Fully Abstract Models of the Probabilistic λ-calculus

Pierre Clairambault
Univ Lyon, CNRS, ENS de Lyon, UCB Lyon 1, LIP, France
pierre.clairambault@ens-lyon.fr

Hugo Paquet
Department of Computer Science and Technology, University of Cambridge, UK
hugo.paquet@cl.cam.ac.uk

Abstract

We compare three models of the probabilistic λ-calculus: the *probabilistic Böhm trees* of Leventis, the *probabilistic concurrent games* of Winskel et al., and the *weighted relational model* of Ehrhard et al. Probabilistic Böhm trees and probabilistic strategies are shown to be related by a precise *correspondence theorem*, in the spirit of existing work for the pure λ-calculus. Using Leventis' theorem (probabilistic Böhm trees characterise observational equivalence), we derive a full abstraction result for the games model. Then, we relate probabilistic strategies to the weighted relational model, using an *interpretation-preserving functor* from the former to the latter. We obtain that the relational model is also fully abstract.

2012 ACM Subject Classification Theory of computation → Denotational semantics; Probabilistic Computation.

Keywords and phrases Game Semantics, Lambda-calculus, Probabilistic programming, Relational model, Full abstraction.

Digital Object Identifier 10.4230/LIPIcs.CSL.2018.15

Acknowledgements This work was performed within the framework of the LABEX MILYON (ANR-10-LABX-0070) of Université de Lyon, within the program "Investissements d'Avenir" (ANR-11-IDEX-0007) operated by the French National Research Agency (ANR).

1 Introduction

The interest in probabilistic programs in recent years, driven in particular by applications in machine learning and statistical modelling, has triggered the need for theoretical foundations, going beyond the pioneering work of Kozen [14] and Saheb-Djahromi [21]. Although a variety of approaches exist, we focus on the *probabilistic λ-calculus* Λ^+, which extends the pure (untyped) λ-calculus with a probabilistic choice operator. The extension is natural and applications are quick to arise — see for instance [3]. But in order for Λ^+ to become a useful formal model for probabilistic computation, the extensive classical theory of the λ-calculus must be readapted.

Among the existing research in this direction, we are especially interested in the work of Ehrhard, Pagani and Tasson [11], and of Leventis [16, 17]. In [11], the authors define an operational semantics for Λ^+ and study a model in the category of *probabilistic coherence spaces*, an existing model [9] of Probabilistic PCF. They prove an adequacy theorem for Λ^+, and this result also applies to the *weighted relational model*, of which probabilistic coherence spaces are a refinement.

27th EACSL Annual Conference on Computer Science Logic (CSL 2018).
Editors: Dan Ghica and Achim Jung; Article No. 15; pp. 15:1–15:17
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

More recently, the PhD thesis of Leventis [16] offers a thorough exploration of the syntactical aspects of the calculus. In particular the author defines a notion of *probabilistic Böhm tree*, and redevelops the Böhm theory of the λ-calculus in a probabilistic setting. This includes Böhm 's separation theorem: probabilistic Böhm trees, in their *infinitely extensional* form, precisely characterise observational equivalence in Λ^+.

In this paper, we propose an alternative model in the framework of concurrent games, integrating ideas from our earlier work on a concurrent games model of probabilistic PCF [5] and from Ker, Ong and Nickau's fully abstract semantics of the pure untyped λ-calculus [13].

In [13], an exact correspondence is proved between strategies and infinitely extensional Böhm trees. Drawing inspiration from that work, we relate probabilistic strategies and probabilistic Böhm trees, but unlike [13], the correspondence is not bijective, because of the additional branching information contained in probabilistic strategies. By quotienting out this information, we derive from Leventis' theorem a full abstraction result for the games model.

Finally, we study a functor from the probabilistic games model to the weighted relational model. This functor is a *time-forgetting* operation on strategies, in the spirit of [1]. Note that proving the functoriality of such operations is usually challenging even without probabilities, see for example Melliès' work [19] — here, we address this by leveraging a "deadlock-free lemma" proved for concurrent strategies in [5]. We show that this functor preserves the interpretation of Λ^+, with significant consequences: Ehrhard et al.'s adequacy result can be lifted to strategies, and the full abstraction result obtained for games via probabilistic Böhm trees can be shown to hold also for the weighted relational model, so far only known to be adequate[1].

