

On Higher-Order Probabilistic Subrecursion

Flavien Breuvert, Ugo Dal Lago, Agathe Herrou

► **To cite this version:**

Flavien Breuvert, Ugo Dal Lago, Agathe Herrou. On Higher-Order Probabilistic Subrecursion. FoS-SaCS 2017 - 20th International Conference on Foundations of Software Science and Computation Structures , Apr 2017, Uppsala, Sweden. 10203, pp.370-386, LNCS. <10.1007/978-3-662-54458-7_22>. <hal-01639379>

HAL Id: hal-01639379

<https://hal.inria.fr/hal-01639379>

Submitted on 20 Nov 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

On Higher-Order Probabilistic Subrecursion^{*}

Flavien Breuvert¹, Ugo Dal Lago² and Agathe Herrou³

¹ INRIA Sophia Antipolis, flavien.breuvert@inria.fr

² University of Bologna & INRIA Sophia Antipolis, ugo.dallago@unibo.it

³ ENS de Lyon, agathe.herrou@ens-lyon.fr

Abstract. We study the expressive power of subrecursive probabilistic higher-order calculi. More specifically, we show that endowing a very expressive deterministic calculus like Gödel's \mathbb{T} with various forms of probabilistic choice operators may result in calculi which are not equivalent as for the class of distributions they give rise to, although they all guarantee *almost-sure* termination. Along the way, we introduce a probabilistic variation of the classic reducibility technique, and we prove that the simplest form of probabilistic choice leaves the expressive power of \mathbb{T} essentially unaltered. The paper ends with some observations about functional expressivity: expectedly, all the considered calculi represent precisely the functions which \mathbb{T} itself represents.

1 Introduction

Probabilistic models are more and more pervasive in computer science and are among the most powerful modeling tools in many areas like computer vision [20], machine learning [19] and natural language processing [17]. Since the early times of computation theory [8], the very concept of an algorithm has been itself generalised from a purely deterministic process to one in which certain elementary computation steps can have a probabilistic outcome. This has further stimulated research in computation and complexity theory [11], but also in programming languages [21].

Endowing programs with probabilistic primitives (e.g. an operator which models sampling from a distribution) poses a challenge to programming language semantics. Already for a minimal, imperative probabilistic programming language, giving a denotational semantics is nontrivial [16]. When languages also have higher-order constructs, everything becomes even harder [14] to the point of disrupting much of the beautiful theory known in the deterministic case [1]. This has stimulated research on denotational semantics of higher-order probabilistic programming languages, with some surprising positive results coming out recently [9, 4].

Not much is known about the expressive power of *probabilistic* higher-order calculi, as opposed to the extensive literature on the same subject about *deterministic* calculi (see, e.g. [24, 23]). What happens to the class of representable

^{*} The authors are partially supported by ANR project 14CE250005 ELICA and ANR project 12IS02001 PACE.

functions if one enriches, say, a deterministic λ -calculus \mathbb{X} with certain probabilistic choice primitives? Are the expressive power or the good properties of \mathbb{X} somehow preserved? These questions have been given answers in the case in which \mathbb{X} is the pure, untyped, λ -calculus [6]: in that case, universality continues to hold, mimicking what happens in Turing machines [22]. But what if \mathbb{X} is one of the many typed λ -calculi ensuring strong normalisation for typed terms [12]?

Let us do a step back, first: when should a higher-order probabilistic program be considered terminating? The question can be given a satisfactory answer being inspired by, e.g., recent works on probabilistic termination in imperative languages and term rewrite systems [18, 2]: one could ask the probability of divergence to be 0, i.e., *almost sure termination*, or the stronger *positive almost sure termination*, in which one requires the average number of evaluation steps to be finite. That almost sure termination is a desirable property, even in a probabilistic setting can be seen in the field of languages like CHURCH and ANGLICAN, in which programs are often assumed to be almost surely terminating, e.g. when doing inference by MH algorithms [13].

In this paper, we initiate a study on the expressive power of terminating higher-order calculi, in particular those obtained by endowing Gödel's \mathbb{T} with various forms of probabilistic choice operators. In particular, three operators will be analysed in this paper:

- A binary probabilistic operator \oplus such that for every pair of terms M, N , the term $M \oplus N$ evaluates to either M or N , each with probability $\frac{1}{2}$. This is a rather minimal option which, however, guarantees universality if applied to the untyped λ -calculus [6] (and, more generally, to universal models of computation [22]).
- A combinator \mathbf{R} , which evaluates to any natural number $n \geq 0$ with probability $\frac{1}{2^{n+1}}$. This is the natural generalisation of \oplus to sampling from a distribution having *countable* rather than *finite* support. This apparently harmless generalisation (which is absolutely non-problematic in a universal setting) has dramatic consequences in a subrecursive scenario, as we will discover soon.
- A combinator \mathbf{X} such that for every pair of values V, W , the term $\mathbf{X}\langle V, W \rangle$ evaluates to either W or $V(\mathbf{X}\langle V, W \rangle)$, each with probability $\frac{1}{2}$. The operator \mathbf{X} can be seen as a probabilistic variation on \mathbb{PCF} 's fixpoint combinator. As such, \mathbf{X} is potentially problematic to termination, giving rise to infinite trees.

This way, various calculi can be obtained, like \mathbb{T}^\oplus , namely a minimal extension of \mathbb{T} , or the full calculus $\mathbb{T}^{\oplus, \mathbf{R}, \mathbf{X}}$, in which the three operators are all available. In principle, the only obvious fact about the expressive power of the above mentioned operators is that both \mathbf{R} and \mathbf{X} are at least as expressive as \oplus : binary choice can be easily expressed by either \mathbf{R} or \mathbf{X} . Less obvious but still easy to prove is the equivalence between \mathbf{R} and \mathbf{X} in presence of a recursive operator (see Section 3.3). But how about, say, \mathbb{T}^\oplus vs. $\mathbb{T}^{\mathbf{R}}$?

Traditionally, the expressiveness of such languages is evaluated by looking at the set of functions $f : \mathbb{N} \rightarrow \mathbb{N}$ defined by typable programs $M : \mathbf{NAT} \rightarrow \mathbf{NAT}$. However, in a probabilistic setting, any program $M : \mathbf{NAT} \rightarrow \mathbf{NAT}$ computes a function from natural numbers to *distributions* of natural numbers. In order to

fit usual criteria, we need to fix a notion of observation of which there are at least two, corresponding to two randomised programming paradigms, namely *Las Vegas* and *Monte Carlo* observations. The main question, then, consists in understanding how the obtained classes relate to each other, and to the class of \mathbb{T} -representable functions. Along the way, however, we manage to understand how to capture the expressive power of probabilistic calculi *per se*. This paper’s contributions can be summarised as follows:

- We first take a look at the full calculus $\mathbb{T}^{\oplus, \mathbb{R}, \mathbb{X}}$, and prove that it enforces almost-sure termination, namely that the probability of termination of any typable term is 1. This is done by appropriately adapting the well-known reducibility technique [12] to a probabilistic operational semantics. We then observe that while $\mathbb{T}^{\oplus, \mathbb{R}, \mathbb{X}}$ cannot be *positively* almost surely terminating, \mathbb{T}^{\oplus} indeed is. This already shows that there must be a gap in expressivity. This is done in Section 3.
- In Section 4, we look more precisely at the expressive power of \mathbb{T}^{\oplus} , proving that the mere presence of probabilistic choice does not add much to the expressive power of \mathbb{T} : in a sense, probabilistic choice can be “lifted up” to the ambient deterministic calculus.
- We look at other fragments of $\mathbb{T}^{\oplus, \mathbb{R}, \mathbb{X}}$ and at their expressivity. More specifically, we prove that (the equiexpressive) $\mathbb{T}^{\mathbb{R}}$ and $\mathbb{T}^{\mathbb{X}}$ represent precisely what \mathbb{T}^{\oplus} can do *at the limit*, in a sense which will be made precise in Section 3. This result, which is the most challenging, is given in Section 5.
- Section 6 is devoted to proving that both for *Monte Carlo* and for *Las Vegas* observations, the class of representable functions of $\mathbb{T}^{\mathbb{R}}$ coincides with the \mathbb{T} -representable ones.

