

Simple Types for Probabilistic Termination

Willem Heijltjes  

Department of Computer Science, University of Bath, UK

Georgina Majury  

Department of Computer Science, University of Bath, UK

Abstract

We present a new typing discipline to guarantee the probability of termination in probabilistic lambda-calculi. The main contribution is a particular naturality and simplicity: our probabilistic types are as simple types, but generated from probabilities as base types, representing a least probability of termination. Simple types are recovered by restricting probabilities to one.

Our vehicle is the Probabilistic Event Lambda-Calculus by Dal Lago, Guerrieri, and Heijltjes, which presents a solution to the issue of confluence in probabilistic lambda-calculi. Our probabilistic type system provides an alternative solution to that using counting quantifiers by Antonelli, Dal Lago, and Pistone, for the same calculus.

The problem that both type systems address is to give a lower bound on the probability that terms head-normalize. Following the recent Functional Machine Calculus by Heijltjes, our development takes the (simplified) Krivine machine as primary, and proceeds via an extension of the calculus with sequential composition and identity on the machine. Our type system then gives a natural account of termination probability on the Krivine machine, reflected back onto head-normalization for the original calculus. In this way we are able to avoid the use of counting quantifiers, while improving on the termination bounds given by Antonelli, Dal Lago, and Pistone.

2012 ACM Subject Classification Theory of computation → Lambda calculus; Theory of computation → Type theory; Theory of computation → Probabilistic computation

Keywords and phrases lambda-calculus, probabilistic termination, simple types

Digital Object Identifier 10.4230/LIPIcs.CSL.2025.31

Acknowledgements We would like to thank Melissa Antonelli, Ugo Dal Lago, and Paolo Pistone for the constructive conversation about both our approaches, and the anonymous referees for their helpful commentary.

1 Introduction

While the study of probabilistic computation can be traced to the 1950s [11], the first study of the probabilistic λ -calculus in particular is considered to be by Saheb-Djahromi in the late 70s [30]. In the near half century since, many variations on higher-order probabilistic computation have been considered [10, 12, 19, 20, 26, 29]. In recent years, perhaps due to the potential for applications in machine learning and modelling of probabilistic systems, the area has seen a return to popularity [5, 6, 16, 17, 24, 31]. An important computational phenomenon in its own right, the study of probabilistic choice can also provide a “foot in the door” for understanding how more general effects might manifest, leading for instance to the recent *Functional Machine Calculus* (FMC) as a confluent λ -calculus with effects [3, 18] that will play a central role in our development.

In this paper we consider the problem of *probabilistic termination*, the probability that a given reduction mechanism reaches the normal form of a term, a key consideration for probabilistic computation and the subject of significant recent attention [4, 7, 8, 15, 22].¹

¹ Note that this objective is distinct from *almost-sure* termination: it considers *exact* probabilities, not those converging in the limit. Almost-sure termination is more commonly studied for iterative constructs [21, 27]; extending the lambda-calculus with an almost-surely terminating iterator is the subject of future work.



Termination being one of the prime considerations of type systems in general, the question of type systems for probabilistic termination is pertinent [1, 7]. Our contribution in this paper is a new typing discipline for probabilistic termination, following the recent line of work on the *probabilistic event λ -calculus* [9] and the FMC. Their prime features are *confluence* for probabilistic computation and other effects, and the encoding of multiple reduction strategies within a single calculus.

One of the greatest challenges facing the study of probabilistic computation is confluence. In most variants the outcome of duplicating a probabilistic term is strategy dependent. Consider the instruction “write down the result of flipping a coin twice”: it is ambiguous whether “twice” refers to “write” or “flipping”, and the choice of interpretation changes the possible outcomes. In most of the literature results are therefore either restricted to a single strategy, such as call-by-value or call-by-name, or rely on a modified definition of reduction. This raises the question whether it is possible to have confluence for probabilistic computation, while expressing different strategies, and rewriting without restriction on contexts, properties which lend elegance to the λ -calculus. The *probabilistic event λ -calculus* [9] proposes a solution to this problem by decomposing a probabilistic sum $M \oplus N$ into a *choice* operation $M a N$, understood as a conditional “if a then M else N ”, and a probabilistic *generator* $\boxed{a}.P$, representing a coin toss whose outcome is bound to the boolean variable a in the term P , which then determines the choice for $M a N$. When in argument position $\boxed{a}.M$ acts as a wrapper, similar to *thunking* in call-by-push value [25] and the bang-calculus [14, 15], protecting the operator from evaluation. These factors combine to return confluence to the calculus, allowing for arbitrary evaluation strategies with an unrestricted β -reduction.

The probabilistic event λ -calculus (PEA) was introduced in [9] with a simple type system, corresponding to that on the lambda-calculus. This system ignored the probabilistic elements of the calculus, and thus gave little insight into the properties of this extension. In a deterministic calculus a type system provides various safety guarantees. These may be qualitative: termination, outputs, composability; or quantitative: run time, term size. In the probabilistic setting, however, it is natural to wonder what guarantees can be made, when the behaviour of a single term can vary between iterations. If, as in the type system mentioned, the probabilities are ignored, strong results can be obtained, albeit on a fairly uninteresting fragment of the calculus. Once probability is acknowledged by the type system the guarantees made become similarly qualified: probability of termination, almost-sure termination, expected run time. Here, we consider the first of these.

One proposed type system for the PEA [1], labelled $C\lambda_{\rightarrow}$, introduces *counting quantifiers* as a Curry–Howard correspondent to an intuitionistic counting propositional logic. These provide a mechanism to express “proportion of truth” by quantifying the number of satisfying assignments to a formula within a given model, in the way that existential and universal quantification may be understood via the existence of a satisfying assignment, respectively the satisfaction of all assignments. By taking the assignments to a formula to describe the branches of a probabilistic computation, counting quantifiers may be used to describe probability bounds. By this mechanism, $C\lambda_{\rightarrow}$ provides a lower bound on the probability of head normalisation. However, as illustrated in [1], the bounds provided by $C\lambda_{\rightarrow}$ are not tight, even in situations where this would be expected (see Example 8).

In this paper we present an alternative approach to the same problem, for the same calculus, using a different inspiration, to provide improved termination bounds. Deriving from the PEA, the FMC provides an additional feature that, to us, seemed crucial to a natural account of probabilistic termination via the type system. First, the “machine” in question is the (simplified) *Krivine machine* [23], whose evaluation is closely related to head

normalisation. The FMC here adds an intriguing aspect: it introduces *sequential composition* and *identity* on the machine, where the latter, the imperative *skip*, provides a notion of *successful termination*, notably absent from the machine for standard λ -calculus. Moreover, this becomes the primary interpretation of types: a type derivation in the FMC is a proof that the machine successfully terminates. Our main question for this work was how to adapt this to capture *probabilistic termination*. The answer, described in Section 6, is an extended *sequential probabilistic event λ -calculus* $\text{SPE}\Lambda$, with a type system $\text{SPE}\Lambda^{\mathbb{Q}\Rightarrow}$ that naturally describes probabilistic termination of the machine, as well as probabilistic termination of head reduction.

The type system $\text{SPE}\Lambda^{\mathbb{Q}\Rightarrow}$ for the extended sequential calculus reflects back onto a natural type system for the original probabilistic event λ -calculus, $\text{PE}\Lambda^{\mathbb{Q}\rightarrow}$, which gives probabilistic head normalization in the following way: types generalize simple types by replacing the base type with a *probability*. Formally,

$$A, B ::= p \mid A \rightarrow B$$

where $p \in [0, 1] \cap \mathbb{Q}$ is a rational number between zero and one inclusive. The intuition is that the base type for simple types, o , may be understood as signifying successful evaluation, even if it is uninhabited. After all, a constant base type such as for booleans or integers signifies the successful return of a corresponding value. This further matches the interpretation in the FMC, where the type o is inhabited by *skip*, and corresponds to termination of the machine without producing a result. Our approach, then, is to replace *certain* termination with a *probability* of termination, replacing the base type o with a probability p . We consider the striking simplicity and natural intuition of this approach one of our main contributions.

2 The probabilistic event lambda-calculus

We recall the *probabilistic event lambda-calculus* $\text{PE}\Lambda$ from [9]. Assume countable sets of *variables*, ranged over by x, y, z , and of *events*, ranged over by a, b, c . The former are *term* variables, to be instantiated by terms of the calculus, and the latter are *boolean* variables, to be instantiated by \top (*true*) or \perp (*false*). The $\text{PE}\Lambda$ extends the λ -calculus with a *generator* $\boxed{a}.M$, which flips a coin and binds the result (\top or \perp) to a , and a *choice* or *conditional* $M a N$, which evaluates to M if a is *true*, and to N otherwise. A traditional probabilistic sum of terms, $M \oplus N$, may be encoded as $\boxed{a}.M a N$. Generators are normally *fair*, with equal probability for \top or \perp , though on occasion we may need a *biased* generator \boxed{a}_p which chooses \top with probability p , and \perp with $1 - p$.

► **Definition 1.** Terms are given by the following grammar,

$$M, N ::= x \mid \lambda x. M \mid M N \mid \boxed{a}. M \mid M a N$$

with, from left to right: a variable; an abstraction, which binds x in M ; an application; a generator, which binds a in M ; and a choice.

The *free variables* and *free events* of a term M are written $\text{fv}(M)$ and $\text{fe}(M)$ respectively. Substitution is written prefix: $\{N/x\}M$ is the capture-avoiding substitution of N for x in M . For an event a we define two *projection* functions π_{\top}^a and π_{\perp}^a , which apply the effect of instantiating a with *true* and *false* respectively, to a term M .

31:4 Simple Types for Probabilistic Termination

► **Definition 2.** The projection functions π_i^a for an event a and $i \in \{\perp, \top\}$ are given as follows, where $a \neq b$.

$$\begin{aligned} \pi_i^a x &= x & \pi_\top^a(M a N) &= \pi_\top^a M & \pi_i^a(M b N) &= (\pi_i^a M) b (\pi_i^a N) \\ \pi_i^a(\lambda x. M) &= \lambda x. \pi_i^a M & \pi_\perp^a(M a N) &= \pi_\perp^a N & \pi_i^a(\boxed{b}. M) &= \boxed{b}. \pi_i^a M \\ \pi_i^a(M N) &= (\pi_i^a M) (\pi_i^a N) \end{aligned}$$

We use the standard notions of *context*, *head context*, and *applicative context* to define the reduction relations. Note that a head context is of the form $\lambda x_1 \dots \lambda x_n. \{\} M_1 \dots M_m$ where m and n are potentially zero.