In Section 2, we present Λ^+ and its operational semantics; we also recall Leventis' work on probabilistic Böhm trees and define concurrent probabilistic strategies, hinting at the correspondence between the two. In Section 3, we outline the construction of a category of concurrent games and probabilistic strategies, and the reflexive object that it contains. We then study, in Section 4, the correspondence between probabilistic strategies and probabilistic Böhm trees, and prove full abstraction for the games model. Finally, in Section 5, we collapse probabilistic strategies down to weighted relations, thus showing full abstraction for the relational model.

2 The Probabilistic λ-calculus

2.1 Syntax

The set Λ^+ of terms of the probabilistic λ-calculus is defined by the following grammar, where p ranges over the interval $[0, 1]$ and x over an infinite set Var:

$$M, N ::= x \mid \lambda x.M \mid MN \mid M +_p N.$$

Write Λ_0^+ for the set of **closed terms**, *i.e.* those with no free variables.

The operator $+_p$ represents probabilistic choice, so that a term of the form $M +_p N$ has two possible reduction steps: to M, with probability p, and to N, with probability $1 - p$. Accordingly, the reduction relation we consider is a Markov process over the set Λ^+, and corresponds to a probabilistic variant of the standard **head-reduction**. It is defined inductively:

[1] Independently and using a different method, Leventis and Pagani have obtained an alternative proof of full abstraction, but this work is so far unpublished.

$$\frac{}{(\lambda x.M)N \xrightarrow{1} M[N/x]} \qquad \frac{}{M +_p N \xrightarrow{p} M} \qquad \frac{}{M +_p N \xrightarrow{1-p} N}$$

$$\frac{M \xrightarrow{p} M'}{\lambda x.M \xrightarrow{p} \lambda x.M'} \qquad \frac{M \xrightarrow{p} M' \qquad M \neq \lambda x.P}{MN \xrightarrow{p} M'N}$$

For $M, N \in \Lambda^+$, there may be many reduction paths from M to N. The **weight** of a path $\pi : M \xrightarrow{p_1} \dots \xrightarrow{p_n} N$ is the product of the transition probabilities: $w(\pi) = \prod_{i=1}^{n} p_i$. The **probability of M reducing to N** is then defined as $\Pr(M \to N) = \sum_{\pi : M \to^* N} w(\pi)$.

The normal forms for this reduction are terms of the form $\lambda x_0 \dots x_{n-1}.\ y\ M_0 \dots M_{k-1}$, where $n, k \in \mathbb{N}$ and $M_i \in \Lambda^+$ for all i. Such terms are called **head-normal forms** (hnfs). A pure λ-term has at most one hnf called – if it exists – *its* hnf, though of course, this does not hold in the presence of probabilities.

Given a set $\mathcal{H}$ of hnfs, we set $\Pr(M \to \mathcal{H}) = \sum_{H \in \mathcal{H}} \Pr(M \to H)$. The **probability of convergence** of a term M, denoted $\Pr_{\Downarrow}(M)$, is the probability of M reducing to some hnf: $\Pr_{\Downarrow}(M) = \Pr(M \to \{H \in \Lambda^+ \mid H \text{ hnf}\})$. Finally we say that two terms M and N are **observationally equivalent**, written $M =_{\mathrm{obs}} N$, if for all contexts $C[\ ]$, $\Pr_{\Downarrow}(C[M]) = \Pr_{\Downarrow}(C[N])$.

2.2 Probabilistic Böhm trees

Infinitely extensional Böhm trees for pure λ-terms

There are several notions of infinite normal forms for pure λ-terms, including *e.g.* the **Böhm trees** [2] and the **Lévy-Longo trees**, among others. The normal forms for the probabilistic λ-terms considered in this paper build on the **infinitely extensional Böhm trees** (also called **Nakajima trees**), which provide a notion of infinitely η-expanded normal form.

The infinitely extensional Böhm tree of M is in general an infinite tree, which can be defined as the limit of a sequence of finite-depth approximants. In fact those approximants will suffice for the purposes of this paper: given a λ-term M and $d \in \mathbb{N}$, the tree $\mathrm{BT}^d(M)$ is $\bot$ if $d = 0$ or if M has no head-normal form, and

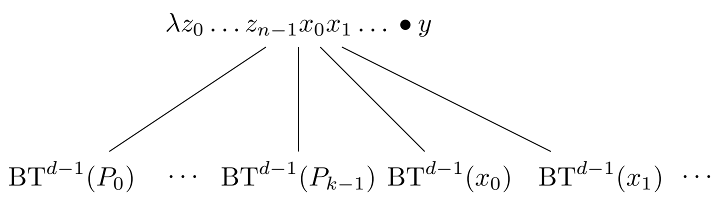

if $d > 0$ and M has hnf $\lambda z_0 \dots z_{n-1}.y\ P_0 \dots P_{k-1}$.