Due to lack of space, most proofs are elided. An extended version of this paper with more details is available [3].

2 Probabilistic Choice Operators, Informally

Any term of Gödel’s \mathbb{T} can be seen as a purely deterministic computational object whose dynamics is finitary, due to the well-known strong normalisation theorem (see, e.g., [12]). In particular, the apparent non-determinism due to multiple redex occurrences is completely harmless because of confluence. In this paper, indeed, we even neglect this problem, and work with a reduction strategy, namely weak call-by-value reduction (keeping in mind that all what we will say also holds in call-by-name). Evaluation of a \mathbb{T} -term M of type \mathbb{NAT} can be seen as a finite sequence of terms ending in the normal form \mathbf{n} of M (see Figure 1). More generally, the unique normal form of any \mathbb{T} term M will be denoted as $\llbracket M \rrbracket$. Noticeably, \mathbb{T} is computationally very powerful. In particular, the \mathbb{T} -representable functions from \mathbb{N} to \mathbb{N} coincide with the functions which are provably total in Peano’s arithmetic [12].

As we already mentioned, the most natural way to enrich deterministic calculi and turn them into probabilistic ones consists in endowing their syntax with one or more probabilistic choice operators. Operationally, each of them models

the essentially stochastic process of sampling from a distribution and proceeding depending on the outcome. Of course, one has many options here as for *which* of the various operators to grab. The aim of this work is precisely to study to which extent this choice have consequences on the overall expressive power of the underlying calculus.

Suppose, for example, that \mathbb{T} is endowed with the binary probabilistic choice operator \oplus described in the Introduction, whose evaluation corresponds to tossing a fair coin and choosing one of the two arguments accordingly. The presence of \oplus has indeed an impact on the dynamics of the underlying calculus: the evaluation of any term M is not deterministic anymore, but can be modelled as a finitely branching tree (see, e.g., Figure 3 for such a tree). The fact that all branches of this tree have finite height (and the tree is thus finite) is intuitive, and a proof of it can be given by adapting the well-known reducibility proof of termination for \mathbb{T} . In this paper, we in fact prove much more, and establish that \mathbb{T}^\oplus can be embedded into \mathbb{T} .

If \oplus is replaced by \mathbf{R} , the underlying tree is not finitely branching anymore, but, again, there is not (at least apparently) any infinitely long branch, since each of them can somehow be seen as a \mathbb{T} computation (see Figure 2 for an example). What happens to the expressive power of the obtained calculus? Intuition tells us that the calculus should not be too expressive viz. \mathbb{T}^\oplus . If \oplus is replaced by \mathbf{X} , on the other hand, the underlying tree *is* finitely branching, but its height can be infinite. Actually, \mathbf{X} and \mathbf{R} are easily shown to be equiexpressive in presence of higher-order recursion, as we show in Section 3.3. On the other hand, for \mathbf{R} and \oplus , no such encoding is available. Nonetheless, $\mathbb{T}^{\mathbf{R}}$ can still be somehow encoded into \mathbb{T} (the embedding being correct only “at the limit”) as we will detail in Section 5. From this embedding, we can show that applying Monte Carlo or Las Vegas algorithms to $\mathbb{T}^{\oplus, \mathbf{X}, \mathbf{R}}$ do not add any expressive power to that \mathbb{T} . This is done in Section 6.

3 The Full Calculus $\mathbb{T}^{\oplus, \mathbf{R}, \mathbf{X}}$

All along this paper, we work with a calculus $\mathbb{T}^{\oplus, \mathbf{R}, \mathbf{X}}$ whose *terms* are the ones generated by the following grammar:

$$M, N, L ::= x \mid \lambda x.M \mid M N \mid \langle M, N \rangle \mid \pi_1 \mid \pi_2 \\ \mid \mathbf{rec} \mid \mathbf{0} \mid \mathbf{S} \mid M \oplus N \mid \mathbf{R} \mid \mathbf{X}.$$

Please observe the presence of the usual constructs from the untyped λ -calculus, but also of primitive recursion, constants for natural numbers, pairs, and the three choice operators we have described in the previous sections.

As usual, terms are taken modulo α -equivalence. Terms in which no variable occurs free are said *closed*, and are collected in the set $\mathbb{T}_C^{\oplus, \mathbf{R}, \mathbf{X}}$. A *value* is simply a closed term from the following grammar:

$$U, V ::= \lambda x.M \mid \langle U, V \rangle \mid \pi_1 \mid \pi_2 \mid \mathbf{rec} \mid \mathbf{0} \mid \mathbf{S} \mid \mathbf{S} V \mid \mathbf{X}.$$

and the set of values is $\mathbb{T}_V^{\oplus, \mathbf{R}, \mathbf{X}}$. *Extended values* are (not necessarily closed) terms generated by the same grammar as values with the addition of variables. Closed terms that are not values are called *reducible* and their set is denoted $\mathbb{T}_R^{\oplus, \mathbf{R}, \mathbf{X}}$. The expression $\langle M_1, \dots, M_n \rangle$ stands for $\langle M_1, \langle M_2, \dots \rangle \rangle$. A *context* is a term with a unique hole:

$$C := (\cdot) \mid \lambda x. C \mid C M \mid M C \mid \langle C, M \rangle \mid \langle M, C \rangle \mid C \oplus M \mid M \oplus C.$$

We write $\mathbb{T}_{(\cdot)}^{\oplus, \mathbf{R}, \mathbf{X}}$ for the set of all such contexts.

Termination of Gödel's \mathbb{T} is guaranteed by the presence of types, which we also need here. *Types* are expressions generated by the following grammar

$$A, B ::= \text{NAT} \mid A \rightarrow B \mid A \times B.$$

Environmental contexts are expressions of the form $\Gamma = x_1 : A_1, \dots, x_n : A_n$, while *typing judgments* are of the form $\Gamma \vdash M : A$. *Typing rules* are given in Figure 5. From now on, only typable terms will be considered. We denote by $\mathbb{T}^{\oplus, \mathbf{R}, \mathbf{X}}(A)$ the set of terms of type A , and similarly for $\mathbb{T}_C^{\oplus, \mathbf{R}, \mathbf{X}}(A)$ and $\mathbb{T}_V^{\oplus, \mathbf{R}, \mathbf{X}}(A)$. We use the shortcut \mathbf{n} for values of type NAT: $\mathbf{0}$ is already part of the language of terms, while $\mathbf{n} + \mathbf{1}$ is simply $\mathbf{S} \mathbf{n}$.

