to probability theory, making it unlikely that their semantics can be seen as a bridge between the Markov and linear semantics, which is the case for the models presented in Section 5.

Goubault-Larrecq [14] has defined a CBPV domain semantics to a language that mixes probability and non-determinism, a long-standing challenge in the theory of programming languages. His focus is in understanding how to make probability interact with non-determinism in a sound way. He studies the full-abstraction of his semantics but does not deal with connections to linear logic.

Acknowledgements The support of the National Science Foundation under grant CCF-2008083 is gratefully acknowledged. I would also like to thank Arthur Azevedo de Amorim, Justin Hsu, Michael Roberts, Christopher Lam and Deepak Garg for their useful comments on earlier versions of this paper.

A Typing Rules and Denotational Semantics LL and MK

$$\begin{array}{c}
 \text{VAR} \\
 \hline
 \Gamma, x : \tau \vdash x : \tau \\
 \\
 \text{LET} \\
 \hline
 \frac{\Gamma \vdash M : \tau_1 \quad \Gamma, x : \tau_1 \vdash N : \tau}{\Gamma \vdash \text{let } x = M \text{ in } N : \tau} \\
 \\
 \text{PAIR} \\
 \hline
 \frac{\Gamma \vdash M : \tau_1 \quad \Gamma \vdash N : \tau_2}{\Gamma \vdash (M, N) : \tau_1 \times \tau_2} \\
 \\
 \text{PRIMITIVE} \\
 \hline
 \frac{\Gamma \vdash M : \tau_1 \quad f : \tau_1 \rightarrow \tau_2}{\Gamma \vdash f(M) : \tau_2} \\
 \\
 \text{PROJ1} \\
 \hline
 \frac{\Gamma \vdash M : \tau_1 \times \tau_2}{\Gamma \vdash \pi_1 M : \tau_1} \\
 \\
 \text{PROJ2} \\
 \hline
 \frac{\Gamma \vdash M : \tau_1 \times \tau_2}{\Gamma \vdash \pi_2 M : \tau_2}
 \end{array}$$

Fig. 4: Typing rules MK

$$\begin{array}{c}
 \text{AXIOM} \\
 \hline
 x : \tau \vdash x : \tau \\
 \\
 \text{UNIT} \\
 \hline
 \cdot \vdash \text{unit} : 1 \\
 \\
 \text{ABSTRACTION} \\
 \hline
 \frac{\Gamma, x : \tau_1 \vdash t : \tau_2}{\Gamma \vdash \lambda x. t : \tau_1 \multimap \tau_2} \\
 \\
 \text{APPLICATION} \\
 \hline
 \frac{\Gamma_1 \vdash t : \tau_1 \multimap \tau_2 \quad \Gamma_2 \vdash u : \tau_1}{\Gamma_1, \Gamma_2 \vdash t u : \tau_2} \\
 \\
 \text{TENSOR} \\
 \hline
 \frac{\Gamma_1 \vdash t : \tau_1 \quad \Gamma_2 \vdash u : \tau_2}{\Gamma_1, \Gamma_2 \vdash t \otimes u : \tau_1 \otimes \tau_2} \\
 \\
 \text{LET TENSOR} \\
 \hline
 \frac{\Gamma_1 \vdash t : \tau_1 \otimes \tau_2 \quad \Gamma_2, x : \tau_1, y : \tau_2 \vdash u : \tau}{\Gamma_1, \Gamma_2 \vdash \text{let } x \otimes y = t \text{ in } u : \tau}
 \end{array}$$