In order to deal with issues of α-renaming, we adopt the same convention as Leventis [16], whereby the infinite sequence of abstracted variables at the root of a tree of depth $d > 0$ is labelled $x_0^d, x_1^d, \dots$ so that any tree is determined by the pair $(y, (T_n)_{n \in \mathbb{N}})$ of its head variable and sequence of subtrees.

Leventis' probabilistic trees

Infinitely extensional Böhm trees for the λ-calculus have striking properties: they characterise observational equivalence of terms, and as a model they yield the *maximal consistent sensible* λ-theory (see [2] for details). In his PhD thesis, Leventis [16] proposes a notion of *probabilistic* Böhm tree which plays the same role for Λ^+. Intuitively, because a term of the form $\lambda x_0 \dots x_{n-1}.z\ P_0 \dots P_{k-1}\ +_p\ \lambda y_0 \dots y_{m-1}.w\ Q_0 \dots Q_{l-1}$ has two hnfs, it may be represented by a probability distribution over trees of the form of that above. Accordingly, two different kinds of trees are considered: **value trees**, representing head-normal forms (without probability distribution at top-level), and **probabilistic Böhm trees**, representing general terms:

▶ **Definition 1.** For each $d \in \mathbb{N}$, the sets $\mathcal{PT}^d$ of **probabilistic Böhm trees of depth** d and $\mathcal{VT}^d$ of **value trees of depth** d are defined as:

$$\begin{aligned}
\mathcal{VT}^0 &= \emptyset, \\
\mathcal{VT}^{d+1} &= \left\{ (y, (T_n)_{n\in\mathbb{N}}) \mid y \in \mathrm{Var} \text{ and } \forall n \in \mathbb{N},\ T_n \in \mathcal{PT}^d \right\} \text{ and} \\
\mathcal{PT}^d &= \left\{ T : \mathcal{VT}^d \to [0,1] \mid \textstyle\sum_{t\in\mathcal{VT}^d} T(t) \leq 1 \right\}.
\end{aligned}$$

We can then assign trees to individual terms:

▶ **Definition 2.** Given $M \in \Lambda^+$ and $d \in \mathbb{N}$, its **probabilistic Böhm tree of depth** d is the tree $\mathrm{PT}^d(M) \in \mathcal{PT}^d$ defined as follows:

$$\begin{aligned}
\mathrm{PT}^d(M) : \mathcal{VT}^d &\longrightarrow [0,1] \\
t &\longmapsto \Pr(M \to \{H \text{ hnf} \mid \mathrm{VT}^d(H) = t\})
\end{aligned}$$

where for any hnf $H = \lambda z_0 \dots z_{n-1}.y\ P_0 \dots P_{k-1}$, the **value tree of depth** d **of** H is defined as

$$\mathrm{VT}^d(H) = \left(y, \left(\mathrm{PT}^{d-1}(P_0), \dots, \mathrm{PT}^{d-1}(P_{k-1}), \mathrm{PT}^{d-1}(x_n^d), \dots\right)\right).$$

Consider for example the term $M_1 = \lambda xy.x\ (y +_{\frac{1}{3}} (\lambda z.z))$, a head-normal form. Figure 1a outlines the first steps in the construction of its value tree of depth d, for some fixed $d \geq 2$; note that we use the symbol δ_t to denote the distribution in which t has probability 1, and all other trees 0.

Infinitely extensional probabilistic Böhm trees precisely characterise observational equivalence in Λ^+; writing $M =_{\mathrm{PT}} N$ if for every $d \in \mathbb{N}$, $\mathrm{PT}^d(M) = \mathrm{PT}^d(N)$, we have:

▶ **Theorem 3** (Leventis [16]). *For any $M, N \in \Lambda^+$, $M =_{obs} N$ if and only if $M =_{PT} N$.*

So infinitely extensional probabilistic Böhm trees provide a *fully abstract* interpretation of the probabilistic λ-calculus. We will see now that similar trees arise as *probabilistic strategies* when interpreting λ-terms in a denotational games model.