► **Definition 3.** Contexts C , head contexts H , and applicative contexts A are defined as follows. A context C with the hole replaced by M , capturing variables, is written $C\{M\}$.

$$\begin{aligned} C ::= \{\} & \mid \lambda x. C \mid C M \mid M C & H ::= \lambda x. H & \mid A \\ & \mid \boxed{a}. C \mid C a M \mid M a C & A ::= A M & \mid \{\} \end{aligned}$$

A *probability* p, q is a rational number between 0 and 1 inclusive. Probabilistic reduction will return a *multi-(sub-)distribution* [2], a finite multiset of weighted terms \mathcal{M} written as $[p_1 \cdot M_1, \dots, p_n \cdot M_n]$ whose *weight* $|\mathcal{M}| = \sum_{i \leq n} p_i$ is (at most) one.² For simplicity, we will refer to these as *distributions*, and we convert implicitly between terms M and the singleton distribution $[1 \cdot M]$. Multiset union is written $\mathcal{S} + \mathcal{T}$, and the empty multiset as \emptyset . A distribution \mathcal{M} is scaled to $p\mathcal{M}$ by multiplying each weight in \mathcal{M} by p . The underlying probability (sub-)distribution of \mathcal{M} , the finite function from terms to probabilities obtained by collecting like terms $p \cdot M$ and $q \cdot M$ as $(p+q) \cdot M$, is written $[\mathcal{M}]$.

► **Definition 4.** Beta-reduction \rightarrow_β and head β -reduction \rightarrow_{β_h} are given by closing the beta-rule below under all contexts C respectively under head contexts H , and implicitly return a singleton distribution.

$$(\lambda x. M)N \rightarrow_\beta \{N/x\}M$$

Projective reduction \rightarrow_π is the following reduction relation from terms to distributions.

$$H\{\boxed{a}. M\} \rightarrow_\pi [\frac{1}{2} \cdot H\{\pi_\top^a M\}, \frac{1}{2} \cdot H\{\pi_\perp^a M\}]$$

Head reduction $(\rightarrow_h) = (\rightarrow_{\beta_h}) \cup (\rightarrow_\pi)$ is the union of head β -reduction and projective reduction. Reduction is lifted to distributions of terms in the expected way: if $M \rightarrow N$ then $[p \cdot M] + \mathcal{M} \rightarrow [p \cdot N] + \mathcal{M}$. We write \twoheadrightarrow for the reflexive-transitive closure of a reduction relation \rightarrow .

The PEA features a second notion of reduction, *permutative reduction* \rightarrow_p [9], which gives a more fine-grained evaluation of probabilistic sums. It is this reduction for which confluence is particularly significant. The effect of permutative reduction is to bridge the gap between the decomposed operators \boxed{a} and $M a N$ and the standard probabilistic sum $M \oplus N$, encoded as $\boxed{a}. M a N$, by internalizing the reduction of $\boxed{a}. M$ to the sum of its two projections [9, Proposition 29]:

$$\boxed{a}. M \twoheadrightarrow_p \boxed{a}. (\pi_\top^a M) a (\pi_\perp^a M) \quad (\text{if } a \in \text{fe}(M)) .$$

² We use multi-distributions rather than distributions to accommodate the reduction measure for the proof of head normalization in Appendix A.2.

$$\begin{array}{c}
\frac{}{\Gamma, x: A \vdash x: A} \quad \frac{\Gamma, x: A \vdash M: B}{\Gamma \vdash \lambda x. M: A \rightarrow B} \quad \frac{\Gamma \vdash M: A \rightarrow B \quad \Gamma \vdash N: A}{\Gamma \vdash MN: B} \\
\frac{\Gamma \vdash M: A \quad \Gamma \vdash N: A}{\Gamma \vdash MaN: A} \quad \frac{\Gamma \vdash M: A}{\Gamma \vdash \boxed{a}. M: A}
\end{array}$$

■ **Figure 1** The simple type system PEA^\rightarrow .

In this paper we will work with projective reduction, since we are interested in termination of head reduction. We will, on occasion, consider terms where probabilistic sums are of the traditional form $M \oplus N = \boxed{a}. MaN$ with $a \notin \text{fe}(M), \text{fe}(N)$. As per the above, these are effectively the normal forms of permutative reduction.

3 Probabilistic types

Before introducing our probabilistic type system, we recall *simple types* for the PEA [9].

► **Definition 5.** Simple types are given by the following grammar.

$$A, B ::= o \mid A \rightarrow B$$

A typing context Γ is a finite function from term variables to types, written as a sequence $x_1: A_1, \dots, x_n: A_n$. A typing judgment $\Gamma \vdash M: A$ assigns the type A to the term M in the context Γ . The simply-typed probabilistic event λ -calculus PEA^\rightarrow is given by the typing rules in Figure 1.

Simple types ignore the probabilistic constructs of the calculus, *generator* and *choice*, requiring only that branches of a choice have equal types. This gives the expected result: typed terms are strongly normalizing, since every possible branch of the computation is typed. For *probabilistic* termination, we wish to capture that a given fraction of all branches of a computation, terminates – where our notion of *termination* is given by *head normalization*.

Our approach derives from the following idea: if the base type o of simple types, even if not inhabited, denotes *certainty* of termination, then we may generalise this to *probability* of termination by replacing base types with arbitrary probabilities. This yields the following notion of types.

► **Definition 6.** Probabilistic types are given by the following grammar, where $p \in [0, 1] \cap \mathbb{Q}$.

$$A, B ::= p \mid A \rightarrow B$$

The intuitive meaning of, for example, assigning a term M a type $A \rightarrow B \rightarrow p$ is: given inputs of type A and B , M terminates with probability at least p . The identity term $I = \lambda x. x$ may be assigned any type $p \rightarrow p$: given an input N that terminates with probability p , the term IN does so as well. The type system will include an axiom that any term, in particular a non-terminating one such as $\Omega = (\lambda x. xx)(\lambda x. xx)$, may be assigned a termination probability of zero. This is expressed by a type $A_1 \rightarrow \dots \rightarrow A_n \rightarrow 0$: for any inputs A_1 through A_n , the term will terminate with probability at least zero.

Note that these types generalise simple types with a single, uninhabited base type o . A probabilistic system with multiple base types, say integers and booleans, would pair the return type with a probability, for instance an integer with $\frac{1}{2}$ probability, or a boolean with $\frac{1}{4}$ probability. The extension of the calculus with sequencing, in Sections 5 and 6, will give a more concrete account of how probabilities interact with return values and return types.

$$\begin{array}{c}
 \frac{}{E \mid \Gamma, x : A \vdash x : A}^{\text{var}} \quad \frac{}{E \mid \Gamma \vdash M : \bar{A} \rightarrow 0}^{\text{zero}} \quad \frac{E \mid \Gamma \vdash M : \bar{A} \rightarrow q}{E \mid \Gamma \vdash M : \bar{A} \rightarrow p}^{\text{low } (p < q)} \\
 \\
 \frac{E \mid \Gamma, x : A \vdash M : B}{E \mid \Gamma \vdash \lambda x. M : A \rightarrow B}^{\text{abs}} \quad \frac{E \mid \Gamma \vdash M : A \rightarrow B \quad E \mid \Gamma \vdash N : A}{E \mid \Gamma \vdash MN : B}^{\text{app}} \\
 \\
 \frac{E, a \mapsto i \mid \Gamma \vdash M_i : A}{E, a \mapsto i \mid \Gamma \vdash M_{\top} a M_{\perp} : A}^{\text{chc}} \quad \frac{E, a \mapsto \top \mid \Gamma \vdash M : \bar{A} \rightarrow p \quad E, a \mapsto \perp \mid \Gamma \vdash M : \bar{A} \rightarrow q}{E \mid \Gamma \vdash \boxed{a}. M : \bar{A} \rightarrow \frac{1}{2}p + \frac{1}{2}q}^{\text{gen}}
 \end{array}$$

■ **Figure 2** The probabilistic type system $\text{PE}\Lambda^{\mathbb{Q}\rightarrow}$.

The term $\boxed{a}. M a \Omega$ gives a fair probabilistic choice between M and Ω . For the computation to be well-typed regardless of the choice, M and Ω should have the same input types, so that they can be applied to the same arguments. Hence, if M has type $A \rightarrow B \rightarrow p$, we may choose $A \rightarrow B \rightarrow 0$ for Ω . Then the type for $\boxed{a}. M a \Omega$ should be $A \rightarrow B \rightarrow \frac{1}{2}p$: given arguments of type A and B , the computation chooses M with probability $\frac{1}{2}$ and so terminates with probability at least $\frac{1}{2}p$.

These considerations motivate the shape of our probabilistic type system, with one further aspect to explain. An arbitrary generator term $\boxed{a}. M$ reduces with a fair probability to either $\pi_{\top}^a M$ or $\pi_{\perp}^a M$, which may then terminate with different probabilities. However, using projections would mean the type system loses the property of being inductive on terms. We will instead record the assignment of a truth value to a in an additional context E , that we call an *event valuation*. The type derivation then projects a term $M a N$ to M or to N according to the value of a in the event valuation E .

► **Definition 7.** *The probabilistically typed probabilistic event λ -calculus $\text{PE}\Lambda^{\mathbb{Q}\rightarrow}$ is given by the typing rules in Figure 2, using the following definitions. An event valuation E is a finite function from events to $\{\perp, \top\}$, written as a sequence $a_1 \mapsto i_1, \dots, a_n \mapsto i_n$. A typing context Γ is a finite function from variables to types, written $x_1 : A_1, \dots, x_n : A_n$. A probabilistic typing judgement $E \mid \Gamma \vdash M : A$ assigns a type A to the term M in the context of E and Γ . A sequence of antecedents $A_1 \rightarrow \dots \rightarrow A_n \rightarrow p$ is abbreviated with vector notation as $\bar{A} \rightarrow p$.*

The typing rules are syntax-driven, except for the low rule to lower a given bound. The rule is not essential to the type system, since it is already implied in the meaning of types as giving a lower bound; but for the same reason, since it is implied, including it brings more clarity than omitting it. Note further that the typing rule **gen** assumes a fair coin toss for the generator \boxed{a} , which gives the resulting probability of $\frac{1}{2}p + \frac{1}{2}q$. An unfair toss \boxed{a}_r with ratio r would give a probability $rp + (1-r)q$.