3.1 Operational Semantics

While evaluating terms in a deterministic calculus ends up in a *value*, the same process leads to a *distribution* of values when performed on terms in a probabilistic calculus. Formalising all this requires some care, but can be done following one of the many definitions from the literature (e.g., [6]).

Given a countable set X , a *distribution* \mathcal{L} on X is a probabilistic subdistribution over elements of X :

$$\mathcal{L}, \mathcal{M}, \mathcal{N} \in \mathfrak{D}(X) = \left\{ f : X \rightarrow [0, 1] \mid \sum_{x \in X} f(x) \leq 1 \right\}.$$

We are especially concerned with distributions over terms here. In particular, a *distribution of type* A is simply an element of $\mathfrak{D}(\mathbb{T}^{\oplus, \mathbf{R}, \mathbf{X}}(A))$. The set $\mathfrak{D}(\mathbb{T}_V^{\oplus, \mathbf{R}, \mathbf{X}})$ is ranged over by metavariables like $\mathcal{U}, \mathcal{V}, \mathcal{W}$. We will use the pointwise order \leq on distributions, which turns them into an ω CPO. Moreover, we use the following notation for Dirac's distributions over terms: $\{M\} := \left\{ \begin{array}{l} M \mapsto 1 \\ N \mapsto 0 \text{ if } M \neq N \end{array} \right\}$. The support of a distribution is indicated as $|\mathcal{M}|$; we also define the reducible and value supports fragments as $|\mathcal{M}|_R := |\mathcal{M}| \cap \mathbb{T}_R^{\oplus, \mathbf{R}, \mathbf{X}}$ and $|\mathcal{M}|_V := |\mathcal{M}| \cap \mathbb{T}_V^{\oplus, \mathbf{R}, \mathbf{X}}$. Notions like \mathcal{M}^R and \mathcal{M}^V have an obvious and natural meaning: for any $\mathcal{M} \in \mathfrak{D}(X)$ and $Y \subseteq X$, then $\mathcal{M}^Y(x) = \mathcal{M}(x)$ if $x \in \mathbb{T}_Y^{\oplus, \mathbf{R}, \mathbf{X}}$ and $\mathcal{M}^Y(x) = 0$ otherwise.

As syntactic sugar, we use integral notations to manipulate distributions, *i.e.*, for any family of distributions $(\mathcal{N}_M)_{M \in \mathbb{T}^{\oplus, \mathbf{R}, \mathbf{X}}} : \mathfrak{D}(\mathbb{T}^{\oplus, \mathbf{R}, \mathbf{X}})^{\mathbb{T}^{\oplus, \mathbf{R}, \mathbf{X}}}$, the expression $\int_{\mathcal{M}} \mathcal{N}_M . dM$ stands for $\sum_{M \in \mathbb{T}^{\oplus, \mathbf{R}, \mathbf{X}}} \mathcal{M}(M) \cdot \mathcal{N}_M$ (by abuse of notation, we may define \mathcal{N}_M only for $M \in |\mathcal{M}|$, since the others are not used anyway). The notation

can be easily adapted, e.g., to families of real numbers $(p_M)_{M \in \mathbb{T}^{\oplus, \mathbb{R}, \mathbb{X}}}$ and to other kinds of distributions. We indicate as $C(\llbracket \mathcal{M} \rrbracket)$ the push-forward distribution $\int_{\mathcal{M}} \{C(\llbracket M \rrbracket)\} dM$ induced by a context C , and as $\sum \mathcal{M}$ the norm $\int_{\mathcal{M}} 1 dM$ of \mathcal{M} . Remark, finally, that we have the useful equality $\mathcal{M} = \int_{\mathcal{M}} \{M\} dM$.

Reduction rules of $\mathbb{T}^{\oplus, \mathbb{R}, \mathbb{X}}$ are given by Figure 6. For reasons of simplicity, the relation \rightarrow indicates both a subset of $\mathbb{T}_C^{\oplus, \mathbb{R}, \mathbb{X}} \times \mathfrak{D}(\mathbb{T}_C^{\oplus, \mathbb{R}, \mathbb{X}})$ and a relation on $\mathfrak{D}(\mathbb{T}_C^{\oplus, \mathbb{R}, \mathbb{X}}) \times \mathfrak{D}(\mathbb{T}_C^{\oplus, \mathbb{R}, \mathbb{X}})$. Notice that the reduction \rightarrow is deterministic. We can easily define \rightarrow^n as the n^{th} exponentiation of \rightarrow and \rightarrow^* as the reflexive and transitive closure of \rightarrow taking the latter as a relation on distributions. In probabilistic systems, we might want to consider infinite reductions such as the ones induced by $\mathbb{X}(\lambda x.x, \mathbf{0})$, which reduces to $\{\mathbf{0}\}$, but in an infinite number of steps. Remark that for any value V , and whenever $\mathcal{M} \rightarrow \mathcal{N}$, it holds that $\mathcal{M}(V) \leq \mathcal{N}(V)$. As a consequence, we can proceed as follows:

Definition 1. Let M be a term and let $(\mathcal{M}_n)_{n \in \mathbb{N}}$ be the unique distribution family such that $M \rightarrow^n \mathcal{M}_n$. The evaluation of M is the value distribution

$$\llbracket M \rrbracket := \{V \mapsto \lim_{n \rightarrow \infty} \mathcal{M}_n(V)\} \in \mathfrak{D}(\mathbb{T}_V^{\oplus, \mathbb{R}, \mathbb{X}}).$$

The success of M is its probability of normalisation, which is formally defined as the norm of its evaluation, i.e., $\text{Succ}(M) := \sum \llbracket M \rrbracket$. $\mathcal{M}_n^{\Delta V}$ stands for $\{V \mapsto \mathcal{M}_n(V) - \mathcal{M}_{n-1}(V)\}$, the distributions of values reachable in exactly n steps. The average reduction length from M is then

$$[M] := \sum_n \left(n \cdot \sum \mathcal{M}_n^{\Delta V} \right) \in \mathbb{N} \cup \{+\infty\}$$

Notice that, by Rule (r- ϵ), the evaluation is continuous: $\llbracket \mathcal{M} \rrbracket = \int_{\mathcal{M}} \llbracket M \rrbracket dM$. Any closed term M of type $\text{NAT} \rightarrow \text{NAT}$ represents a function $g : \mathbb{N} \rightarrow \mathfrak{D}(\mathbb{N})$ iff for every n, m it holds that $g(n)(m) = \llbracket M \mathbf{n} \rrbracket(\mathbf{m})$.




3.2 Almost-Sure Termination

We now have all the necessary ingredients to specify a quite powerful notion of probabilistic computation. When, precisely, should such a process be considered *terminating*? Do all probabilistic branches (see figures 1-4) need to be finite? Or should we stay more liberal? The literature on the subject is pointing to the notion of *almost-sure termination*: a probabilistic computation should be considered terminating if the set of infinite computation branches, although not necessarily empty, has null probability [18, 10, 15]. This has the following incarnation in our setting:

Definition 2. A term M is said to be almost-surely terminating (AST) iff $\text{Succ}(M) = 1$.