2.3 Strategies and event structures

Moving towards our game semantics of Λ^+, we will first introduce our probabilistic strategies as a more economical, syntax-free presentation of probabilistic Böhm trees. The usual correspondence between Böhm trees and innocent strategies [12, 13] is thus naturally extended to the probabilistic and nondeterministic case.

First, we notice that the precise name given to variables in *e.g.* Figure 1a does not matter. Techniques like De Bruijn levels or indices do not apply here since we abstract infinitely many variables at each level – however, a variable occurrence is uniquely identified by a *pointer* to the node where it was abstracted, along with a *number* n, expressing that the variable

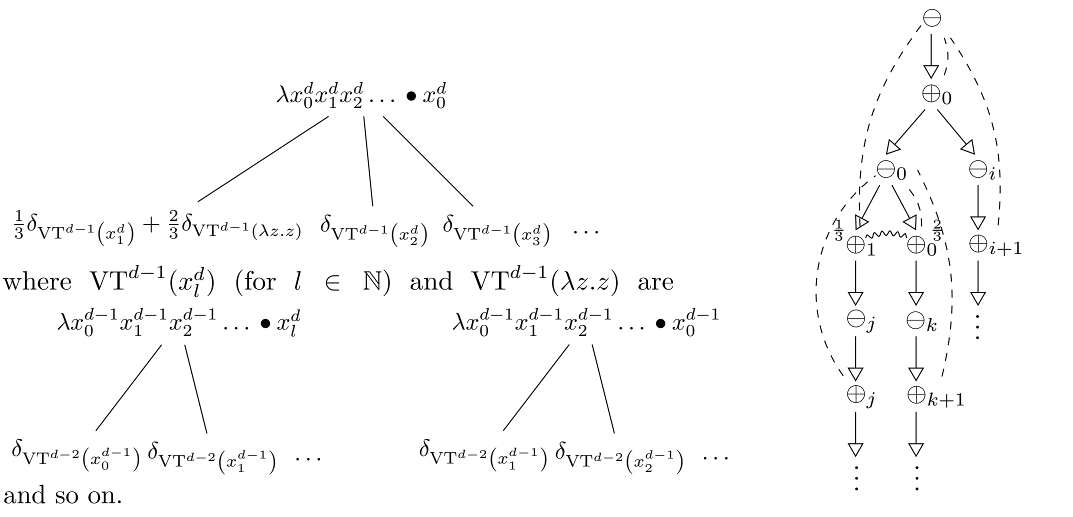

(a) As a value tree of depth $d \geq 2$. **(b)** As a probabilistic strategy.

Figure 1 Two interpretations of the term $M_1 = \lambda xy.x\ (y +_{\frac{1}{3}} (\lambda z.z))$.

was the $(n+1)$-th introduced at this node. For example, the variable x_0^d is expressed with a pointed to the initial node, along with number 0. As a consequence of this representation, we can omit the abstractions: at each node, there are always countably many variables being introduced, and their name does not matter as they will be referred to differently.

Next, we split each node of the Böhm tree into two: first a node intuitively carrying the abstractions (the target of pointers – we refer to these nodes as *negative*), and one carrying the variable occurrence (the source of pointers – we refer to those as *positive*). Besides bringing us closer to games, this allows us to easily distinguish the two kinds of branching of probabilistic Böhm trees. The different arguments of a variable node form a *negative* branching: each comes with its own (implicit) distinct set of fresh variables, and a subtree (by convention, we annotate by n the negative node corresponding to the nth argument). In contrast, for a probabilistic choice such as $\frac{1}{3}\delta_{\mathrm{VT}^{d-1}(x_1^d)} + \frac{2}{3}\delta_{\mathrm{VT}^{d-1}(\lambda z.z)}$ in Figure 1a, the two subtrees start by defining the same variables – so instead we represent this using a *positive* branching, where we further annotate the first node of each branch with its probability.

Altogether, and ignoring the wiggly line $\rightsquigarrow$ for now, the reader may check that the diagram of Figure 1b matches the Böhm tree of Figure 1a according to these conventions (the correspondence will be made formal in Section 4). Read from top to bottom, these diagrams have an interactive flavour: they describe the actions of a player $\oplus$ depending on those of its opponent $\ominus$. Our formalisation in terms of *strategies* will follow this intuition.