► **Example 8.** To demonstrate our type system we reprise the example $t[a, b]$ from [1, Section 6.2], written here as the term $\boxed{a}. \boxed{b}. \boxed{c}. T$ below. In Figure 3 we derive the following type.

$$\boxed{a}. \boxed{b}. \boxed{c}. T : 1 \rightarrow \frac{3}{8} \quad \text{where} \quad T = ((I c \Omega) b \Omega) a (\Omega b I)$$

We will discuss a number of aspects of our type system. First, as a particularly natural feature, observe that the rules **var**, **abs** and **app** make it conservative over standard simply-typed λ -calculus – which, interestingly, is agnostic to the choice of base types.

$$\begin{aligned}
D_{\perp,\perp,\perp} &= \frac{\frac{\frac{\frac{}{a \mapsto \perp, b \mapsto \perp, c \mapsto \perp \mid x: 1 \vdash x: 1}^{\text{var}}}{a \mapsto \perp, b \mapsto \perp, c \mapsto \perp \mid \vdash \lambda x. x: 1 \rightarrow 1}^{\text{abs}}}{a \mapsto \perp, b \mapsto \perp, c \mapsto \perp \mid \vdash I c \Omega: 1 \rightarrow 1}^{\text{chc}}}{a \mapsto \perp, b \mapsto \perp, c \mapsto \perp \mid \vdash (I c \Omega) b \Omega: 1 \rightarrow 1}^{\text{chc}}}{a \mapsto \perp, b \mapsto \perp, c \mapsto \perp \mid \vdash ((I c \Omega) b \Omega) a (\Omega b I): 1 \rightarrow 1}^{\text{chc}} \\
D_{\perp,\perp} &= \frac{D_{\perp,\perp,\perp} \quad \frac{}{a \mapsto \perp, b \mapsto \perp, c \mapsto \top \mid \vdash T: 1 \rightarrow 0}^{\text{zero}}}{a \mapsto \perp, b \mapsto \perp \mid \vdash \boxed{c}. T: 1 \rightarrow \frac{1}{2}}^{\text{gen}} \\
D_{\perp} &= \frac{D_{\perp,\perp} \quad \frac{}{a \mapsto \perp, b \mapsto \top \mid \vdash \boxed{c}. T: 1 \rightarrow 0}^{\text{zero}}}{a \mapsto \perp \mid \vdash \boxed{b}. \boxed{c}. T: 1 \rightarrow \frac{1}{4}}^{\text{chc}} \\
D_{\top,\top,i} &= \frac{\frac{\frac{\frac{}{a \mapsto \top, b \mapsto \top, c \mapsto i \mid x: 1 \vdash x: 1}^{\text{var}}}{a \mapsto \top, b \mapsto \top, c \mapsto i \mid \vdash \lambda x. x: 1 \rightarrow 1}^{\text{abs}}}{a \mapsto \top, b \mapsto \top, c \mapsto i \mid \vdash \Omega b I: 1 \rightarrow 1}^{\text{chc}}}{a \mapsto \top, b \mapsto \top, c \mapsto i \mid \vdash ((I c \Omega) b \Omega) a (\Omega b I): 1 \rightarrow 1}^{\text{chc}} \\
D_{\top,\top} &= \frac{D_{\top,\top,\perp} \quad D_{\top,\top,\top}}{a \mapsto \top, b \mapsto \top \mid \vdash \boxed{c}. T: 1 \rightarrow 1}^{\text{gen}} \\
D_{\top} &= \frac{\frac{}{a \mapsto \top, b \mapsto \perp \mid \vdash \boxed{c}. T: 1 \rightarrow 0}^{\text{zero}} \quad D_{\top,\top}}{a \mapsto \top \mid \vdash \boxed{b}. \boxed{c}. T: 1 \rightarrow \frac{1}{2}}^{\text{gen}} \\
D &= \frac{D_{\perp} \quad D_{\top}}{a \mapsto \perp \mid \vdash \boxed{b}. \boxed{c}. T: 1 \rightarrow \frac{1}{4} \quad a \mapsto \top \mid \vdash \boxed{b}. \boxed{c}. T: 1 \rightarrow \frac{1}{2}}^{\text{gen}} \\
&\quad \mid \vdash \boxed{a}. \boxed{b}. \boxed{c}. T: 1 \rightarrow \frac{3}{8}
\end{aligned}$$

■ **Figure 3** Typing derivations for Example 8, where $T = ((I c \Omega) b \Omega) a (\Omega b I)$.

31:8 Simple Types for Probabilistic Termination

Second, a key difference with the simple type system of Figure 1 is that probabilistic branching occurs in the *generator* rule, whereas for simple types it is the *choice* rule. This is due to aiming for *head normalization* instead of *strong normalization*. Consider the example term $\boxed{a}. M a (\Omega a N)$. Head reduction projects it to M and N , removing Ω , which will not be reduced. This is reflected in the probabilistic type system, where the branches of the generator typing rule project to M and N via the event valuation. The simple type system, reflecting strong normalization, would require also Ω to be typed – which of course it cannot.

For terms of the form $M \oplus N = \boxed{a}. M a N$ this difference is moot: both type systems branch similarly for this construct, to M and N . This leaves as only distinction the generalisation of types themselves, from a single base type o to probabilities p . In accordance with the meaning of a base type as the probability of termination, simple types are then recovered by restricting to base type $p = 1$. This rules out the rules **zero** and **low**, assigning a zero-weighted type $\bar{A} \rightarrow 0$ and reducing the weight of a type. It is easy to observe that the remaining rules preserve the restriction $p = 1$, to give the following proposition.

► **Proposition 9.** *For the fragment below left, the simply-typed PEA coincides with the probabilistically-typed PEA restricted to the types below right.*

$$M, N ::= x \mid \lambda x. M \mid M N \mid \boxed{a}. M a N \qquad A, B ::= 1 \mid A \rightarrow B$$

Our main result for the probabilistically-typed PEA is that a type $\bar{A} \rightarrow p$ guarantees head normalization with probability at least p . Formally, this is stated by a head reduction to a distribution of which a proportion of at least p is in head-normal form.

► **Theorem 10.** *For closed M , if $M : \bar{A} \rightarrow p$ then $M \rightarrow_{\text{h}} \mathcal{N}_0 + \mathcal{N}_1$ where all terms in \mathcal{N}_0 are head normal and $|\mathcal{N}_0| \geq p$.*

The result will follow directly from the corresponding Theorem 22 for the expanded probabilistic calculus with sequential composition, introduced in Section 5.

4 Comparison with counting quantifiers

In this section we will give a close comparison with the type system $\text{C}\lambda_{\rightarrow}$ of Antonelli, Dal Lago, and Pistone [1], and demonstrate that our approach gives tighter bounds on the probability of termination.

The first distinction between $\text{C}\lambda_{\rightarrow}$ and $\text{PEA}^{\text{Q}\rightarrow}$ is that the former uses indexed event variables x_a^i for $i \in \mathbb{N}$ instead of events a . However, there is no formal need for this; since the indices i are static, we may replace x_a^i and x_a^j for distinct i and j simply with distinct events a and b . This simplification extends to the semantics. Probabilities in $\text{C}\lambda_{\rightarrow}$ are given by boolean formulas \mathfrak{b} over the variables x_a^i , indicating a subset of the space $(2^{\mathbb{N}})^X$ where X is a finite set of events. However, since the indices i are fixed and bounded, it is sufficient to consider finite spaces 2^X . The elements of this set are the event valuations E with domain X , and a boolean formula \mathfrak{b} over X indicates the set of event valuations $\llbracket \mathfrak{b} \rrbracket_X = \{E \in 2^X \mid E \models \mathfrak{b}\}$, where $E \models \mathfrak{b}$ is characterized syntactically as expected:

$$\frac{}{E, a \mapsto \top \models a} \quad \frac{}{E, a \mapsto \perp \models \neg a} \quad \frac{}{E \models \top} \quad \frac{E \models \mathfrak{b} \quad E \models \mathfrak{c}}{E \models \mathfrak{b} \wedge \mathfrak{c}} \quad \frac{E \models \mathfrak{b}}{E \models \mathfrak{b} \vee \mathfrak{c}}$$

For a set $\mathcal{E} \subseteq 2^X$ the measure $\mu_X(\mathcal{E})$ is given by

$$\mu_X(\mathcal{E}) = \frac{|\mathcal{E}|}{2^{|X|}}$$

where $|S|$ denotes the size of a set S . Then $\mu(\mathfrak{b})$ is $\mu_X(\llbracket \mathfrak{b} \rrbracket_X)$ where X is the domain of \mathfrak{b} . This coincides with the measure $\mu(\mathfrak{b})$ over $(2^{\mathbb{N}})^X$ in [1], as the latter makes no essential use of the infinity offered by \mathbb{N} .

The second difference is the syntax of types: $\text{C}\lambda_{\rightarrow}$ introduces probabilities through counting quantifiers C^p , where $\text{PE}\Lambda^{\mathbb{Q}\rightarrow}$ has probabilities as base types. Types are nevertheless isomorphic: $\text{C}\lambda_{\rightarrow}$ types \mathfrak{s} are of the form $C^p(\mathfrak{s}_1 \rightarrow \dots \rightarrow \mathfrak{s}_n \rightarrow o)$, with a fixed outer counting quantifier, and map 1-to-1 onto $\text{PE}\Lambda^{\mathbb{Q}\rightarrow}$ types A of the form $A_1 \rightarrow \dots \rightarrow A_n \rightarrow p$ by associating the probability p instead with the consequent o . Formally, we encode types \mathfrak{s} and typing contexts Γ of $\text{C}\lambda_{\rightarrow}$ into our setting as follows.

$$\begin{aligned} \llbracket C^q(o) \rrbracket &= q & \llbracket x_1 : \mathfrak{s}_1, \dots, x_n : \mathfrak{s}_n \rrbracket &= x_1 : \llbracket \mathfrak{s}_1 \rrbracket, \dots, x_n : \llbracket \mathfrak{s}_n \rrbracket \\ \llbracket C^q(\mathfrak{s} \Rightarrow \tau) \rrbracket &= \llbracket \mathfrak{s} \rrbracket \rightarrow \llbracket C^q(\tau) \rrbracket \end{aligned}$$

Having connected boolean formulas \mathfrak{b} to event valuations E , and $\text{C}\lambda_{\rightarrow}$ types to $\text{PE}\Lambda^{\mathbb{Q}\rightarrow}$ types, we may state the following conservativity result of $\text{PE}\Lambda^{\mathbb{Q}\rightarrow}$ over $\text{C}\lambda_{\rightarrow}$.

► **Proposition 11.** *If $E \in \llbracket \mathfrak{b} \rrbracket_X$ then*

$$\Gamma \vdash^X M : \mathfrak{b} \rightsquigarrow \mathfrak{s} \quad \text{implies} \quad E \mid \llbracket \Gamma \rrbracket \vdash M : \llbracket \mathfrak{s} \rrbracket.$$

Proof. By induction on the typing derivation for $\Gamma \vdash^X M : \mathfrak{b} \rightsquigarrow \mathfrak{s}$. ◀

► **Corollary 12.** *For a given closed term M the type system $\text{PE}\Lambda^{\mathbb{Q}\rightarrow}$ gives the same or higher termination bounds than $\text{C}\lambda_{\rightarrow}$.*

In the reverse direction, Example 8 and its counterpart in [1, Section 6.2] show that the two type systems do not give the exact same bounds, and in some cases $\text{PE}\Lambda^{\mathbb{Q}\rightarrow}$ gives a strictly higher bound. The reason is that $\text{PE}\Lambda^{\mathbb{Q}\rightarrow}$ locates branching between alternatives at the **gen**-rule for $\boxed{a}.M$, where $\text{C}\lambda_{\rightarrow}$ branches for *choice* terms $M a N$ or in a contraction rule. Crucially, the **gen**-rule allows branches with different termination bounds. We illustrate this further by attempting to simulate the rule for $\boxed{a}.M$ in $\text{C}\lambda_{\rightarrow}$.

$$\frac{E, a \mapsto \top \mid \llbracket \Gamma \rrbracket \vdash M : \bar{A} \rightarrow p \quad E, a \mapsto \perp \mid \llbracket \Gamma \rrbracket \vdash M : \bar{A} \rightarrow q}{E \mid \llbracket \Gamma \rrbracket \vdash \boxed{a}.M : \bar{A} \rightarrow \frac{1}{2}p + \frac{1}{2}q} \text{gen}$$