This section is concerned with proving that $\mathbb{T}^{\oplus, \mathbb{R}, \mathbb{X}}$ indeed guarantees almost-sure termination. This will be done by adapting Girard-Tait's reducibility technique.

The following is a crucial intermediate step towards Theorem 1, the main result of this section.



Lemma 1. For any M, N , it holds that $\llbracket M N \rrbracket = \llbracket \llbracket M \rrbracket \llbracket N \rrbracket \rrbracket$. In particular, if the application $M N$ is almost-surely terminating, so are M and N .

Theorem 1. The full system $\mathbb{T}^{\oplus, \mathbf{R}, \mathbf{X}}$ is almost-surely terminating (AST), i.e.,

$$\forall M \in \mathbb{T}^{\oplus, \mathbf{R}, \mathbf{X}}, \quad \text{Succ}(M) = 1.$$

Proof. The proof⁴ is based on the notion of a *reducible* term which is given as follows by induction on the structure of types:

$$\begin{aligned} \text{Red}_{\text{NAT}} &:= \{M \in \mathbb{T}^{\oplus, \mathbf{R}, \mathbf{X}}(\text{NAT}) \mid M \text{ is AST}\}; \\ \text{Red}_{A \rightarrow B} &:= \{M \mid \forall V \in \text{Red}_A \cap \mathbb{T}_V^{\oplus, \mathbf{R}, \mathbf{X}}, (M V) \in \text{Red}_B\}; \\ \text{Red}_{A \times B} &:= \{M \mid (\pi_1 M) \in \text{Red}_A, (\pi_2 M) \in \text{Red}_B\}. \end{aligned}$$

Then we can observe that:

- *The reducibility candidates over Red_A are \rightarrow -saturated:* by induction on A we can indeed show that if $M \rightarrow \mathcal{M}$ then $|\mathcal{M}| \subseteq \text{Red}_A$ iff $M \in \text{Red}_A$.
- *The reducibility candidates over Red_A are precisely the AST terms M such that $\llbracket M \rrbracket \subseteq \text{Red}_A$:* this goes by induction on A . Trivial for $A = \text{NAT}$. Let $M \in \text{Red}_{B \rightarrow C}$: remark that there is a value $V \in \text{Red}_B$, thus $(M V) \in \text{Red}_C$ and $(M V)$ is AST by IH; using Lemma 1 we get M AST and it is easy to see that if $U \in |\llbracket M \rrbracket|$ then $U \in |\mathcal{M}|$ for some $M \rightarrow^* \mathcal{M}$ so that $U \in \text{Red}_{B \rightarrow C}$ by saturation. Conversely, let M be AST with $|\llbracket M \rrbracket| \subseteq \text{Red}_{B \rightarrow C}$ and let $V \in \text{Red}_B$ be a value: by IH, for any $U \in |\llbracket M \rrbracket| \subseteq \text{Red}_{B \rightarrow C}$ we have $(U V)$ AST with an evaluation supported by elements of Red_C ; by Lemma 1 $\llbracket M V \rrbracket = \llbracket \llbracket M \rrbracket V \rrbracket$ meaning that $(M V)$ is AST and has an evaluation supported by elements of Red_C , so that we can conclude by IH. Similar for products.
- *Every term M such that $x_1 : A_1, \dots, x_n : A_n \vdash M : B$ is a candidate in the sense that if $V_i \in \text{Red}_{A_i}$ for every $1 \leq i \leq n$, then $M[V_1/x_1, \dots, V_n/x_n] \in \text{Red}_B$:* by induction on the type derivation. The only difficult cases are those for the application and for \mathbf{X} (the one for **rec** is just an induction on its third argument).
 - We need to show that if $M \in \text{Red}_{A \rightarrow B}$ and $N \in \text{Red}_A$ then $(M N) \in \text{Red}_B$. But since $N \in \text{Red}_A$, this means that it is AST and for every $V \in |\llbracket N \rrbracket|$, $(M V) \in \text{Red}_B$. In particular, by Lemma 1, we have $\llbracket M N \rrbracket = \llbracket M \llbracket N \rrbracket \rrbracket$ so that $(M N)$ is AST and $|\llbracket M N \rrbracket| \subseteq \bigcup_{V \in |\llbracket N \rrbracket|} |\llbracket M V \rrbracket| \subseteq \text{Red}_B$.
 - We need to show that for any value $U \in \text{Red}_{A \rightarrow A}$ and $V \in \text{Red}_A$ if holds that $(\mathbf{X} \langle U, V \rangle) \in \text{Red}_A$. By an easy induction on n , $(U^n V) \in \text{Red}_A$. Moreover, by an easy induction on n we have $\llbracket \mathbf{X} \langle U, V \rangle \rrbracket = \frac{1}{2^{n+1}} \llbracket U^n (\mathbf{X} \langle U, V \rangle) \rrbracket + \sum_{i \leq n} \frac{1}{2^{i+1}} \llbracket U^i V \rrbracket$ so that at the limit $\llbracket \mathbf{X} \langle U, V \rangle \rrbracket = \sum_{i \in \mathbb{N}} \frac{1}{2^{i+1}} \llbracket U^i V \rrbracket$. We can then conclude that $(\mathbf{X} \langle U, V \rangle)$ is AST (since each of the $(U^i V) \in \text{Red}_B$ are AST and $\sum_i \frac{1}{2^{i+1}} = 1$) and that $|\llbracket M N \rrbracket| = \bigcup_i |\llbracket U^i V \rrbracket| \subseteq \text{Red}_A$.

⁴ Another proof of almost sure termination using reducibility candidate can be found in [25].

This concludes the proof. \square

Almost-sure termination could however be seen as too weak a property: there is no guarantee about the average computation length. For this reason, another stronger notion is often considered, namely *positive* almost-sure termination:

Definition 3. *A term M is said to be positively almost-surely terminating (or PAST) iff the average reduction length $[M]$ is finite.*

Gödel's \mathbb{T} , when paired with \mathbf{R} , is combinatorially too powerful to guarantee positive almost sure termination:

Theorem 2. $\mathbb{T}^{\oplus, \mathbf{R}, \mathbf{X}}$ *is not positively almost-surely terminating.*

Proof. The naive exponential function applied to \mathbf{R} is computing, with probability $\frac{1}{2^{n+1}}$ the number 2^{n+1} in time 2^{n+1} . This is already a counterexample, because it clearly has infinite average termination time. \square

3.3 On Fragments of $\mathbb{T}^{\oplus, \mathbf{R}, \mathbf{X}}$: a Roadmap

The calculus $\mathbb{T}^{\oplus, \mathbf{R}, \mathbf{X}}$ contains at least four fragments, namely Gödel's \mathbb{T} and the three fragments \mathbb{T}^{\oplus} , $\mathbb{T}^{\mathbf{R}}$ and $\mathbb{T}^{\mathbf{X}}$ corresponding to the three probabilistic choice operators we consider. It is then natural to ask how these fragments relate to each other as for their respective expressive power. At the end of this paper, we will have a very clear picture in front of us.