Fig. 5: Typing rules LL

$$\begin{array}{c}
 \text{VAR} \\
 \hline
 \frac{}{\Gamma \times \tau \xrightarrow{\text{delete} \times \text{id}_\tau} 1 \times \tau \cong \tau} \\
 \\
 \text{PAIR} \\
 \hline
 \frac{\Gamma \xrightarrow{M} \tau_1 \quad \Gamma \xrightarrow{N} \tau_2}{\Gamma \xrightarrow{\text{copy}} \Gamma \times \Gamma \xrightarrow{M \times N} \tau_1 \times \tau_2} \\
 \\
 \text{PROJ} \\
 \hline
 \frac{\Gamma \xrightarrow{M} \tau_1 \times \tau_2}{\Gamma \xrightarrow{M; (\text{id}_{\tau_1} \times \text{delete})} \tau_1 \times 1 \cong \tau_1} \\
 \\
 \text{LET} \\
 \hline
 \frac{\Gamma \xrightarrow{M} \tau_1 \quad \Gamma \times \tau_1 \xrightarrow{N} \tau}{\Gamma \xrightarrow{\text{copy}} \Gamma \times \Gamma \xrightarrow{(\text{id}_\Gamma \times M); N} \tau} \\
 \\
 \text{PRIMITIVE} \\
 \hline
 \frac{\Gamma \xrightarrow{M} \tau_1 \quad \tau_1 \xrightarrow{f} \tau_2}{\Gamma \xrightarrow{M} \tau_1 \xrightarrow{f} \tau_2}
 \end{array}$$

Fig. 6: Denotational semantics for MK

$$\begin{array}{c}
\text{AXIOM} \\
\frac{}{\tau \xrightarrow{id_\tau} \tau}
\end{array}
\qquad
\begin{array}{c}
\text{TENSOR} \\
\frac{\Gamma_1 \xrightarrow{t_1} \underline{\tau_1} \quad \Gamma_2 \xrightarrow{t_2} \underline{\tau_2}}{\Gamma_1, \Gamma_2 \xrightarrow{t_1 \otimes t_2} \underline{\tau_1} \otimes \underline{\tau_2}}
\end{array}$$

$$\begin{array}{c}
\text{LET TENSOR} \\
\frac{\Gamma_1 \xrightarrow{t} \underline{\tau_1} \otimes \underline{\tau_2} \quad \Gamma_2 \otimes \underline{\tau_1} \otimes \underline{\tau_2} \xrightarrow{u} \underline{\tau}}{\Gamma_1 \otimes \Gamma_2 \xrightarrow{(id \otimes t); u} \underline{\tau}}
\end{array}
\qquad
\begin{array}{c}
\text{ABSTRACTION} \\
\frac{\Gamma \otimes \underline{\tau_1} \xrightarrow{t} \underline{\tau_2}}{\Gamma \xrightarrow{cur([\![t]\!]])} \underline{\tau_1} \multimap \underline{\tau_2}}
\end{array}$$

$$\begin{array}{c}
\text{APPLICATION} \\
\frac{\Gamma_1 \xrightarrow{t} \underline{\tau_1} \multimap \underline{\tau_2} \quad \Gamma_2 \xrightarrow{u} \underline{\tau_1}}{\Gamma_1 \otimes \Gamma_2 \xrightarrow{(t \otimes u); ev} \underline{\tau_2}}
\end{array}$$

Fig. 7: Denotational semantics for LL

B Commutative Diagrams

$$\begin{array}{ccc}
(F(X) \otimes_{\mathbf{D}} F(Y)) \otimes_{\mathbf{D}} F(Z) & \xrightarrow{\alpha} & F(X) \otimes_{\mathbf{D}} (F(Y) \otimes_{\mathbf{C}} F(Z)) \\
\downarrow \mu \otimes id & & \downarrow id \otimes \mu \\
F(X \otimes_{\mathbf{C}} Y) \otimes_{\mathbf{D}} F(Z) & & F(X) \otimes_{\mathbf{D}} F(Y \otimes_{\mathbf{C}} Z) \\
\downarrow \mu & & \downarrow \mu \\
F((X \otimes_{\mathbf{C}} Y) \otimes_{\mathbf{C}} Z) & \xrightarrow{F\alpha} & F(X \otimes_{\mathbf{C}} (Y \otimes_{\mathbf{C}} Z))
\end{array}$$