Probabilistic Böhm trees as probabilistic event structures

Now, we formalise the representation introduced above as a *probabilistic strategy* in the sense of [24], *i.e.* a form of probabilistic event structure. In this section we only provide this as a static representation, and leave the mechanism to *compose* strategies for Section 3. Our strategies (such as the one of Figure 1b) involve a partial order: the *dependency relation* (going from top to bottom); a relation $\rightsquigarrow$ indicating conflict and generated by probabilistic choice; and an annotation for probabilities. These are naturally formalised as *probabilistic concurrent strategies* [24] (though for the purposes of this paper we will only make use of *sequential* such strategies). We first recall the definition of event structures.

▶ **Definition 4.** An **event structure** [22] is a tuple $(E, \leq, \mathrm{Con})$ where E is a set of **events**, $\leq$ a partial order indicating **causal dependency**, and Con a non-empty set of **consistent** finite subsets of E, such that

$[e] = \{e' \mid e' \leq e\}$ is finite for all $e \in E$

$\{e\} \in \mathrm{Con}$ for all $e \in E$

$Y \subseteq X \in \mathrm{Con} \implies Y \in \mathrm{Con}$

$X \in \mathrm{Con}$ and $e \leq e' \in X \implies X \cup \{e\} \in \mathrm{Con}$.

The event structures we consider additionally have a polarity function $\mathrm{pol} : E \to \{+, -\}$ indicating for each event whether it is a move of Player (+) or Opponent (−). We call them **event structures with polarity** (esps).

We fix some notation. Write $e \rightarrowtriangle e'$ for **immediate causality**, *i.e.* $e < e'$ with no events in between. Write $\mathcal{C}(E)$ for the set of finite **configurations** of E, *i.e.* those finite $x \subseteq E$ such that $x \in \mathrm{Con}$ and x is down-closed: if $e \leq e' \in x$ then $e \in x$. If E has polarity, we sometimes annotate an event e to specify its polarity, as in e^+, e^-. If $x, y \in \mathcal{C}(E)$, write $x \subseteq^+ y$ (resp. $x \subseteq^- y$) if $x \subseteq y$ and every event in $y \setminus x$ has positive (resp. negative) polarity.

Ignoring probabilities and pointers, the diagram of Figure 1b is an esp: $\leq$ is the transitive reflexive closure of $\rightarrowtriangle$, and consistent sets are those finite sets whose down-closure do not contain two events related by the *immediate conflict* $\sim\!\sim\!\sim$. We now equip esps with probabilities, which comes in the form of a $[0, 1]$-valued function called a *valuation*.

For the forest-like event structures required to represent probabilistic λ-terms, it suffices to fix, for each Opponent event, a probability distribution on the Player events that immediately follow, as in Figure 1b. But to compose them we apply the more general machinery of [24], where valuations assign a coefficient to each *configuration* and not simply to each event. For $x \in \mathcal{C}(E)$, the coefficient $v(x)$ is the probability that the configuration x will be *reached* in an execution, provided the Opponent moves in x occur. The following definition is from [24]:

▶ **Definition 5.** A **probabilistic event structure with polarity** consists of an esp $(E, \leq, \mathrm{Con}, \mathrm{pol})$ and a **valuation**, that is, a map $v : \mathcal{C}(E) \to [0, 1]$ satisfying

- $v(\emptyset) = 1$;
- if $x \subseteq^- y$, then $v(x) = v(y)$; and
- if $y \subseteq^+ x_1, \ldots, x_n$, then

$$v(y) \geq \sum_I (-1)^{|I|+1} \, v\left(\bigcup_{i \in I} x_i\right)$$

where I ranges over non-empty subsets of $\{1, \ldots, n\}$ such that $\bigcup_{i \in I} x_i$ is a configuration.

Leaving aside pointers the diagram of Figure 1b represents a probabilistic esp, setting the valuation of a configuration x to be $\frac{1}{3}$ (resp. $\frac{2}{3}$) if it contains the event annotated with $\frac{1}{3}$ (resp. $\frac{2}{3}$), and 1 otherwise – a configuration cannot contain both labelled events.