The first result we can give is the equivalence between the apparently dual fragments $\mathbb{T}^{\mathbf{R}}$ and $\mathbb{T}^{\mathbf{X}}$. The embeddings are in fact quite simple:

Proposition 1. $\mathbb{T}^{\mathbf{R}}$ *and* $\mathbb{T}^{\mathbf{X}}$ *are both equiexpressive with* $\mathbb{T}^{\oplus, \mathbf{R}, \mathbf{X}}$.

Proof. The calculus $\mathbb{T}^{\mathbf{R}}$ embeds the full system $\mathbb{T}^{\oplus, \mathbf{R}, \mathbf{X}}$ via the encoding:⁵

$$M \oplus N := \mathbf{rec}\langle \lambda z.N, \lambda xyz.M, \mathbf{R} \rangle \mathbf{0}; \quad \mathbf{X} := \lambda xy.\mathbf{rec}\langle y, \lambda z.x, \mathbf{R} \rangle.$$

The fragment $\mathbb{T}^{\mathbf{X}}$ embeds the full system $\mathbb{T}^{\oplus, \mathbf{R}, \mathbf{X}}$ via the encoding:

$$M \oplus N := \mathbf{X}\langle \lambda xy.M, \lambda y.N \rangle \mathbf{0}; \quad \mathbf{R} := \mathbf{X}\langle \mathbf{S}, \mathbf{0} \rangle.$$

In both cases, the embedding is compositional and preserves types. That the two embeddings are correct can be proved easily, see [3]. \square

Notice how simulating \mathbf{X} by \mathbf{R} requires the presence of recursion, while the converse is not true. The implications of this fact are intriguing, but lie outside the scope of this work.

In the following, we will no longer consider $\mathbb{T}^{\mathbf{X}}$ nor $\mathbb{T}^{\oplus, \mathbf{R}, \mathbf{X}}$ but only $\mathbb{T}^{\mathbf{R}}$, keeping in mind that all these are equiexpressive due to Proposition 1. The rest of this paper, thus, will be concerned with understanding the relative expressive power

⁵ Notice that the dummy abstractions on z and the $\mathbf{0}$ at the end ensure the correct reduction order by making $\lambda z.N$ a value.

of the three fragments \mathbb{T} , \mathbb{T}^\oplus , and \mathbb{T}^R . Can any of the (obvious) strict *syntactical* inclusions between them be turned into a strict *semantic* inclusion? Are the three systems equiexpressive?

In order to compare probabilistic calculi to deterministic ones, several options are available. The most common one is to consider notions of observations over the probabilistic outputs; this will be the purpose of Section 6. In this section, we will look at whether it is possible to deterministically represent the distributions computed by the probabilistic calculus at hand. We say that the distribution $\mathcal{M} \in \mathfrak{D}(\mathbb{N})$ is *finitely represented* by⁶ $f : \mathbb{N} \rightarrow \mathbb{B}$, if there exists a q such that for every $k \geq q$ it holds that $f(k) = 0$ and

$$\mathcal{M} = \{\mathbf{k} \mapsto f(k)\}.$$

Moreover, the definition can be extended to families of distributions $(\mathcal{M}_n)_n$ by requiring the existence of $f : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{B}$, $q : \mathbb{N} \rightarrow \mathbb{N}$ such that for all $k \geq q(n)$, $f(n, k) = 0$ and

$$\forall n, \quad \mathcal{M}_n = \{\mathbf{k} \mapsto f(n, k)\}.$$

In this case, we say that the representation is *parameterized*.

We will see in Section 4 that the distributions computed by \mathbb{T}^\oplus are exactly the (parametrically) finitely representable by \mathbb{T} terms. In \mathbb{T}^R , however, distributions are more complex (infinite, non-rational). That is why only a characterisation in terms of approximations is possible. More specifically, a distribution $\mathcal{M} \in \mathfrak{D}(\mathbb{N})$ is said to be *functionally represented* by two functions $f : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{B}$ and $g : \mathbb{N} \rightarrow \mathbb{N}$ iff for every $n \in \mathbb{N}$ and for every $k \geq g(n)$ it holds that $f(n, k) = 0$ and

$$\sum_{k \in \mathbb{N}} \left| \mathcal{M}(k) - f(n, k) \right| \leq \frac{1}{n}.$$

In other words, the distribution \mathcal{M} can be approximated arbitrarily well, and uniformly, by finitely representable ones. Similarly, we can define a parameterised version of this definition at first order.

In Section 5, we show that distributions generated by \mathbb{T}^R terms are indeed uniform limits over those of \mathbb{T}^\oplus ; using our result on \mathbb{T}^\oplus this give their (parametric) functional representability in the deterministic \mathbb{T} .

4 Binary Probabilistic Choice

This section is concerned with two theoretical results on the expressive power of \mathbb{T}^\oplus . The main feature of \mathbb{T}^\oplus is that its terms are *positively* almost surely terminating. This is a corollary of the following theorem (whose proof [3] proceeds again by reducibility).

⁶ Here \mathbb{B} stands for the set of dyadic numbers, i.e. rationals in the form $\frac{n}{2^m}$ (where $m, n \in \mathbb{N}$) and \mathbf{BIN} for their representation in system \mathbb{T} , encoded as pairs of natural numbers.

Theorem 3. *For any term $M \in \mathbb{T}^\oplus$, $M \rightarrow^* \llbracket M \rrbracket$.*

Now, if $M \rightarrow^n \llbracket M \rrbracket$, then $\llbracket M \rrbracket$ can be at most n since the distribution $\mathcal{M}_m^{\Delta V}$ of values reachable in exactly m steps (see Definition 1) will be 0 for every $m > n$. But this means that typable terms normalise in finite time:

Corollary 1. *Any term $M \in \mathbb{T}^\oplus$ is positively almost-surely terminating.*

But this is not the only consequence. In fact, the finiteness of $\llbracket M \rrbracket$ and the fact that \mathbb{T}^\oplus is sufficiently expressive allow for a finite representation of \mathbb{T}^\oplus -distributions by \mathbb{T} -definable functions. To prove it, let us consider an extension of \mathbb{T} with a single memory-cell c of type NAT . This memory-cell is used to store some “random coins” simulating probabilistic choices. The operator \oplus can be encoded as follows:

$$(M \oplus N)^* \quad := \quad \text{if } (\text{mod}_2 c) \text{ then } (c := \text{div}_2 c ; M^*) \text{ else } (c := \text{div}_2 c ; N^*)$$

Notice that conditionals and modulo arithmetic are easily implementable in \mathbb{T} . From Theorem 3, we know that for any $M \in \mathbb{T}^\oplus(\text{NAT})$, there is $n \in \mathbb{N}$ such that $M \rightarrow^n \llbracket M \rrbracket$, and since the evaluation of M can thus involve at most n successive probabilistic choices, we have that

$$\llbracket M \rrbracket(\mathbf{k}) = \frac{\#\{m < 2^n \mid k = \llbracket c := \mathbf{m} ; M^* \rrbracket\}}{2^n}.$$

By way of a state-passing transformation, we can enforce $(c := \mathbf{m} ; M^*)$ into a term of \mathbb{T} . But then, the whole $\#\{m < 2^n \mid k = \llbracket c := \mathbf{m} ; M^* \rrbracket\}$ can be represented as a \mathbb{T} -term $k : \mathbb{N} \vdash N : \mathbb{N}$ which finitely represents the distribution $\llbracket M \rrbracket$.