$$\begin{array}{ccc}
1 \otimes_{\mathbf{D}} F(X) & \xrightarrow{\epsilon \otimes id} & F(1) \otimes_{\mathbf{D}} F(X) \\
\downarrow \iota_{\mathbf{D}} & & \downarrow \mu \\
F(X) & \xleftarrow{F(\iota_{\mathbf{C}})} & F(1 \otimes_{\mathbf{C}} X)
\end{array}
\qquad
\begin{array}{ccc}
F(X) \otimes_{\mathbf{D}} 1 & \xrightarrow{id \otimes \epsilon} & F(X) \otimes_{\mathbf{D}} F(1) \\
\downarrow r_{\mathbf{D}} & & \downarrow \mu \\
F(X) & \xleftarrow{F(r_{\mathbf{C}})} & F(X \otimes_{\mathbf{C}} 1)
\end{array}$$

Fig. 8: Lax monoidal diagrams

C Monoidal Monads and Their Algebras

An important theorem from the categorical probability literature is that Markov categories are an abstraction of programming in the Kleisli category of monoidal affine monads, where affinity means that $T1 \cong 1$.

Theorem 14 ([12]). *Let $(\mathbf{C}, \times, 1)$ be a cartesian category and $T : \mathbf{C} \rightarrow \mathbf{C}$ a monoidal (affine) monad. The Kleisli category \mathbf{C}_T is a Markov category.*

The monoidal product of \mathbf{C}_T is \times with unit 1, the copy operation is given by $\Delta_X; \eta_X : X \rightarrow T(X \times X)$ and the deletion operation is given by $T1 \cong 1$ and 1 being terminal.

Furthermore, under certain conditions, the Eilenberg-Moore category \mathbf{C}^T for monoidal monads is symmetric monoidal closed. The monoidal unit is given by TI , the monoidal product is given by the coequalizer depicted in Figure 9 and the closed structure is given by the equalizer depicted in Figure 10.

Theorem 15. *Let \mathbf{C} be a symmetric monoidal closed category with equalizers, reflexive co-equalizers and $T : \mathbf{C} \rightarrow \mathbf{C}$ a monoidal monad. The category \mathbf{C}^T is also symmetric monoidal closed.*

$$\begin{array}{ccccccc}
 T(TX \otimes TY) & \xrightarrow{T\kappa} & TT(X \otimes Y) & \xrightarrow{\mu} & T(X \otimes Y) & \longrightarrow & X \otimes_T Y \\
 & & \searrow & & \nearrow & & \\
 & & & & T(x \otimes y) & &
 \end{array}$$

Fig. 9: Symmetric Monoidal Structure in \mathbf{C}^T

$$\begin{array}{ccccccc}
 X \multimap_T Y & \longrightarrow & X \multimap Y & \xrightarrow{s} & TX \multimap TY & \xrightarrow{id_{TX} \multimap y} & TX \multimap Y \\
 & & \searrow & & \nearrow & & \\
 & & & & x \multimap id_Y & &
 \end{array}$$

Fig. 10: Closed Structure in \mathbf{C}^T

Even though, in general, in order to define the monoidal product one requires a coequalizer, for our purposes we are only interested in products of the form $TA \otimes_T TB$ which, luckily, are easier to characterize, since the equality $TX \otimes_T TY = T(X \otimes Y)$ holds [23].

In this case the lax monoidal transformations $\mu_{X,Y} : TX \otimes_T TY \rightarrow T(X \otimes Y)$ and $\epsilon : FI \rightarrow FI$ are simply the identity morphisms. Besides, by using the universal properties of coequalizers it is possible to show the equality $\tilde{\alpha}_{TX,TY,TZ} = \alpha_{X,Y,Z}$, where $\tilde{\alpha}$ is the associator for the monoidal product \otimes_T .

Theorem 16. *Let \mathbf{C} be a symmetric monoidal category with reflexive co-equalizers and $T : \mathbf{C} \rightarrow \mathbf{C}$ a monoidal monad. The triple (ι, μ, ϵ) is a lax monoidal functor.*

Proof. The proof follows by unfolding the definitions.