Probabilistic strategies are certain probabilistic esps, equipped with a *labelling map* into the *game* they play on. Games are themselves esps, with the following particular shape:

▶ **Definition 6.** An **arena** is an esp A which is

- *forest-shaped*: if $a, b, c \in A$ with $a \leq b$ and $c \leq b$ then $a \leq c$ or $c \leq a$; and

- *alternating*: if $a \rightarrowtriangle b$ then $\mathrm{pol}(a) \neq \mathrm{pol}(b)$.
- *race-free*: if $x \in \mathcal{C}(A)$ has $x \subseteq^- y \in \mathcal{C}(A)$ and $x \subseteq^+ z \in \mathcal{C}(A)$, then $y \cup z \in \mathcal{C}(A)$.

Usually in game semantics, arenas represent *types*. For our untyped language, strategies representing terms all play on a *universal arena* U, introduced soon. For now though, we define a *probabilistic strategy* playing on arbitrary arena A as an esp, labelled by A.

▶ **Definition 7.** A **probabilistic strategy** on A consists of a probabilistic esp S, and a labelling function $\sigma : S \to A$ on events, preserving polarity, and such that:

(1) σ *preserves configurations:* for every $x \in \mathcal{C}(S)$, $\sigma x \in \mathcal{C}(A)$;
(2) σ *is locally injective*: if $s, s' \in x \in \mathcal{C}(S)$ and $\sigma s = \sigma s'$, then $s = s'$;
(3) σ is **receptive**: for $x \in \mathcal{C}(S)$, if $\sigma x \subseteq^- y \in \mathcal{C}(A)$, there is a unique $x \subseteq x' \in \mathcal{C}(S)$ such that $\sigma x' = y$;
(4) σ is **courteous**: for $s, s' \in S$, if $s \rightarrowtriangle_S s'$ and if $\mathrm{pol}(s) = +$ or $\mathrm{pol}(s') = -$, then $\sigma s \rightarrowtriangle_A \sigma s'$.

Conditions (1) and (2) express that σ is a **map of event structures** from S to A. Conditions (3) and (4) are there to restrict the behaviour of Player: they prevent any further constraints from being put on Opponent events than those already specified by the game.

The diagram of Figure 1b presents a probabilistic strategy $\sigma : S \to A$ – or more precisely the diagram presents S, with the pointers being representations of the immediate dependency in A of positive moves (though we do not display A for lack of space).

Winskel [24], building on previous work [20], showed how to *compose* probabilistic strategies and organise them into a category. But his games are affine, and cannot deal with the replication of resources. In recent work [5], we have extended probabilistic strategies with *symmetry*, that augments the expressivity of esps by allowing interchangeable copies of the same event. In the next section we introduce probabilistic strategies with symmetry, and give the interpretation of Λ^+. Because of this replication of resources the interpretation of the term M_1 of Figure 1 will be an *expansion* of Figure 1b, taking into account Opponent's replications – and in general, our correspondence theorem will associate a probabilistic Böhm tree with its expansion in that sense, formulated as a probabilistic strategy.

3 Game semantics for Λ^+

In this section we construct our game semantics for Λ^+. The category of games we use is close to our earlier concurrent games model of probabilistic PCF [5], in which we introduce a universal arena inspired from [13].

3.1 Games and strategies with symmetry

Symmetry in event structures [23] can be presented via *isomorphism families*:

▶ **Definition 8.** An **isomorphism family** on an event structure E is a set $\widetilde{E}$ of bijections between configurations of E, such that:

- $\widetilde{E}$ contains all identity bijections, and is closed under composition and inverse of bijections.
- For every $\theta : x \cong y \in \widetilde{E}$ and $x' \in \mathcal{C}(E)$ such that $x' \subseteq x$, then $\theta|_{x'} \in \widetilde{E}$.
- For every $\theta : x \cong y \in \widetilde{E}$ and every extension $x \subseteq x' \in \mathcal{C}(E)$, there exists a (non-necessarily unique) $y \subseteq y' \in \mathcal{C}(E)$ and an extension $\theta \subseteq \theta'$ such that $\theta' : x' \cong y' \in \widetilde{E}$.

As usual [23], it follows from these axioms that any $\theta \in \widetilde{E}$ is an order-isomorphism, *i.e.* preserves and reflects the order. An **event structure with symmetry** is a pair $(E, \widetilde{E})$, with $\widetilde{E}$ an isomorphism family on E. If E has polarity, then we ask that every $\theta \in \widetilde{E}$ preserves it, and call $(E, \widetilde{E})$ an **event structure with symmetry and polarity** (essp).