The branching at this rule in $\text{C}\lambda_{\rightarrow}$ is captured with a contraction on the two premisses, which requires the probabilities to be equal, i.e. $p = q$. The derivation is as follows, where $d = \top$ and $\mu(d) = 1$ for the counting rule.

$$\frac{\Gamma \vdash^{X \cup \{a\}} M : \mathfrak{b} \wedge a \rightsquigarrow C^p \sigma \quad \Gamma \vdash^{X \cup \{a\}} M : \mathfrak{b} \wedge \neg a \rightsquigarrow C^p \sigma \quad \mathfrak{b} \models (\mathfrak{b} \wedge a) \vee (\mathfrak{b} \wedge \neg a)}{\frac{\Gamma \vdash^{X \cup \{a\}} M : \mathfrak{b} \rightsquigarrow C^p \sigma}{\Gamma \vdash^X \boxed{a}.M : \mathfrak{b} \rightsquigarrow C^p \sigma}}$$

This is the issue illustrated by Example 8 and its counterpart in [1, Section 6.2], for which $\text{PE}\Lambda^{\mathbb{Q}\rightarrow}$ gives the actual termination probability of $\frac{3}{8}$, while $\text{C}\lambda_{\rightarrow}$ gives a best approximation of $\frac{1}{4}$. Reprising the example in $\text{C}\lambda_{\rightarrow}$, the two sub-derivations for $\boxed{b}.\boxed{c}.T$, given in condensed form below, assign probabilities of $\frac{1}{4}$ and $\frac{1}{2}$. These may only be combined in a contraction by lowering the first probability to match the $\frac{1}{4}$ of the first.

$$\frac{\frac{\frac{\vdash^{\{a,b,c\}} T : a \wedge b \wedge c \rightsquigarrow C^1 \sigma}{\vdash^{\{a,b\}} \boxed{c}.T : a \wedge b \rightsquigarrow C^{\frac{1}{2}} \sigma}}{\vdash^{\{a\}} \boxed{b}.\boxed{c}.T : a \rightsquigarrow C^{\frac{1}{4}} \sigma}}{\frac{\frac{\vdash^{\{a,b,c\}} T : \neg a \wedge \neg b \rightsquigarrow C^1 \sigma}{\vdash^{\{a,b\}} \boxed{c}.T : \neg a \wedge \neg b \rightsquigarrow C^1 \sigma}}{\vdash^{\{a\}} \boxed{b}.\boxed{c}.T : \neg a \rightsquigarrow C^{\frac{1}{2}} \sigma}}$$

The above analysis suggests that this issue with $C\lambda_{\rightarrow}$ may be fixed by adopting the generator typing rule of $\text{PEA}^{\mathbb{Q}\rightarrow}$, adjusted appropriately as follows. The key here is that different branches feature the dual atoms a and $\neg a$, not seen in either the simple type system $\text{PEA}^{\mathbb{Q}\rightarrow}$ nor the intersection type system in [1].

$$\frac{\Gamma \vdash^{X \cup \{a\}} M : c \wedge a \multimap C^p \sigma \quad \Gamma \vdash^{X \cup \{a\}} M : d \wedge \neg a \multimap C^q \sigma \quad \mathfrak{b} \models c \vee d}{\Gamma \vdash^X \boxed{a}. M : \mathfrak{b} \multimap C^{\frac{1}{2}p + \frac{1}{2}q} \sigma}$$

5 Sequencing

Two observations about probabilistic λ -calculi motivate the developments in the remainder of this paper. The first is the primary role of *head reduction*. It is well known that head normalization corresponds closely to evaluation on the Krivine Machine [23], and sometimes the machine gives a more natural model of what is being studied. The second is that probabilistic evaluation in λ -calculi needs to account for the difference between call-by-value (cbv) and call-by-name (cbn), to which end additional constructions are introduced, sometimes ad-hoc. The PEA is an example, as is the separate consideration of a cbv- and a cbn-probabilistic sum by Faggian and Ronchi Della Rocca [16], while Antonelli, Dal Lago, and Pistone [1] add to the standard cbn-application a second cbv-application.

These observations prompted us to consider probabilistic termination from the perspective of the *Functional Machine Calculus* (FMC) [18], a λ -calculus with computational effects. Firstly, the Krivine Machine plays a central role in the FMC (indeed it is the “M” in “FMC”), while the FMC provides the machine with a new notion of *successful termination*, absent from the standard λ -calculus. It is then a natural question if and how this may be used to capture the *probability* of successful termination. Secondly, the FMC may express both cbn and cbv behaviour, with the cbn λ -calculus a fragment and the cbv λ -calculus encoded in the syntax. The need for ad-hoc constructs to control reduction behaviour is thus avoided.

The Krivine Machine, simplified by replacing *environments* with *substitution*, evaluates a λ -term in the presence of a stack of input terms. An abstraction $\lambda x. M$ pops the top off the stack, say N , and continues as $\{N/x\}M$, while an application $M N$ pushes its argument N and continues as M . A λ -term may thus be viewed as a language of instruction sequences for this machine: application-*push*, abstraction-*pop*, variable-*execute*.

The FMC then extends the λ -calculus with *sequential composition* $M ; N$ and its unit, the imperative *skip* \star , with the expected semantics: *concatenation* of machine instructions and the *empty* instruction. As in models of imperative languages, *skip* indicates successful termination of the machine. This gives a fragment called the *sequential λ -calculus*.

We adopt these modifications in the PEA to give the *sequential probabilistic event λ -calculus* (SPEA), defined below. Following [18] we render abstraction as $\langle x \rangle. M = \lambda x. M$ and application as $[N]. M = M N$ to emphasise the machine behaviour of *pop* and *push*, retaining the standard syntax as a shorthand. In particular, the new notation clarifies the interaction between *push* and *sequencing*: the following three terms are equivalent, rendered first in standard notation and second with prefix application.

$$(\star N); M \sim (\star; M) N \sim M N \quad ([N].\star); M \sim [N].(\star; M) \sim [N].M$$

The full FMC further generalises to a machine with multiple independent stacks, addressed by a set of *locations*, in which *pop* and *push* are then parameterised to operate on the corresponding stack. This allows us to encode the effects of mutable higher-order store, input/output, and indeed probabilistic computation: the generator $\boxed{a}. M$ of the PEA is, in the FMC, an abstraction $\text{rnd}\langle a \rangle. M$ parameterised to draw from a stream of random values labelled *rnd*. The SPEA is thus a fragment of the FMC with two locations.

► **Definition 13.** *The sequential probabilistic event λ -calculus SPEA is given as follows.*

$$M, N, P ::= x \mid \langle x \rangle. M \mid [N]. M \mid \boxed{a}. M \mid M a N \mid \star \mid N; M$$

Prefixing binds tighter than sequencing, $[N]. M; P = ([N]. M); P$, and sequencing associates right, $M; N; P = M; (N; P)$. *Projections* and *contexts* extend to the SPEA as below; *head contexts* and *applicative contexts* are as for the PEA.

$$\pi_i^a \star = \star \quad \pi_i^a(N; M) = \pi_i^a N; \pi_i^a M \quad C ::= \dots \mid C; M \mid M; C$$

The interaction between sequentiality and the λ -calculus is governed by the following *sequencing* reduction rules. These make sequential composition right-associative, and let the *prefixing* of *push*, *pop*, and the generator propagate past it (as in a standard list concatenation algorithm). The result is to make the first such action on the abstract machine the leading construct.

$$\star; P \rightarrow_\sigma P \quad \begin{array}{l} (N; M); P \rightarrow_\sigma N; (M; P) \\ [N]. M; P \rightarrow_\sigma [N]. (M; P) \end{array} \quad \begin{array}{l} \langle x \rangle. M; P \rightarrow_\sigma \langle x \rangle. (M; P) \quad (x \notin \text{fv}(P)) \\ \boxed{a}. M; P \rightarrow_\sigma \boxed{a}. (M; P) \quad (a \notin \text{fe}(P)) \end{array}$$

The *sequencing* relation \rightarrow_σ is given by closing these rules under all contexts C , and *head sequencing* $\rightarrow_{\sigma h}$ by closing under head contexts H only. The β -reduction rule in SPEA notation is as below left, with \rightarrow_β given by closing under all contexts and $\rightarrow_{\beta h}$ by closing under head contexts. Projective reduction, below right, is as previously.

$$[N]. \langle x \rangle. M \rightarrow_\beta \{N/x\}M \quad H\{\boxed{a}. M\} \rightarrow_\pi \left[\frac{1}{2} \cdot H\{\pi_+^a M\}, \frac{1}{2} \cdot H\{\pi_-^a M\} \right]$$

Head reduction \rightarrow_h is the union of all three head relations:

$$\rightarrow_h = \rightarrow_{\beta h} \cup \rightarrow_{\sigma h} \cup \rightarrow_\pi .$$

We round off by observing the shape of head-normal forms, assuming no free event variables.

► **Proposition 14.** *The head-normal forms of event-closed SPEA-terms are of one of the three forms $H\{\star\}$, $H\{x\}$, and $H\{x; M\}$.*

5.1 Encoding call-by-value

Sequential composition provides an essential element that the λ -calculus lacks, and which is at the heart of the cbv/cbn dichotomy: *control over execution*. The cbv behaviour of an application $M N$ is encoded almost as $N; M$: *first* evaluate the argument N , *then* evaluate the function M (the full encoding, below, includes an extra part to also *execute* M).

We demonstrate the encoding of the cbv-probabilistic λ -calculus $\Lambda_{\oplus}^{\text{cbv}}$ of Faggian and Ronchi Della Rocca [16]. The encoding of cbv λ -terms is standard: see [13, 18, 28]. *Values* V, W and *terms* M, N encode by the translations $-_v$ and $-_t$ respectively, below.

$$\begin{array}{ll} \text{Values } V, W : & x_v = x \\ & (\lambda x. M)_v = \langle x \rangle. M_t \end{array} \quad \begin{array}{ll} \text{Terms } M, N : & V_t = [V_v]. \star \\ & (M N)_t = N_t; M_t; \langle x \rangle. x \\ & (M \oplus N)_t = \boxed{a}. M_t a N_t \end{array}$$

The operational intuition is that a *push* represents a *return value*: a term M_t evaluates until it is of the form $V_t = [V_v]. \star$, at which point V_v is pushed to the stack and the machine terminates. Then β -reduction is simulated as follows.

31:12 Simple Types for Probabilistic Termination

$$\begin{aligned}
((\lambda x.M)V)_t &= [V_v]. \star ; [\langle x \rangle. M_t]. \star ; \langle y \rangle. y \\
&\rightarrow_{\sigma} [V_v]. [\langle x \rangle. M_t]. \langle y \rangle. y \\
&\rightarrow_{\beta} [V_v]. \langle x \rangle. M_t \\
&\rightarrow_{\beta} \{V_v/x\}M_t \\
&= (\{V/x\}M)_t
\end{aligned}$$

The probabilistic reduction rule of $\Lambda_{\oplus}^{\text{cbv}}$ is that of projective reduction, under the given encoding, but it applies in *surface contexts*:

$$S ::= \{ \} \mid MS \mid SM$$

Then the translation of $S\{M \oplus N\}$ indeed does not place the probabilistic redex inside a *push*, the requirement for correct behaviour.