In the long version of this paper [3], a stronger result is proved, namely that for any functional $M \in \mathbb{T}^\oplus(\text{NAT} \rightarrow \text{NAT})$, there are terms $M_\downarrow \in \mathbb{T}(\text{NAT} \rightarrow \text{NAT} \rightarrow \text{NAT})$ and $M_\# \in \mathbb{T}(\text{NAT} \rightarrow \text{NAT})$ such that for all $n \in \mathbb{N}$:

$$\llbracket M \mathbf{n} \rrbracket(\mathbf{k}) = \frac{\#\{m < 2^{\llbracket M_\# \mathbf{n} \rrbracket} \mid k = \llbracket M_\downarrow \mathbf{n} \mathbf{m} \rrbracket\}}{2^{\llbracket M_\# \mathbf{n} \rrbracket}}.$$

The supplementary difficulty, here, comes from the bound $M_\#$ that have to be computed dynamically as it depends on its argument \mathbf{n} .

As a consequence:

Theorem 4. *Distributions generated by \mathbb{T}^\oplus -terms are precisely those which can be finitely generated by parameterized \mathbb{T} -functionals; i.e., for any term $M : \text{NAT} \rightarrow \text{NAT}$, there are two \mathbb{T} -functionals $f : (\mathbb{N} \times \mathbb{N}) \rightarrow \mathbb{B}$ and $q : \mathbb{N} \rightarrow \mathbb{N}$ such that for all n :*

$$\llbracket M \mathbf{n} \rrbracket = \{\mathbf{k} \mapsto f(n, k) \mid k \leq q(n)\}.$$


5 Countable Probabilistic Choice

In this section, we show that \mathbb{T}^\oplus approximates \mathbb{T}^{R} : for any term $M \in \mathbb{T}^{\text{R}}(\text{NAT})$, there is a term $N \in \mathbb{T}^\oplus(\text{NAT} \rightarrow \text{NAT})$ that represents a sequence approximating M

uniformly. We will here make strong use of the fact that M has type \mathbf{NAT} . This is a natural drawback when we understand that the encoding $(\cdot)^\dagger$ on which the result above is based is not direct, but goes through yet another state passing style transformation. Nonetheless, everything can be lifted easily to the first order, achieving the parameterisation of our theorem.

The basic idea behind the embedding $(\cdot)^\dagger$ is to mimic any instance of the operator \mathbf{R} in the source term by some term $\mathbf{0} \oplus (1 \oplus (\cdots (\mathbf{n} \oplus \perp) \cdots))$, where n is *sufficiently large*, and \perp is an arbitrary value of type \mathbf{NAT} . Of course, the semantics of this term is *not* the same as that of \mathbf{R} , due to the presence of \perp ; however, n will be chosen sufficiently large for the difference to be negligible. Notice, moreover, that this term can be generalized into the following parametric form $\mathbf{R}^\ddagger := \lambda x. \text{rec } \langle \perp, (\lambda x. \mathbf{S} \oplus (\lambda y. \mathbf{0})) \rangle, x$.

Once \mathbf{R}^\ddagger is available, a natural candidate for the encoding $(\cdot)^\dagger$ would be to consider something like $M^\ddagger := \lambda z. M[(\mathbf{R}^\ddagger z)/\mathbf{R}]$. In the underlying execution tree, $(M^\ddagger \mathbf{n})$ correctly simulates the first n branches of \mathbf{R} (which has infinite arity), but truncates the rest with garbage terms \perp :



The question is whether the remaining untruncated tree has a “sufficient weight”, i.e., whether there is a minimal bound to the probability to stay in this untruncated tree. However, in general $(\cdot)^\ddagger$ fails on this point, not achieving to approximate M uniformly. In fact, this probability is basically $(1 - \frac{1}{2^n})^d$ where d is its depth. Since in general the depth of the untruncated tree can grow very rapidly on n in a powerful system like \mathbb{T} , there is no hope for this transformation to perform a uniform approximation.

The solution we are using is to have the precision m of $\mathbf{0} \oplus (1 \oplus (\cdots (\mathbf{m} \oplus \perp) \cdots))$ to *dynamically grow* along the computation. More specifically, in the approximants $M^\ddagger \mathbf{n}$, the growing speed of m will increase with n : in the n -th approximant $M^\ddagger \mathbf{n}$, the operator \mathbf{R} will be simulated as $\mathbf{0} \oplus (1 \oplus (\cdots (\mathbf{m} \oplus \perp) \cdots))$ and, somehow, m will be updated to $m + n$. Why does it work? Simply because even for an (hypothetical) infinite and complete execution tree of M , we would stay inside the n^{th} untruncated tree with probability $\prod_{k \geq 0} (1 - \frac{1}{2^{k+n}})$ which is asymptotically above $(1 - \frac{1}{n})$.

Implementing this scheme in \mathbb{T}^\oplus requires a feature which is not available (but which can be encoded), namely ground-type references. We then prefer to show that the just described scheme can be realised in an intermediate language called $\mathbb{T}^{\mathbf{R}}$, whose operational semantics is formulated not on *terms*, but rather on triples in the form (M, m, n) , where M is the term currently being evaluated, m is the current approximation threshold value, and n is the value of which m is incremented whenever \mathbf{R} is simulated. The operational semantics is standard, except for the following rule:

$$\overline{(\bar{\mathbf{R}}, m, n) \rightarrow \left\{ (k, m+n, n) \mapsto \frac{1}{2^{k+1}} \mid k < m \right\}}^{(r-\bar{\mathbf{R}})}$$

Notice how this operator behaves similarly to \mathbf{R} with the exception that it fails when drawing too big of a number (*i.e.*, bigger than the first state m). Notice that the failure is represented by the fact that the resulting distribution does not necessarily sum to 1. The intermediate language $\mathbb{T}^{\bar{\mathbf{R}}}$ is able to approximate $\mathbb{T}^{\mathbf{R}}$ at every order (Theorem 5 below). Moreover, the two memory cells can be shown to be expressible in \mathbb{T}^{\oplus} , again by way of a continuation-passing transformation. Crucially, the initial value of n can be passed as an argument to the encoded term.

For any $M \in \mathbb{T}^{\mathbf{R}}$ we denote $M^* := M[\bar{\mathbf{R}}/\mathbf{R}]$. We say that $(M, m, n) \in \mathbb{T}^{\bar{\mathbf{R}}}$ if $m, n \in \mathbb{N}$ and $M = N^*$ for some $N \in \mathbb{T}^{\mathbf{R}}$. Similarly, $\mathfrak{D}(\mathbb{T}^{\bar{\mathbf{R}}})$ is the set of probabilistic distributions over $\mathbb{T}^{\bar{\mathbf{R}}} \times \mathbb{N}^2$, *i.e.*, over the terms plus states.