We illustrate this definition by presenting as an essp the *universal arena* — the game that Λ^+ strategies will play on. It is an infinitely deep tree, with at every level, ω available moves, corresponding to calls from one of the players to a variable in context. There are ω 'symmetric' copies of each move. Formally:

▶ **Definition 9.** The esp $(U, \leq, \mathrm{Con}, \mathrm{pol})$ is defined as having:

- *events*: $U = (\mathbb{N} \times \mathbb{N})^*$, finite sequences of ordered pairs;
- *causality*: $s \leq t$ if s is a prefix of t;
- *consistency*: no conflicts, $\mathrm{Con} = \mathcal{P}_{\mathrm{fin}}(U)$;
- *polarity*: $\mathrm{pol}(s) = -$ if $|s|$ is even, $+$ if it is odd.

In a pair $(m, n) \in \mathbb{N} \times \mathbb{N}$, m represents the variable address (the subscript in Figure 1b) and n is the copy index of the move (not displayed in Figure 1b).

We now add symmetry to U, following the intuition that different copies of the same move should be interchangeable. The isomorphism family $\widetilde{U}$ is generated by an equivalence relation $\sim$ on events, defined as the smallest equivalence relation satisfying $s \sim s' \implies s \cdot (m, n) \sim s' \cdot (m, n')$ for any $s, s' \in U$ and $m, n, n' \in \mathbb{N}$. Then, a bijection $\theta : x \cong y$ between configurations of U is in $\widetilde{U}$ whenever for all $e \in x$, $e \sim \theta(e)$.

The elements of $\widetilde{U}$ are *reindexing bijections*, which may update the copy indices of moves in a configuration. In the sequel, we will identify strategies differing only by the choice of positive copy indices, hence we need to formally identify the bijections in $\widetilde{U}$ which do not affect Opponent's copy indices. Because of the dual nature of games we must do the same for Player; thus we define $\sim^+$ and $\sim^-$ to be the smallest equivalence relations on U satisfying:

$$\begin{aligned} s \sim^p s' &\implies s \cdot (m, n) \sim^p s' \cdot (m, n) \quad (\text{for } p \in \{+, -\}) \\ s \sim^+ s' \text{ and } |s| \text{ is even} &\implies s \cdot (m, n) \sim^+ s' \cdot (m, n') \\ s \sim^- s' \text{ and } |s| \text{ is odd} &\implies s \cdot (m, n) \sim^- s' \cdot (m, n') \end{aligned}$$

for any s, s', m, n, n'. Just like $\sim$ generates $\widetilde{U}$, the relations $\sim^+$ and $\sim^-$ generate isomorphism families $\widetilde{U}_+$ and $\widetilde{U}_-$, respectively.

In general, the compositional mechanism will require all arenas to come with similar data:

▶ **Definition 10.** A **$\sim$-arena** is a tuple $\mathcal{A} = (A, \widetilde{A}, \widetilde{A}_-, \widetilde{A}_+)$ with A an arena, and $\widetilde{A}$, $\widetilde{A}_-$, and $\widetilde{A}_+$ isomorphism families on A, such that

- $\widetilde{A}_-$ and $\widetilde{A}_+$ are subsets of $\widetilde{A}$;
- if $\theta \in \widetilde{A}_- \cap \widetilde{A}_+$ then θ is an identity bijection;
- if $\theta \in \widetilde{A}_-$ and $\theta \subseteq^- \theta' \in \widetilde{A}$ then $\theta' \in \widetilde{A}_-$ (where the notation $\subseteq^-$ makes sense since bijections preserve polarity);
- if $\theta \in \widetilde{A}_+$ and $\theta \subseteq^+ \theta' \in \widetilde{A}$ then $\theta' \in \widetilde{A}_+$.

In particular, $\sim$-arenas are certain *thin concurrent games*, in the terminology of [7, 8].