5.2 The abstract machine

The small-step operational semantics of the SPEA is given by the following abstract machine. A *state* is a triple (S, M, K) , where M is a term and S and K are stacks of terms. S is the *operand stack*, with the head to the right as SN , and K is the *continuation stack*, with the head to the left as NK . In both cases the empty stack is written ε , and concatenation by juxtaposition, ST . *Transitions* or *steps* are probabilistic: a transition rule is written as below left, expressing that the machine transitions from a state (S, M, K) to (T, N, L) with probability p . We may omit p when $p = 1$. A *run* is a sequence of steps, written as below centre, where probabilities are multiplied, i.e. p below is the product of the probabilities of all steps. A run is *successful* if it terminates with *skip* and an empty continuation stack, as below right; the stack T then holds the *return values* of the computation.

$$\begin{array}{lll}
\text{step: } p \frac{(S, M, K)}{(T, N, L)} & \text{run: } p \frac{(S, M, K)}{(T, N, L)} & \text{successful run: } p \frac{(S, M, K)}{(T, \star, \varepsilon)}
\end{array}$$

► **Definition 15.** *The sequential probabilistic machine (SPM) is given by the following probabilistic transitions.*

$$\begin{array}{lll}
\frac{(S, [N].M, K)}{(SN, M, K)} & \frac{(S, N; M, K)}{(S, N, MK)} & \frac{\frac{1}{2}(S, \boxed{a}.M, K)}{\frac{1}{2}(S, \pi_{\top}^a M, K)} \\
\frac{(SN, \langle x \rangle. M, K)}{(S, \{N/x\}M, K)} & \frac{(S, \star, MK)}{(S, M, K)} & \frac{\frac{1}{2}(S, \boxed{a}.M, K)}{\frac{1}{2}(S, \pi_{\perp}^a M, K)}
\end{array}$$

5.3 Big-step semantics

Running the machine for a given term and input stack gives a distribution of return stacks. We use the following notation, extending from that for distributions over terms.

$$\mathcal{T} = [p_1 \cdot T_1, \dots, p_n \cdot T_n] = [p_i \cdot T_i]_{i \leq n}$$

► **Definition 16.** *The evaluation relation $S, M \Downarrow \mathcal{T}$ is defined inductively by the following rules.*

$$\begin{array}{lll}
\frac{}{S, \star \Downarrow [1 \cdot S]} & \frac{S, \{N/x\}M \Downarrow \mathcal{T}}{SN, \langle x \rangle. M \Downarrow \mathcal{T}} & \frac{R, M \Downarrow [p_i \cdot S_i]_{i \leq n} \quad (S_i, N \Downarrow \mathcal{T}_i)_{i \leq n}}{R, M; N \Downarrow \sum_{i \leq n} p_i \mathcal{T}_i} \\
\frac{}{S, M \Downarrow \emptyset} & \frac{SN, M \Downarrow \mathcal{T}}{S, [N]. M \Downarrow \mathcal{T}} & \frac{S, \pi_{\top}^a M \Downarrow \mathcal{T}_{\top} \quad S, \pi_{\perp}^a M \Downarrow \mathcal{T}_{\perp}}{S, \boxed{a}. M \Downarrow \frac{1}{2} \mathcal{T}_{\top} + \frac{1}{2} \mathcal{T}_{\perp}}
\end{array}$$

We demonstrate that small-step and big-step semantics agree.

► **Proposition 17.** $S, M \Downarrow \mathcal{T}$ if and only if there is a finite collection of n distinct runs

$$p_i \frac{(S, M, \varepsilon)}{(T_i, \star, \varepsilon)} \quad (i \leq n) \quad \text{such that} \quad \mathcal{T} = [p_i \cdot T_i]_{i \leq n} .$$

Proof. (\implies) By induction on $S, M \Downarrow \mathcal{T}$. (\impliedby) By induction each run in the collection of n runs. ◀

6 Sequential probabilistic types

Types for the sequential λ -calculus are of the form below, with the meaning: given an input stack of terms typed by A_1 through A_n , the machine will terminate successfully and return a stack with types B_1 through B_m .

$$A_n \dots A_1 \Rightarrow B_1 \dots B_m$$

For the SPEA, we parameterize this with the *probability* of successful termination.

► **Definition 18.** Sequential probabilistic types are given by the following grammars, where p is a probability.

$$\begin{aligned} A, B, C &::= \overline{A} \stackrel{p}{\Rightarrow} \overline{C} && (\text{types}) \\ \overline{A} &::= A_1 \dots A_n && (\text{type vectors}) \end{aligned}$$

A typing judgement $E \mid \Gamma \vdash M : A$ assigns a term M the type A in the context of E and Γ , and $E \mid \Gamma \vdash S : \overline{A}$ assigns a stack of terms S a type vector \overline{A} . The sequential probabilistic type system $\text{SPEA}^{\mathbb{Q}\Rightarrow}$ is given by the typing rules in Figure 4. We may omit p when $p = 1$.

There are no base types: their rôle is subsumed by types with empty vectors ($\stackrel{p}{\Rightarrow}$). Observe that because stacks are last-in-first-out, the identity term on two elements is $\langle x \rangle. \langle y \rangle. [y]. [x]. \star$, i.e. with the order of x and y reversed between popping and pushing. We match this reversal in types, and assign this term the type $AB \Rightarrow BA$. Since we want identity types to be of the form $\overline{A} \Rightarrow \overline{A}$, in a type $\overline{A} \stackrel{p}{\Rightarrow} \overline{C}$ we consider the antecedent type vector \overline{A} to be reversed, i.e.

$$\overline{A} \Rightarrow \overline{A} = A_n \dots A_1 \Rightarrow A_1 \dots A_n .$$

Probabilistic PEA-types embed into sequential types by $\overline{A} \rightarrow p = \overline{A} \stackrel{p}{\Rightarrow} \varepsilon$. With this identification, for PEA-terms and -types the two type systems coincide.

► **Proposition 19.** For M a PEA-term, $E \mid \Gamma \vdash M : \overline{A} \rightarrow p$ if and only if $E \mid \Gamma \vdash M : \overline{A} \stackrel{p}{\Rightarrow} \varepsilon$.

Every type is inhabited by a closed term. For a type A , define the *zero term* 0_A as follows: for $A = B_1 \dots B_m \stackrel{p}{\Rightarrow} C_1 \dots C_n$,

$$0_A = \langle x_1 \rangle \dots \langle x_m \rangle. [0_{C_1}] \dots [0_{C_n}]. \star .$$

► **Proposition 20** (Type inhabitation). Every type A is inhabited by its zero term, $\vdash 0_A : A$.

Proof. By induction on the type A , using the low rule to lower the termination probability from 1 to an arbitrary p . ◀

31:14 Simple Types for Probabilistic Termination

$$\begin{array}{c}
\frac{}{E \mid \Gamma, x : \bar{A} \stackrel{p}{\Rightarrow} \bar{C} \vdash x : \bar{A} \bar{B} \stackrel{p}{\Rightarrow} \bar{B} \bar{C}} \text{var} \quad \frac{}{E \mid \Gamma \vdash M : \bar{A} \stackrel{0}{\Rightarrow} \bar{C}} \text{zero} \\
\\
\frac{E \mid \Gamma, x : A \vdash M : \bar{B} \stackrel{p}{\Rightarrow} \bar{C}}{E \mid \Gamma \vdash \langle x \rangle . M : \bar{A} \bar{B} \stackrel{p}{\Rightarrow} \bar{C}} \text{abs} \quad \frac{E \mid \Gamma \vdash N : A \quad E \mid \Gamma \vdash M : \bar{A} \bar{B} \stackrel{p}{\Rightarrow} \bar{C}}{E \mid \Gamma \vdash [N] . M : \bar{B} \stackrel{p}{\Rightarrow} \bar{C}} \text{app} \\
\\
\frac{E, a \mapsto i \mid \Gamma \vdash M_i : A}{E, a \mapsto i \mid \Gamma \vdash M_{\top} a M_{\perp} : A} \text{chc} \quad \frac{E, a \mapsto \top \mid \Gamma \vdash M : \bar{A} \stackrel{p}{\Rightarrow} \bar{C} \quad E, a \mapsto \perp \mid \Gamma \vdash M : \bar{A} \stackrel{q}{\Rightarrow} \bar{C}}{E \mid \Gamma \vdash \boxed{a} . M : \bar{A} \stackrel{\frac{1}{2}p + \frac{1}{2}q}{\Rightarrow} \bar{C}} \text{gen} \\
\\
\frac{}{E \mid \Gamma \vdash \star : \bar{A} \stackrel{1}{\Rightarrow} \bar{A}} \text{skip} \quad \frac{E \mid \Gamma \vdash M : \bar{A} \stackrel{p}{\Rightarrow} \bar{B} \quad E \mid \Gamma \vdash N : \bar{B} \stackrel{q}{\Rightarrow} \bar{C}}{E \mid \Gamma \vdash M ; N : \bar{A} \stackrel{pq}{\Rightarrow} \bar{C}} \text{seq} \\
\\
\frac{E \mid \Gamma \vdash M : \bar{A} \stackrel{q}{\Rightarrow} \bar{C}}{E \mid \Gamma \vdash M : \bar{A} \stackrel{p}{\Rightarrow} \bar{C}} \text{low } (p < q) \quad \frac{(E \mid \Gamma \vdash M_i : A_i)_{i \leq n}}{E \mid \Gamma \vdash \varepsilon M_1 \dots M_n : A_1 \dots A_n} \text{stk}
\end{array}$$

■ **Figure 4** The sequential probabilistic type system $\text{SPEA}^{\text{Q}\Rightarrow}$.

Type inhabitation is an important justification for the interpretation of types as a guarantee for successful machine termination: it means that for a typed term $M : A \stackrel{p}{\Rightarrow} \bar{C}$, a suitable input stack $S : \bar{A}$ always exists. Our main theorem, below, formalizes this interpretation: for a term $M : A \stackrel{p}{\Rightarrow} \bar{C}$ and stack $S : \bar{A}$, the machine successfully returns a stack $T : \bar{C}$ with probability at least p .

► **Theorem 21.** *For a typed, closed term $M : \bar{A} \stackrel{p}{\Rightarrow} \bar{C}$ and stack $S : \bar{A}$ there is a finite set of distinct successful runs*

$$p_i \frac{(S, M, \varepsilon)}{(T_i, \star, \varepsilon)} \quad (i \leq n)$$

with sum probability $\sum_{i \leq n} p_i \geq p$.

Proof. See Appendix A.1. ◀

The lower bound given by the theorem is an exact bound when two conditions are met: the typing derivation does not use the low rule to lower the termination bound, and every use of the zero rule to assign a zero probability applies to a term that is in fact non-terminating. In such a case, such as Example 8, types can be seen to give an exact bound.