D Proofs

Theorem 3. Let $\Gamma, x : \underline{\tau}_1 \vdash t : \underline{\tau}$ and $\Delta \vdash u : \underline{\tau}_1$ be well-typed terms, then $\Gamma, \Delta \vdash t\{x/u\} : \underline{\tau}$

Proof. The proof follows by structural induction on the typing derivation $\Gamma, x : \underline{\tau}_1 \vdash t : \underline{\tau}$:

- Axiom: Since $t = x$ then $t\{x/u\} = u$ and $\underline{\tau}_1 = \underline{\tau}$.
- Abstraction: By hypothesis, $\Gamma, x : \underline{\tau}_1, y : \underline{\tau}_2 \vdash t : \underline{\tau}_3$. Since we can assume wlog that $x \neq y$ and that $y \notin \Delta$, $\lambda y. t\{x/u\} = \lambda y. t\{x/u\}$. Therefore we can show that $\Gamma, \Delta \vdash \lambda y. t\{x/u\} : \tau_2 \multimap \tau_3$ by applying the rule Abstraction and by the induction hypothesis.
- Application: $t_1 t_2\{x/u\} = t_1\{x/u\} t_2\{x/u\}$. Since the language LL is linear, only one of t_1 or t_2 will have x as a free variable. By symmetry we can assume that t_1 has x as a free variable and we can prove $\Gamma, \Delta \vdash t_1\{x/u\} t_2 : \underline{\tau}$ by applying the rule Application and by the induction hypothesis.
- Sample: It is easy to prove that $(\text{sample } t \text{ as } y \text{ in } M)\{x/u\} = \text{sample } (t\{x/u\}) \text{ as } y \text{ in } M$

Theorem 7. Let $x_1 : \tau_1, \dots, x_n : \tau_n \vdash t : \tau$ and $\Gamma_i \vdash t_i : \tau_i$ be well-typed terms.
 $\llbracket \Gamma_1, \dots, \Gamma_n \vdash t\{\vec{x}_i/\vec{t}_i\} : \underline{\tau} \rrbracket = (\llbracket \Gamma_1 \vdash t_1 : \underline{\tau}_1 \rrbracket \otimes \dots \otimes \llbracket \Gamma_n \vdash t_n : \underline{\tau}_n \rrbracket \rrbracket; \llbracket \Gamma_1, \dots, \Gamma_n \rrbracket \vdash t : \underline{\tau}.$

Proof. The proof follows by induction on the typing derivation of t .

- Axiom: Since $t = x$ then $t\{x/t_0\} = t_0$ and $\llbracket t\{x/t_0\} \rrbracket = \llbracket t_0 \rrbracket = \llbracket t_0 \rrbracket; id = \llbracket t_0 \rrbracket; \llbracket x \rrbracket.$
- Unit: Since $t = x$ then $t\{x/t_0\} = t_0$ and $\llbracket t\{x/t_0\} \rrbracket = \llbracket t_0 \rrbracket = \llbracket t_0 \rrbracket; id = \llbracket t_0 \rrbracket; \llbracket x \rrbracket.$
- Tensor: We know that $t = t_1 \otimes t_2$. Furthermore, from linearity we know that each free variable appears either in t_1 or in t_2 . Without loss of generality we can assume that $(t_1 \otimes t_2)\{x_1, \dots, x_n/u_1, \dots, u_n\} = (t_1\{x_1, \dots, x_k/u_1, \dots, u_k\}) \otimes (t_2\{x_{k+1}, \dots, x_n/u_{k+1}, \dots, u_n\})$. We can conclude this case from the induction hypothesis and functoriality of \otimes .
- LetTensor: This case follows from the functoriality of \otimes and the induction hypothesis.
- Abstraction: This case follows from unfolding the definitions, using the induction hypothesis and by naturality of cur .
- Application: Analogous to the Tensor case
- Sample: This case is analogous to the Tensor case.