For any m and n , the behaviour of M and (M^*, m, n) are similar, except that (M^*, m, n) will “fail” more often. In other words, all $(M^*, m, n)_{m, n \in \mathbb{N}}$ somehow approximate M from below:

Lemma 2. *For any $M \in \mathbb{T}^{\mathbf{R}}$ and any $m, n \in \mathbb{N}$, $\llbracket M \rrbracket \succeq \llbracket M^*, m, n \rrbracket$, *i.e.*, for every $V \in \mathbb{T}_V^{\mathbf{R}}$, we have*

$$\llbracket M \rrbracket (V) \geq \sum_{p, q} \llbracket M^*, m, n \rrbracket (V^*, p, q).$$

Proof. By an easy induction, one can show that for any $\mathcal{M} \in \mathfrak{D}(\mathbb{T}^{\mathbf{R}})$ and $\mathcal{N} \in \mathfrak{D}(\mathbb{T}^{\bar{\mathbf{R}}})$ if $\mathcal{M} \succeq \mathcal{N}$, $\mathcal{M} \rightarrow \mathcal{L}$ and $\mathcal{N} \rightarrow \mathcal{P}$, then $\mathcal{L} \succeq \mathcal{P}$. This ordering is then preserved at the limit so that we get our result. \square

In fact, the probability of “failure” of any $(M, m, n)_{m, n \in \mathbb{N}}$ can be upper-bounded explicitly. More precisely, we can find an infinite product underapproximating the success rate of (M, m, n) by reasoning inductively over the execution $(M, m, n) \rightarrow^* \llbracket (M, m, n) \rrbracket$, which is possible because of the PAST.

Lemma 3. *For any $M \in \mathbb{T}^{\bar{\mathbf{R}}}$ and any $m, n \geq 1$*

$$\text{Succ}(M, m, n) \geq \prod_{k \geq 0} \left(1 - \frac{1}{2^{m+kn}} \right).$$

Proof. We denote $\#(m, n) := \prod_{k \geq 0} \left(1 - \frac{1}{2^{m+kn}} \right)$ and $\#\mathcal{M} := \int_{\mathcal{M}} \#(m, n) dM dmdn$.

Remark that for any M and any m, n , if $(M, m, n) \rightarrow \mathcal{M}$ then \mathcal{M} is either of the form $\{(N, m, n)\}$ or $\{(N_i, m+n, n) \mapsto \frac{1}{2^{i+1}} \mid i < m\}$ for some N of $(N_i)_{i \leq m}$. Thus we have that if $(M, m, n) \rightarrow \mathcal{N}$ then $\#\mathcal{N} = \#(m, n)$ and that if $\mathcal{M} \rightarrow \mathcal{N}$ then $\#\mathcal{N} = \#\mathcal{M}$. In particular, since $(M, m, n) \rightarrow^* \llbracket M, m, n \rrbracket$ we can conclude

$$\text{Succ}(M, m, n) = \int_{\llbracket M, m, n \rrbracket} 1 dM dmdn \geq \#\llbracket M \rrbracket = \#\mathcal{M} = \#(m, n) = \prod_{k \geq 0} \left(1 - \frac{1}{2^{m+kn}} \right).$$

\square

This gives us an analytic lower bound to the success rate of (M, m, n) . However, it is not obvious that this infinite product is an interesting bound, it is not even clear that it can be different from 0. This is why we will further underapproximate this infinite product to get a simpler expression whenever $m = n$:

Lemma 4. *For any $M \in \mathbb{T}^{\bar{\mathbf{R}}}$ and any $n \geq 4$*

$$\text{Succ}(M, n, n) \geq 1 - \frac{1}{n}.$$

Proof. By Lemma 3 we have that $\text{Succ}(M, n, n) \geq \prod_{k \geq 1} (1 - \frac{1}{2^{kn}})$ which is above the product $\prod_{k \geq 1} (1 - \frac{1}{n^2 k^2})$ whenever $n \geq 4$. This infinite product has been shown by Euler to be equal to $\frac{\sin(\frac{\pi}{n})}{\frac{\pi}{n}}$. By an easy numerical analysis we then obtain that $\frac{\sin(\frac{\pi}{n})}{\frac{\pi}{n}} \geq 1 - \frac{1}{n}$. \square

This lemma can be restated by saying that the probability of “failure” of (M^*, n, n) , *i.e.* the difference between $\llbracket M^*, n, n \rrbracket$ and $\llbracket M \rrbracket$, is bounded by $\frac{1}{n}$. With this we then get our first theorem, which is the uniform approximability of elements of $\mathbb{T}^{\mathbf{R}}$ by those of $\mathbb{T}^{\bar{\mathbf{R}}}$:

Theorem 5. *For any $M \in \mathbb{T}^{\mathbf{R}}$ and any $n \in \mathbb{N}$,*

$$\sum_V \left| \llbracket M \rrbracket(V) - \Sigma_{m', n'} \llbracket M^*, n, n \rrbracket(V^*, m', n') \right| \leq \frac{1}{n}.$$

Proof. By Lemma 2, for each V the difference is positive, thus we can remove the absolute value and distribute the sum. We conclude by using the fact that $\text{Succ}(M) = 1$ and $\text{Succ}(M^*, n, n) \geq 1 - \frac{1}{n}$. \square

The second theorem, *i.e.*, the uniform approximability of ground elements of $\mathbb{T}^{\mathbf{R}}$ by those of \mathbb{T}^{\oplus} , follows immediately:

Theorem 6. *Distributions in $\mathbb{T}^{\mathbf{R}}(\mathbf{NAT})$ can be approximated by \mathbb{T}^{\oplus} -distributions (which are finitely \mathbb{T} -representable), *i.e.*, for any $M \in \mathbb{T}^{\mathbf{R}}(\mathbf{NAT})$, there is $M^\dagger \in \mathbb{T}^{\oplus}(\mathbf{NAT} \rightarrow \mathbf{NAT})$ such that:*

$$\forall n, \sum_k \left| \llbracket M \rrbracket(\mathbf{k}) - \llbracket M^\dagger \mathbf{n} \rrbracket(\mathbf{k}) \right| \leq \frac{1}{n}.$$

Moreover:

- the encoding is parameterisable, in the sense that for all $M \in \mathbb{T}^{\mathbf{R}}(\mathbf{NAT} \rightarrow \mathbf{NAT})$, there is $M^\dagger \in \mathbb{T}^{\oplus}(\mathbf{NAT} \rightarrow \mathbf{NAT} \rightarrow \mathbf{NAT})$ such that $\llbracket (M \mathbf{m})^\dagger \rrbracket = \llbracket M^\dagger \mathbf{m} \rrbracket$ for all $m \in \mathbb{N}$;
- the encoding is such that $\llbracket M \rrbracket(\mathbf{k}) \leq \llbracket M^\dagger \mathbf{n} \rrbracket(\mathbf{k})$ only if $k = 0$.

Proof. It is clear that in an extension of \mathbb{T}^{\oplus} with two global memory cells m, n and with an exception monad, the $\bar{\mathbf{R}}$ operator can be encoded by the term $\bar{\mathbf{R}} := \mathbf{rec} \langle \perp, (\lambda x. \mathbf{S} \oplus (\lambda y. \mathbf{0})), (m := !m + !n) \rangle$, where \perp is raising an error/exception and where $m := !m + !n$ is returning the value of m before changing the memory cell to $m + n$. We can conclude by referring to the usual state passing style encoding of exceptions and state-monads into \mathbb{T} (and thus into \mathbb{T}^{\oplus}). \square

6 Subrecursion

If one wishes to define \mathbb{T}^\oplus -definable or \mathbb{T}^R -definable functions as a set of ordinary set-theoretic functions (say from \mathbb{N} to \mathbb{N}), it is necessary to collapse the random output into a deterministic one. As already acknowledged by the complexity community, there are at least two reasonable ways to do so: by using either Monte Carlo or Las Vegas observations.