Since head reduction as a strategy closely follows machine evaluation, it is no surprise that the termination bound for the machine is also a lower bound for probabilistic head reduction. This is formalised in the following theorem.

► **Theorem 22.** *For a closed M , if $M : \bar{A} \stackrel{p}{\Rightarrow} \bar{C}$ then $M \rightarrow_{\text{h}} \mathcal{N}_0 + \mathcal{N}_1$ where all terms in \mathcal{N}_0 are head normal and $|\mathcal{N}_0| \geq p$.*

Proof. See Appendix A.2. ◀

Observe that our main theorem on probabilistic termination for the PEA, Theorem 10, follows immediately from the corresponding Theorem 22 above by conservativity of the SPEA over the PEA. This is despite the fact that the proof of Theorem 22 makes essential use of the notion of successful termination made available by sequencing, absent from the PEA.

7 Conclusions

To us, what stands out about our approach is the simplicity and the natural intuition of our type system. This manifests in the transparent reasoning in our proofs, which despite using deep techniques such as abstract reducibility, give a direct and clear connection between types, machine behaviour, and reduction.

Compared with the approach in [1], the simplicity of our approach manifests in several advantages. First is to avoid the need for *counting quantifiers*, associating probabilities instead with base types, and the use of simple event valuations over boolean formulas. Second, it is clear that the expression of both call-by-name and call-by-value behaviour is an essential ingredient for a probabilistic calculus. We eschew the introduction of ad-hoc constructs, relying instead on a principled interpretation of *cbv* via sequential composition. Finally, where the main example [1, Example 6.2] requires intersection types to produce the exact bound, our approach does so directly in Example 8, while morally remaining within the realm of simple types, and strictly improving on the bounds provided by $\text{C}\lambda_{\rightarrow}$.

References

- 1 Melissa Antonelli, Ugo Dal Lago, and Paolo Pistone. Curry and howard meet borel. In *Proceedings of the 37th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS'22)*, pages 45:1–45:13. ACM, 2022. doi:10.1145/3531130.3533361.
- 2 Martin Avanzini, Ugo Dal Lago, and Akihisa Yamada. On probabilistic term rewriting. In John P. Gallagher and Martin Sulzmann, editors, *Proc. Functional and Logic Programming - 14th International Symposium, FLOPS 2018*, volume 10818 of *Lecture Notes in Computer Science*, pages 132–148. Springer, 2018. doi:10.1007/978-3-319-90686-7_9.
- 3 Chris Barrett, Willem Heijltjes, and Guy McCusker. The Functional Machine Calculus II: Semantics. In *31st EACSL Annual Conference on Computer Science Logic (CSL 2023)*, volume 252 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 10:1–10:18, 2023. doi:10.4230/LIPIcs.CSL.2023.10.
- 4 Raven Beutner and Luke Ong. On probabilistic termination of functional programs with continuous distributions. In Stephen N. Freund and Eran Yahav, editors, *Proceedings 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation (PLDI)*, pages 1312–1326. ACM, 2021. doi:10.1145/3453483.3454111.
- 5 Flavien Breuvert and Ugo Dal Lago. On intersection types and probabilistic lambda calculi. In *Proceedings of the 20th International Symposium on Principles and Practice of Declarative Programming, PPDP 2018*, pages 8:1–8:13. ACM, 2018. doi:10.1145/3236950.3236968.
- 6 Fredrik Dahlqvist and Dexter Kozen. Semantics of higher-order probabilistic programs with conditioning. *Proceedings of the ACM on Programming Languages*, 4(POPL):57:1–57:29, 2020. doi:10.1145/3371125.
- 7 Ugo Dal Lago, Claudia Faggian, and Simona Ronchi Della Rocca. Intersection types and (positive) almost-sure termination. *Proceedings of the ACM on Programming Languages*, 5(POPL):1–32, 2021. doi:10.1145/3434313.
- 8 Ugo Dal Lago and Charles Grellois. Probabilistic termination by monadic affine sized typing. *ACM Transactions on Programming Languages and Systems*, 41(2):10:1–10:65, 2019. doi:10.1145/3293605.
- 9 Ugo Dal Lago, Giulio Guerrieri, and Willem Heijltjes. Decomposing probabilistic lambda-calculi. In *Proceedings of Foundations of Software Science and Computation Structures (FoSSaCS)*, volume 12077 of *LNCS*, pages 136–156, 2020. doi:10.1007/978-3-030-45231-5_8.
- 10 Ugo Dal Lago and Margherita Zorzi. Probabilistic operational semantics for the lambda calculus. *RAIRO - Theoretical Informatics and Applications*, 46(3):413–450, 2012. doi:10.1051/ita/2012012.

- 11 Karel de Leeuw, Edward F. Moore, Claude E. Shannon, and Norman Shapiro. Computability by probabilistic machines. *Automata studies*, 34:183–198, 1956.
- 12 Ugo De'Liguoro and Adolfo Piperno. Non deterministic extensions of untyped lambda-calculus. *Information and Computation*, 122(2):149–177, 1995. doi:10.1006/inco.1995.1145.
- 13 Rémi Douence and Pascal Fradet. A systematic study of functional language implementations. *ACM Transactions on Programming Languages and Systems*, 20(2):344–387, 1998. doi:10.1145/276393.276397.
- 14 Thomas Ehrhard and Giulio Guerrieri. The bang calculus: An untyped lambda-calculus generalizing call-by-name and call-by-value. In *Proceedings of the 18th International Symposium on Principles and Practice of Declarative Programming (PPDP'16)*, pages 174–187, 2016. doi:10.1145/2967973.2968608.
- 15 Thomas Ehrhard and Christine Tasson. Probabilistic call by push value. *Logical Methods in Computer Science*, 15(1), 2019. doi:10.23638/LMCS-15(1:3)2019.
- 16 Claudia Faggian and Simona Ronchi Della Rocca. Lambda calculus and probabilistic computation. In *34th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2019*, pages 1–13, 2019. doi:10.1109/LICS.2019.8785699.
- 17 Jean Goubault-Larrecq. A probabilistic and non-deterministic call-by-push-value language. In *34th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2019*, pages 1–13. IEEE Computer Society, 2019. doi:10.1109/LICS.2019.8785809.
- 18 Willem Heijltjes. The functional machine calculus. In *Proceedings of the 38th Conference on the Mathematical Foundations of Programming Semantics, MFPS XXXVIII*, volume 1 of *ENTICS*, 2022. doi:10.46298/ENTICS.10513.
- 19 C. Jones and G. D. Plotkin. A probabilistic powerdomain of evaluations. In *Proceedings of the Fourth Annual Symposium on Logic in Computer Science (LICS '89)*, pages 186–195. IEEE Computer Society, 1989. doi:10.1109/LICS.1989.39173.
- 20 Achim Jung and Regina Tix. The troublesome probabilistic powerdomain. *Electronic Notes in Theoretical Computer Science*, 13:70–91, 1998. doi:10.1016/S1571-0661(05)80216-6.
- 21 Benjamin Kaminski. *Advanced Weakest Precondition Calculi for Probabilistic Programs*. PhD thesis, Fakultät für Mathematik, Informatik und Naturwissenschaften, RWTH Aachen University, 2019. doi:10.18154/RWTH-2019-01829.
- 22 Naoki Kobayashi, Ugo Dal Lago, and Charles Grellois. On the termination problem for probabilistic higher-order recursive programs. In *Proceedings of the 34th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2019*, pages 1–14. IEEE, 2019. doi:10.1109/LICS.2019.8785679.
- 23 Jean-Louis Krivine. A call-by-name lambda-calculus machine. *Higher-Order and Symbolic Computation*, 20(3):199–207, 2007. doi:10.1007/s10990-007-9018-9.
- 24 Thomas Leventis. A deterministic rewrite system for the probabilistic λ -calculus. *Mathematical Structures in Computer Science*, 29(10):1479–1512, 2019. doi:10.1017/S0960129519000045.
- 25 Paul Blain Levy. *Call-by-push-value: A functional/imperative synthesis*, volume 2 of *Semantic Structures in Computation*. Springer Netherlands, 2003. doi:10.1007/978-94-007-0954-6.
- 26 Udi Manber and Martin Tompa. Probabilistic, nondeterministic, and alternating decision trees. In *14th Annual ACM Symposium on Theory of Computing*, pages 234–244, 1982. doi:10.1145/800070.802197.
- 27 Annabelle McIver and Carroll Morgan. Abstraction and refinement in probabilistic systems. *SIGMETRICS Perform. Eval. Rev.*, 32(4):41–47, March 2005. doi:10.1145/1059816.1059824.
- 28 A.J. Power and Hayo Thielecke. Closed Freyd- and κ -categories. In *International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 1644 of *Lecture Notes in Computer Science*, pages 625–634. Springer, 1999. doi:10.1007/3-540-48523-6_59.
- 29 Norman Ramsey and Avi Pfeffer. Stochastic lambda calculus and monads of probability distributions. In *Conference Record of POPL 2002: The 29th SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '02, pages 154–165, 2002. doi:10.1145/503272.503288.

- 30 Nasser Saheb-Djahromi. Probabilistic LCF. In *Mathematical Foundations of Computer Science 1978, Proceedings, 7th Symposium*, volume 64 of *Lecture Notes in Computer Science*, pages 442–451. Springer, 1978. doi:10.1007/3-540-08921-7_92.
- 31 Davide Sangiorgi and Valeria Vignudelli. Environmental bisimulations for probabilistic higher-order languages. In *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2016*, pages 595–607, 2016. doi:10.1145/2837614.2837651.

A Proofs for Section 6

A.1 Probabilistic machine termination

Our main theorem for the typed SPEA will be that a type $\bar{A} \stackrel{p}{\Downarrow} \bar{B}$ guarantees a probability of termination of the machine of at least p , given an input stack of type \bar{A} , and returning a stack of type \bar{B} . The proof is a direct application of abstract reducibility. For each type we define a set $\text{RUN}(A)$ that holds the terms with the above property, and we proceed to prove that every term of type A belongs to this set.

We write $\mathcal{D}_p(X)$ for the set of finite distributions \mathcal{X} of weight $|\mathcal{X}| \geq p$ over a set X .

► **Definition 23.** *The set $\text{RUN}(-)$ is defined by mutual induction on types A and type vectors \bar{A} as a set of closed terms respectively of closed stacks, as follows.*

$$\begin{aligned} \text{RUN}(\bar{A} \stackrel{p}{\Downarrow} \bar{B}) &= \{ M \mid \forall S \in \text{RUN}(\bar{A}). \exists \mathcal{T} \in \mathcal{D}_p(\text{RUN}(\bar{B})). S, M \Downarrow \mathcal{T} \} \\ \text{RUN}(A_1 \dots A_n) &= \{ \varepsilon M_1 \dots M_n \mid M_i \in \text{RUN}(A_i) \} \end{aligned}$$

For the proof of Lemma 25, the main reducibility lemma, we need the following notation, as well as an additional lemma. We write σ for a substitution map $\{M_i/x_i\}_{i \leq n}$, where σM is the application of σ to M . The set $\text{RUN}(-)$ then extends to contexts Γ as follows. Note that the definition implies that σ is closed (i.e. each substituting term is closed).

$$\text{RUN}(x_1 : A_1, \dots, x_n : A_n) = \{ \sigma \mid \sigma x_i \in \text{RUN}(A_i), 1 \leq i \leq n \}$$

For an event valuation $E = (a_k \mapsto i_k)_{k \leq n}$ we write $\pi_E M$ for the projection on each a_i in E .

$$\pi_E M = \pi_{i_1}^{a_1} (\dots (\pi_{i_n}^{a_n} M) \dots)$$

We *expand* the stacks in a distribution \mathcal{T} by prepending a stack S as $S\mathcal{T}$, where $S[p_i \cdot T_i]_{i \leq n} = [p_i \cdot S T_i]_{i \leq n}$.