References

1. de Amorim, A.A., Gaboardi, M., Hsu, J., Katsumata, S.y.: Probabilistic relational reasoning via metrics. In: Symposium on Logic in Computer Science (LICS) (2019)

2. Aumann, R.J.: Borel structures for function spaces. *Illinois Journal of Mathematics* (1961)
3. Barthe, G., Fournet, C., Grégoire, B., Strub, P.Y., Swamy, N., Zanella-Béguelin, S.: Probabilistic relational verification for cryptographic implementations. In: *Principles of Programming Languages (POPL)* (2014)
4. Benton, N., Wadler, P.: Linear logic, monads and the lambda calculus. In: *Symposium on Logic in Computer Science (LICS)* (1996)
5. Borceux, F.: *Handbook of Categorical Algebra: Volume 2, Categories and Structures*, vol. 2. Cambridge University Press (1994)
6. Chakarov, A., Sankaranarayanan, S.: Probabilistic program analysis with martingales. In: *International Conference on Computer Aided Verification (CAV)* (2013)
7. Clerc, F., Danos, V., Dahlqvist, F., Garnier, I.: Pointless learning. In: *International Conference on Foundations of Software Science and Computation Structures (FoSSaCS)* (2017)
8. Dahlqvist, F., Kozen, D.: Semantics of higher-order probabilistic programs with conditioning. In: *Principles of Programming Languages (POPL)* (2019)
9. Danos, V., Ehrhard, T.: Probabilistic coherence spaces as a model of higher-order probabilistic computation. *Information and Computation* **209**(6), 966–991 (2011)
10. Ehrhard, T.: On the linear structure of cones. In: *Logic in Computer Science (LICS)* (2020)
11. Ehrhard, T., Pagani, M., Tasson, C.: Measurable cones and stable, measurable functions: a model for probabilistic higher-order programming. In: *Principles of Programming Languages (POPL)* (2017)
12. Fritz, T.: A synthetic approach to markov kernels, conditional independence and theorems on sufficient statistics. *Advances in Mathematics* **370**, 107239 (2020)
13. Geoffroy, G.: Extensional denotational semantics of higher-order probabilistic programs, beyond the discrete case (unpublished) (2021)
14. Goubault-Larrecq, J.: A probabilistic and non-deterministic call-by-push-value language. In: *Logic in Computer Science (LICS)* (2019)
15. Heunen, C., Kammar, O., Staton, S., Yang, H.: A convenient category for higher-order probability theory. In: *Logic in Computer Science (LICS)* (2017)
16. Huang, M., Fu, H., Chatterjee, K., Goharshady, A.K.: Modular verification for almost-sure termination of probabilistic programs. *Proceedings of the ACM on Programming Languages (OOPSLA)* (2019)
17. Levy, P.B.: Call-by-push-value. Ph.D. thesis (2001)
18. Maraist, J., Odersky, M., Turner, D.N., Wadler, P.: Call-by-name, call-by-value, call-by-need and the linear lambda calculus. *Theoretical Computer Science* (1999)
19. McBride, C., Paterson, R.: Applicative programming with effects. *Journal of functional programming* **18**(1), 1–13 (2008)
20. McIver, A., Morgan, C., Kaminski, B.L., Katoen, J.P.: A new proof rule for almost-sure termination. *Proceedings of the ACM on Programming Languages (POPL)* (2017)
21. Mellies, P.A.: Categorical semantics of linear logic. *Panoramas et synthèses* **27**, 15–215 (2009)
22. Scibior, A., Kammar, O., Vakar, M., Staton, S., Yang, H., Cai, Y., Ostermann, K., Moss, S., Heunen, C., Ghahramani, Z.: Denotational validation of higher-order bayesian inference. *Proceedings of the ACM on Programming Languages* (2018)
23. Seal, G.J.: Tensors, monads and actions. arXiv preprint arXiv:1205.0101 (2012)
24. Tasson, C., Ehrhard, T.: Probabilistic call by push value. *Logical Methods in Computer Science* (2019)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