As the careful reader may have foreseen, the finite (parametric) representation of \mathbb{T}^\oplus -distributions into \mathbb{T} is collapsing both observations into \mathbb{T} -definable functions. One only need to explore the finite representation, the resulting process suffers from an exponential blow-up, which is easily absorbed by \mathbb{T} , in which all elementary functions (and much more than that!) can be expressed.

Theorem 7 (Monte Carlo). *Let $f : \mathbb{N} \rightarrow \mathbb{N}$ and $M : \text{NAT} \rightarrow \text{NAT}$ a \mathbb{T}^R -term such that $(M \mathbf{m})$ evaluates into $f(m)$ with probability $p \geq \frac{1}{2} + \frac{1}{g(m)}$ for a \mathbb{T} -definable function g . Then f is \mathbb{T} -definable.*

Theorem 8 (Las Vegas). *Let $f : \mathbb{N} \rightarrow \mathbb{N}$ and $M : \text{NAT} \rightarrow \text{NAT}$ a \mathbb{T}^R -term such that $(M \mathbf{m})$ evaluate either to $\mathbf{0}$ (representing a failure) or to $\mathbf{f(m)} + \mathbf{1}$, the later happening with probability $p \geq \frac{1}{g(m)}$ for some \mathbb{T} -definable function g . Then f is \mathbb{T} -definable.*

7 Conclusions

This paper is concerned with the impact of adding various forms of probabilistic choice operators to a higher-order subrecursive calculus in the style of Gödel's \mathbb{T} . The presented results help in understanding the relative expressive power of various calculi which can be obtained this way, by showing some separation and equivalence results.

The probabilistic choice operators we consider here are just examples of how one can turn a deterministic calculus like \mathbb{T} into a probabilistic model of computation. The expressiveness of $\mathbb{T}^{\oplus, R, X}$ is sufficient to encode most reasonable probabilistic operators, but what can we say about their own expressive power? For example, what about a ternary operator in which either of the first two operators is chosen with a probability *which depends* on the value of the third operator? A general theory of probabilistic choice operators and of their expressive power is still lacking.

Another research direction to which this paper hints at consists in studying the logical and proof-theoretical implications of endowing a calculus like \mathbb{T} with probabilistic choice operators. What is even more exciting, however, is the application of the ideas presented here to polynomial time computation. This would allow to go towards a characterization of expected polynomial time computation, thus greatly improving on the existing works on the implicit complexity of probabilistic systems [5, 7], which only deals with worst-case execution time. The authors are currently engaged in that.

References

1. Henk P. Barendregt. *The Lambda Calculus, Its Syntax and Semantics*. Studies in logic and the foundations of mathematics. North-Holland, 1981.
2. Olivier Bournez and Florent Garnier. Proving positive almost-sure termination. In *Proc. of RTA*, volume 3467 of *LNCS*, pages 323–337, 2005.
3. Flavien Breuvert, Ugo Dal Lago, and Agathe Herrou. On probabilistic subrecursion (long version). Available at <http://arxiv.org/abs/1701.04786>, 2016.
4. Raphaëlle Crubillé and Ugo Dal Lago. On probabilistic applicative bisimulation and call-by-value λ -calculi. In *Proc. of ESOP*, volume 8410 of *LNCS*, pages 209–228, 2014.
5. Ugo Dal Lago and Paolo Parisen Toldin. A higher-order characterization of probabilistic polynomial time. *Inf. Comput.*, 241:114–141, 2015.
6. Ugo Dal Lago and Margherita Zorzi. Probabilistic operational semantics for the lambda calculus. *RAIRO - Theor. Inf. and Applic.*, 46(3):413–450, 2012.
7. Ugo Dal Lago, Sara Zuppiroli, and Maurizio Gabbriellini. Probabilistic recursion theory and implicit computational complexity. *Sci. Ann. Comp. Sci.*, 24(2):177–216, 2014.
8. Karel De Leeuw, Edward F. Moore, Claude E. Shannon, and Norman Shapiro. Computability by probabilistic machines. *Automata studies*, 34:183–198, 1956.
9. Thomas Ehrhard, Michele Pagani, and Christine Tasson. Probabilistic Coherence Spaces are Fully Abstract for Probabilistic PCF. In P. Sewell, editor, *Proc. of POPL*. ACM, 2014.
10. Luis María Ferrer Fioriti and Holger Hermanns. Probabilistic termination: Soundness, completeness, and compositionality. In *Proc. of POPL*, pages 489–501, 2015.
11. John Gill. Computational complexity of probabilistic turing machines. *SIAM J. Comput.*, 6(4):675–695, 1977.
12. Jean-Yves Girard, Paul Taylor, and Yves Lafont. *Proofs and Types*. Cambridge University Press, 1989.
13. Noah D. Goodman, Vikash K. Mansinghka, Daniel M. Roy, Keith Bonawitz, and Joshua B. Tenenbaum. Church: a language for generative models. In *UAI*, pages 220–229, 2008.
14. Achim Jung and Regina Tix. The troublesome probabilistic powerdomain. *Electr. Notes Theor. Comput. Sci.*, 13:70–91, 1998.
15. Benjamin Lucien Kaminski and Joost-Pieter Katoen. On the hardness of almost-sure termination. In *Proc. of MFCS*, volume 9234 of *LNCS*, pages 307–318, 2015.
16. Dexter Kozen. Semantics of probabilistic programs. *J. Comput. Syst. Sci.*, 22(3):328–350, 1981.
17. Christopher D Manning and Hinrich Schütze. *Foundations of statistical natural language processing*, volume 999. MIT Press, 1999.
18. Annabelle McIver and Carroll Morgan. *Abstraction, Refinement and Proof for Probabilistic Systems*. Monographs in Computer Science. Springer, 2005.
19. Judea Pearl. *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan Kaufmann, 1988.
20. Simon J. D. Prince. *Computer Vision: Models, Learning, and Inference*. Cambridge University Press, New York, NY, USA, 2012.
21. N. Saheb-Djahromi. Probabilistic LCF. In *Proc. of MFCS*, volume 64 of *LNCS*, pages 442–451, 1978.
22. Eugene S Santos. Probabilistic Turing machines and computability. *Proceedings of the American Mathematical Society*, 22(3):704–710, 1969.

23. Morten Heine Sørensen and Pawel Urzyczyn. *Lectures on the Curry-Howard Isomorphism*. Elsevier Science Inc., New York, NY, USA, 2006.
24. Richard Statman. The typed lambda-calculus is not elementary recursive. *Theor. Comput. Sci.*, 9:73–81, 1979.
25. Sam Staton, Hongseok Yang, Chris Heunen, Ohad Kammar, and Frank Wood. Semantics for probabilistic programming: higher-order functions, continuous distributions, and soft constraints. In *Proc. of LICS*, pages 525–534, 2016.