► **Lemma 24.** *If $S, M \Downarrow \mathcal{T}$ then $RS, M \Downarrow R\mathcal{T}$ for any stack R .*

Proof. By induction on the definition of \Downarrow . ◀

The main reducibility lemma is then the following.

► **Lemma 25.** *If $E \mid \Gamma \vdash M : A$ and $\sigma \in \text{RUN}(\Gamma)$ then $\pi_E(\sigma M) \in \text{RUN}(A)$.*

Proof. By induction on the typing derivation for M . Note that since σ is closed, $\pi_E(\sigma M) = \sigma(\pi_E(M))$. We cover three key cases (sequencing, abstraction, and generator); the remaining are similar.

Sequencing case:

$$\frac{E \mid \Gamma \vdash N : \bar{A} \stackrel{p}{\Downarrow} \bar{B} \quad E \mid \Gamma \vdash M : \bar{B} \stackrel{q}{\Downarrow} \bar{C}}{E \mid \Gamma \vdash N ; M : \bar{A} \stackrel{pq}{\Downarrow} \bar{C}} \text{seq}$$

31:18 Simple Types for Probabilistic Termination

Let $N' = \pi_E(\sigma N)$ and $M' = \pi_E(\sigma M)$. Given $R \in \text{RUN}(\overline{A})$, by induction we have $R, N' \Downarrow [p_i \cdot S_i]_{i \leq n}$ with each $S_i \in \text{RUN}(\overline{B})$ and $\sum_{i \leq n} p_i \geq p$. Again by induction, for each $i \leq n$ we have $S_i, M' \Downarrow \mathcal{T}_i$ with $\mathcal{T}_i \in \mathcal{D}_q(\text{RUN}(\overline{C}))$. The definition of \Downarrow gives $R, (N'; M') \Downarrow \mathcal{T}$ for $\mathcal{T} = \sum_{i \leq n} p_i \mathcal{T}_i$. Finally, observe that $\mathcal{T} \in \mathcal{D}_{pq}(\text{RUN}(\overline{C}))$ since $|\mathcal{T}| = \sum_{i \leq n} p_i |\mathcal{T}_i| \geq pq$ and since each \mathcal{T}_i is a distribution over $\text{RUN}(\overline{C})$.

Abstraction case:

$$\frac{E \mid \Gamma, x : A \vdash M : \overline{B} \stackrel{p}{\Downarrow} \overline{C}}{E \mid \Gamma \vdash \langle x \rangle . M : A \overline{B} \stackrel{p}{\Downarrow} \overline{C}}_{\text{abs}}$$

Let $\sigma \in \text{RUN}(\Gamma)$ and $S N \in \text{RUN}(\overline{B} A)$. We need to demonstrate a distribution $\mathcal{T} \in \mathcal{D}_p(\text{RUN}(\overline{C}))$ such that $S N, \pi_E(\sigma(\langle x \rangle . M)) \Downarrow \mathcal{T}$. Let $M' = \pi_E(\sigma M)$ and note that since σ is not defined on x and σ and N are closed, $\pi_E(\sigma\{N/x\}M) = \{N/x\}M'$ and $\pi_E(\sigma(\langle x \rangle . M)) = \langle x \rangle . M'$. The inductive hypothesis for the premise $E \mid \Gamma, x : A \vdash M : \overline{B} \stackrel{p}{\Downarrow} \overline{C}$, with $\sigma\{N/x\} \in \text{RUN}(\Gamma, x : A)$, gives the desired $\mathcal{T} \in \mathcal{D}_p(\text{RUN}(\overline{C}))$ with evaluation $S, \{N/x\}M' \Downarrow \mathcal{T}$. Then by the definition of evaluation, $S N, M' \Downarrow \mathcal{T}$, and hence $\langle x \rangle . M' \in \text{RUN}(A \overline{B} \stackrel{p}{\Downarrow} \overline{C})$.

Generator case:

$$\frac{E, a \mapsto \top \mid \Gamma \vdash M : \overline{A} \stackrel{p}{\Downarrow} \overline{C} \quad E, a \mapsto \perp \mid \Gamma \vdash M : \overline{A} \stackrel{q}{\Downarrow} \overline{C}}{E \mid \Gamma \vdash \boxed{a} . M : \overline{A} \stackrel{\frac{1}{2}p + \frac{1}{2}q}{\Downarrow} \overline{C}}_{\text{gen}}$$

Let $\sigma \in \text{RUN}(\Gamma)$ and $S \in \text{RUN}(\overline{A})$. Let $M_i = \pi_{E, a \mapsto i}(\sigma M)$ for $i \in \top, \perp$. By the inductive hypothesis, $M_\top \in \text{RUN}(\overline{A} \stackrel{p}{\Downarrow} \overline{C})$ and $M_\perp \in \text{RUN}(\overline{A} \stackrel{q}{\Downarrow} \overline{C})$, which gives $S, M_\top \Downarrow \mathcal{T}_\top$ and $S, M_\perp \Downarrow \mathcal{T}_\perp$ for some $\mathcal{T}_\top \in \mathcal{D}_p(\text{RUN}(\overline{C}))$ and $\mathcal{T}_\perp \in \mathcal{D}_q(\text{RUN}(\overline{C}))$. Let $r = \frac{1}{2}(p + q)$; then $\mathcal{T} = \frac{1}{2}(\mathcal{T}_\top + \mathcal{T}_\perp) \in \mathcal{D}_r(\text{RUN}(\overline{C}))$ and by definition of evaluation, $S, \pi_E(\sigma(\boxed{a} . M)) \Downarrow \mathcal{T}$. We conclude that $\pi_E(\sigma(\boxed{a} . M)) \in \text{RUN}(\overline{A} \stackrel{r}{\Downarrow} \overline{C})$. \blacktriangleleft

Our main theorem, the probability of machine termination, is a direct consequence.

► **Theorem 21 (restatement).** *For a typed, closed term $M : \overline{A} \stackrel{p}{\Downarrow} \overline{C}$ and stack $S : \overline{A}$ there is a finite set of distinct successful runs*

$$p_i \frac{(S, M, \varepsilon)}{(T_i, \star, \varepsilon)} \quad (i \leq n)$$

with sum probability $\sum_{i \leq n} p_i \geq p$.

Proof. By Lemma 25 we have $M \in \text{RUN}(\overline{A} \stackrel{p}{\Downarrow} \overline{C})$ and $S \in \text{RUN}(\overline{A})$. By the definition of $\text{RUN}(-)$ we then have $S, M \Downarrow \mathcal{T}$ with $|\mathcal{T}| \geq p$. Proposition 17 then gives the desired set of machine runs. \blacktriangleleft

A.2 Probabilistic head normalization

Finally, we will relate machine termination to head reduction. For a redex in a head context $\lambda x_1 \dots \lambda x_n . \{ \} M_1 \dots M_m$ the machine runs as follows: after the abstractions consume the top n elements off the stack, and the applications push the terms M_i onto it, then the redex itself is the first part of the term to be evaluated on the machine. Machine evaluation thus corresponds tightly to head reduction, with the same order of evaluation of redexes.

However, in this correspondence, the machine halts with a *variable*, while successful termination in our setting requires a *skip* with an empty continuation stack. We will thus use a different approach.

$$\begin{array}{c}
\frac{}{S, \star \Downarrow_0 [1 \cdot S]} \quad \frac{S, \{N/x\}M \Downarrow_w \mathcal{T}}{S N, \langle x \rangle. M \Downarrow_{w+1} \mathcal{T}} \quad \frac{R, M \Downarrow_w [p_i \cdot S_i]_{i \leq n} \quad (S_i, N \Downarrow_{v_i} \mathcal{T}_i)_{i \leq n}}{R, M ; N \Downarrow_{(w + \sum_{i \leq n} v_i)} \sum_{i \leq n} p_i \mathcal{T}_i} \\
\frac{}{S, M \Downarrow_0 \emptyset} \quad \frac{S N, M \Downarrow_w \mathcal{T}}{S, [N]. M \Downarrow_{w+1} \mathcal{T}} \quad \frac{S, \pi_{\top}^{\alpha}(M) \Downarrow_v \mathcal{T}_{\top} \quad S, \pi_{\perp}^{\alpha}(M) \Downarrow_w \mathcal{T}_{\perp}}{S, \boxed{a}. M \Downarrow_{v+w+1} \frac{1}{2} \mathcal{T}_{\top} + \frac{1}{2} \mathcal{T}_{\perp}}
\end{array}$$

■ **Figure 5** The weighted probabilistic evaluation relation.

The main idea is that reduction *shortens* the runs of the machine, by removing consecutive *push* and *pop* operations, or in the case of projective reduction, removing a generator. By annotating the evaluation relation to count abstractions, applications, and generators we may then observe that this measure reduces under head reduction.

► **Definition 26.** The weighted evaluation relation $S, M \Downarrow_w \mathcal{T}$ is given in Figure 5, where the weight w is a natural number. We extend it to a distribution of terms M by the following rule, carrying a multiset of weights \mathcal{W} .

$$\frac{(S, M_i \Downarrow_{w_i} \mathcal{T}_i)_{i \leq n}}{S, [p_i \cdot M_i]_{i \leq n} \Downarrow_{[w_i]_{i \leq n}} \sum_{i \leq n} p_i \mathcal{T}_i}$$

The core lemma establishes that the weight in the evaluation relation decreases for $\rightarrow_{\beta h}$ and \rightarrow_{π} reduction steps, and is stable under $\rightarrow_{\sigma h}$. However, this does not apply to terms introduced by the zero-rule, as $S, M \Downarrow_0 \emptyset$, which are the potentially non-terminating terms. Crucially for our purpose, when the evaluation returns a non-empty distribution, head-reduction on the term progresses towards a head-normal form.

► **Lemma 27.** Let $S, M \Downarrow_w \mathcal{T}$ where $\mathcal{T} \neq \emptyset$.

- If $M \rightarrow_{\sigma h} N$ then $S, N \Downarrow_w \mathcal{T}$.
- If $M \rightarrow_{\beta h} N$ then $S, N \Downarrow_v \mathcal{T}$ with $v < w$.
- If $M \rightarrow_{\pi} \mathcal{N}$ then $S, \mathcal{N} \Downarrow_{\mathcal{W}} \mathcal{T}$ with $v < w$ for every v in \mathcal{W} .

Proof. We first consider the reduction steps themselves, and then their closure under head contexts. The proof is by induction on \Downarrow . We consider three cases; the remaining are similar.

Beta $[N]. \langle x \rangle. M \rightarrow_{\beta} \{N/x\}M$ Since \mathcal{T} is non-empty, the derivation must be as below left, as none of the inferences may be replaced by a zero-rule. The case is then as follows.

$$\frac{\frac{S, \{N/x\}M \Downarrow_w \mathcal{T}}{S N, \langle x \rangle. M \Downarrow_{w+1} \mathcal{T}}}{S, [N]. \langle x \rangle. M \Downarrow_{w+2} \mathcal{T}} \rightarrow_{\beta} S, \{N/x\}M \Downarrow_w \mathcal{T}$$

Projection $\boxed{a}. M \rightarrow_{\pi} [\frac{1}{2} \cdot \pi_{\top}^{\alpha} M, \frac{1}{2} \cdot \pi_{\perp}^{\alpha} M]$ Since \mathcal{T} is non-empty, the derivation for the redex is the first below. The derivation for the reduct, second below, uses the evaluation rule for distributions.

$$\frac{S, \pi_{\top}^{\alpha} M \Downarrow_v \mathcal{T}_{\top} \quad S, \pi_{\perp}^{\alpha} M \Downarrow_w \mathcal{T}_{\perp}}{S, \boxed{a}. M \Downarrow_{v+w+1} \frac{1}{2} \mathcal{T}_{\top} + \frac{1}{2} \mathcal{T}_{\perp}}$$

$$\frac{S, \pi_{\top}^{\alpha} M \Downarrow_v \mathcal{T}_{\top} \quad S, \pi_{\perp}^{\alpha} M \Downarrow_w \mathcal{T}_{\perp}}{S, [\frac{1}{2} \cdot \pi_{\top}^{\alpha} M, \frac{1}{2} \cdot \pi_{\perp}^{\alpha} M] \Downarrow_{[v,w]} \frac{1}{2} \mathcal{T}_{\top} + \frac{1}{2} \mathcal{T}_{\perp}}$$

31:20 Simple Types for Probabilistic Termination

Sequence-generator $\boxed{a}.M;P \rightarrow_\sigma \boxed{a}.(M;P)$ ($a \notin \text{fe}(P)$) The derivations are as follows, with premises stacked for space. In the second derivation $w' = w + \sum_{i \leq n} u_i$ and $v' = v + \sum_{n < i \leq m} u_i$.

$$\frac{\frac{R, \pi_\top^\alpha M \Downarrow_w [p_i \cdot S_i]_{i \leq n} \quad R, \pi_\perp^\alpha M \Downarrow_v [p_i \cdot S_i]_{n < i \leq m}}{R, \boxed{a}.M \Downarrow_{w+v+1} \frac{1}{2} [p_i \cdot S_i]_{i \leq m} \quad (S_i, P \Downarrow_{u_i} \mathcal{T}_i)_{i \leq m}}}{R, \boxed{a}.M;P \Downarrow_{(w+v+1 + \sum_{i \leq m} u_i)} \sum_{i \leq m} \frac{1}{2} p_i \mathcal{T}_i}$$

$$\frac{\frac{R, \pi_\top^\alpha M \Downarrow_w [p_i \cdot S_i]_{i \leq n} \quad (S_i, P \Downarrow_{u_i} \mathcal{T}_i)_{i \leq n}}{R, \pi_\top^\alpha M;P \Downarrow_{w'} \sum_{i \leq n} p_i \mathcal{T}_i} \quad \frac{R, \pi_\perp^\alpha M \Downarrow_v [p_i \cdot S_i]_{n < i \leq m} \quad (S_i, P \Downarrow_{u_i} \mathcal{T}_i)_{n < i \leq m}}{R, \pi_\perp^\alpha M;P \Downarrow_{v'} \sum_{n < i \leq m} p_i \mathcal{T}_i}}{R, \boxed{a}.(M;P) \Downarrow_{(w+v+1 + \sum_{i \leq m} u_i)} \sum_{i \leq m} \frac{1}{2} p_i \mathcal{T}_i}$$

This concludes the reduction steps. The remaining cases consider the closure under head contexts. For sequencing reduction, which leaves the weight unchanged, this is immediate, and the cases for beta-steps are straightforward and omitted. For projective reduction, let $M \rightarrow_\pi [\frac{1}{2} \cdot N_\top \frac{1}{2} \cdot N_\perp]$, and consider the remaining cases.

Application $[P].M \rightarrow_\pi [\frac{1}{2} \cdot [P].N_\top \frac{1}{2} \cdot [P].N_\perp]$ The derivation for $[P].M$ is as follows.

$$\frac{S P, M \Downarrow_w \mathcal{T}}{S, [P].M \Downarrow_{w+1} \mathcal{T}}$$

By induction, for the premise of this derivation we get one for the distribution $\mathcal{N} = [\frac{1}{2} \cdot N_\top, \frac{1}{2} \cdot N_\perp]$, below, where it follows that $\mathcal{T} = \frac{1}{2} \mathcal{T}_\top + \frac{1}{2} \mathcal{T}_\perp$ and $u, v < w$.

$$\frac{S P, N_\top \Downarrow_u \mathcal{T}_\top \quad S P, N_\perp \Downarrow_v \mathcal{T}_\perp}{S P, \mathcal{N} \Downarrow_{[u,v]} \mathcal{T}}$$

Then for $\mathcal{P} = [\frac{1}{2} \cdot [P].N_\top \frac{1}{2} \cdot [P].N_\perp]$ we get the following derivation.

$$\frac{\frac{S P, N_\top \Downarrow_u \mathcal{T}_\top}{S, [P].N_\top \Downarrow_{u+1} \mathcal{T}_\top} \quad \frac{S P, N_\perp \Downarrow_v \mathcal{T}_\perp}{S, [P].N_\perp \Downarrow_{v+1} \mathcal{T}_\perp}}{S, \mathcal{P} \Downarrow_{[u+1, v+1]} \mathcal{T}}$$

Abstraction $\langle x \rangle.M \rightarrow_\pi [\frac{1}{2} \cdot \langle x \rangle.N_\top \frac{1}{2} \cdot \langle x \rangle.N_\perp]$ The derivation for $\langle x \rangle.M$ is as follows.

$$\frac{S, \{P/x\}M \Downarrow_w \mathcal{T}}{S P, \langle x \rangle.M \Downarrow_{w+1} \mathcal{T}}$$

Given the reduction for M , we also have the following, where we abbreviate the reduct as \mathcal{P} .

$$\{P/x\}M \rightarrow_\pi [\frac{1}{2} \cdot \{P/x\}N_\top \frac{1}{2} \cdot \{P/x\}N_\perp] = \mathcal{P}$$

By induction this gives us the following derivation for \mathcal{P} , where again $\mathcal{T} = \frac{1}{2} \mathcal{T}_\top + \frac{1}{2} \mathcal{T}_\perp$ and $u, v < w$.

$$\frac{S, \{P/x\}N_\top \Downarrow_u \mathcal{T}_\top \quad S, \{P/x\}N_\perp \Downarrow_v \mathcal{T}_\perp}{S, \mathcal{P} \Downarrow_{[u,v]} \mathcal{T}}$$

Then for $\mathcal{N} = [\frac{1}{2} \cdot \langle x \rangle . N_{\top} \frac{1}{2} \cdot \langle x \rangle . N_{\perp}]$ we get the required derivation.

$$\frac{\frac{S, \{P/x\}N_{\top} \Downarrow_u \mathcal{T}_{\top}}{S P, \langle x \rangle . N_{\top} \Downarrow_{u+1} \mathcal{T}_{\top}} \quad \frac{S, \{P/x\}N_{\perp} \Downarrow_v \mathcal{T}_{\perp}}{S P, \langle x \rangle . N_{\perp} \Downarrow_{v+1} \mathcal{T}_{\perp}}}{S, \mathcal{N} \Downarrow_{[u+1, v+1]} \mathcal{T}} \quad \blacktriangleleft$$

We apply the above lemma to relate machine evaluation to head reduction in the following way: if evaluation returns a distribution of weight p , then head-reduction terminates with probability at least p .

► **Lemma 28.** *If $S, \mathcal{M} \Downarrow_{\mathcal{W}} \mathcal{T}$ then $\mathcal{M} \rightarrow_{\text{h}} \mathcal{N}_0 + \mathcal{N}_1$ where all terms in \mathcal{N}_0 are head normal and $|\mathcal{N}_0| \geq |\mathcal{T}|$.*

Proof. Let $\mathcal{M} = [p_i \cdot M_i]_{i \leq n}$ evaluate by the following derivation.

$$\frac{(S, M_i \Downarrow_{w_i} \mathcal{T}_i)_{i \leq n}}{S, [p_i \cdot M_i]_{i \leq n} \Downarrow_{[w_i]_{i \leq n}} \sum_{i \leq n} p_i \mathcal{T}_i}$$

First, we head-reduce any M_i where $\mathcal{T}_i \neq \emptyset$. This process terminates by induction on \mathcal{W} , using Lemma 27: a β -step or projective step reduces \mathcal{W} , while sequencing reduction alone is terminating and does not increase it.

Then the distribution $\mathcal{N}_0 = [p_i \cdot M_i \mid i \leq n, \mathcal{T}_i \neq \emptyset]$ is head-normal. The weights of \mathcal{N}_0 and \mathcal{T} are as follows.

$$|\mathcal{N}_0| = \sum [p_i \mid i \leq n, \mathcal{T}_i \neq \emptyset] \quad |\mathcal{T}| = \sum_{i \leq n} p_i |\mathcal{T}_i|$$

It follows that $|\mathcal{N}_0| \geq |\mathcal{T}|$. ◀

Our final theorem then ties everything together: typing gives successful evaluation on the machine, which in turn gives termination of head reduction.

► **Theorem 22 (restatement).** *For a closed M , if $M : \bar{A} \stackrel{p}{\rightarrow} \bar{C}$ then $M \rightarrow_{\text{h}} \mathcal{N}_0 + \mathcal{N}_1$ where all terms in \mathcal{N}_0 are head normal and $|\mathcal{N}_0| \geq p$.*

Proof. We provide M with an input stack consisting only of zero terms 0_A . Let S be the stack of zero terms for \bar{A} . By Proposition 20 zero-terms inhabit their types, so that $S : \bar{A}$.

By Lemma 25 we have $S \in \text{RUN}(\bar{A})$ and $M \in \text{RUN}(\bar{A} \stackrel{p}{\rightarrow} \bar{C})$. By the definition of $\text{RUN}(-)$ this gives an evaluation $S, M \Downarrow \mathcal{T}$ where $|\mathcal{T}| \geq p$.

For the corresponding weighted derivation $S, M \Downarrow_w \mathcal{T}$, Lemma 28 gives a head reduction $M \rightarrow_{\text{h}} \mathcal{N}_0 + \mathcal{N}_1$ where all terms in \mathcal{N}_0 are head normal and $|\mathcal{N}_0| \geq |\mathcal{T}| \geq p$. ◀