▶ **Lemma 11.** *$\mathcal{U} = (U, \widetilde{U}, \widetilde{U}_-, \widetilde{U}_+)$ is a $\sim$-arena.*

Strategies are in turn equipped with symmetry:

▶ **Definition 12.** A **probabilistic essp** is an essp $\mathcal{S}$ with a valuation $v : \mathcal{C}(S) \to [0,1]$, such that for every $\theta : x \cong y$ in $\widetilde{S}$, $v(x) = v(y)$. A **probabilistic $\sim$-strategy** on a $\sim$-arena A consists of a probabilistic essp $\mathcal{S}$, and a labelling $\sigma : \mathcal{S} \to \mathcal{A}$, such that:

(1) the underlying map $\sigma : S \to A$ is a strategy;
(2) σ preserves symmetry: if $\theta : x \cong y \in \widetilde{S}$ then $\sigma\theta : \sigma x \cong \sigma y$ defined as $\{(\sigma\, s, \sigma\, s') \mid (s, s') \in \theta\}$, is in $\widetilde{A}$ (that is, it is a **map of essps** $(S, \widetilde{S}) \to (A, \widetilde{A})$);
(3) σ is **$\sim$-receptive**: if $\theta \in \widetilde{S}$ and $\sigma\theta \subseteq^- \psi \in \widetilde{A}$, there is a unique $\theta \subseteq \theta' \in \widetilde{S}$ *s.t.* $\sigma\theta' = \psi$.
(4) $\mathcal{S}$ is **thin**: for $\theta : x \cong y$ in $\widetilde{S}$ with $x \subseteq^+ x \cup \{s\}$, there is a *unique* $t \in S$ *s.t.* $\theta \cup \{(s,t)\} \in \widetilde{S}$.

Finally, before we define our category of games and strategies with symmetry, let us say what it means for strategies to be the same *up to Player copy indices:*

▶ **Definition 13.** Probabilistic $\sim$-strategies $\sigma : \mathcal{S} \to \mathcal{A}$ and $\tau : \mathcal{T} \to \mathcal{A}$ are **weakly isomorphic** if there is an isomorphism of essps $\varphi : \mathcal{S} \to \mathcal{T}$, such that for any $x \in \mathcal{C}(S)$, $v_S(x) = v_T(\varphi x)$, and moreover the diagram

$$\begin{tikzcd} \mathcal{S} \arrow[r, "\varphi"] \arrow[d, "\sigma"'] & \mathcal{T} \arrow[dl, "\tau"] \\ \mathcal{A} & \end{tikzcd}$$

commutes *up to positive symmetry*, in the sense that for any $x \in \mathcal{C}(S)$, the set $\{(\sigma e, \tau(\varphi e)) \mid e \in x\}$ is (the graph of) a bijection in $\widetilde{A}_+$.

3.2 The category PG

We now define a category with objects the $\sim$-arenas, and morphisms probabilistic $\sim$-strategies.

Let us first define some constructions on games: if $\mathcal{A}$ is a $\sim$-arena, its **dual** $\mathcal{A}^\perp$ is the $\sim$-arena obtained by reversing the polarity of events in A, and swapping the positive and negative isomorphism families. If $\mathcal{A}$ and $\mathcal{B}$ are $\sim$-arenas, their **parallel composition** $\mathcal{A} \parallel \mathcal{B}$ is the tuple $(A \parallel B, \widetilde{A} \parallel \widetilde{B}, \widetilde{A}_- \parallel \widetilde{B}_-, \widetilde{A}_+ \parallel \widetilde{B}_+)$, where $A \parallel B$ is the esp with events $A + B$ (the tagged disjoint union), componentwise causal dependency and polarity, and consistent sets those of the form $X_A \parallel X_B$ for $X_A \in \mathrm{Con}_A$ and $X_B \in \mathrm{Con}_B$; and where the parallel composition $\widetilde{A} \parallel \widetilde{B}$ of isomorphism families $\widetilde{A}$ and $\widetilde{B}$ comprises bijections of the form $\theta : x_A \parallel x_B \cong y_A \parallel y_B$, defined as $\theta(1,a) = (1, \theta_A(a))$ and $\theta(2,b) = (2, \theta_B(b))$ for some $\theta_A : x_A \cong y_A$ and $\theta_B : x_B \cong y_B$ in the component iso families. Note that we will often make use of the parallel composition $\parallel_{i \in I} \mathcal{A}_i$ of a family of $\sim$-arenas; it is defined analogously.

With that in place, a **probabilistic $\sim$-strategy from $\mathcal{A}$ to $\mathcal{B}$** is a probabilistic $\sim$-strategy on the $\sim$-arena $\mathcal{A}^\perp \parallel \mathcal{B}$. Given $\sigma : \mathcal{S} \to \mathcal{A}^\perp \parallel \mathcal{B}$ and $\tau : \mathcal{T} \to \mathcal{B}^\perp \parallel \mathcal{C}$, we can form their **interaction** as the pullback

$$\begin{tikzcd} & \mathcal{T} \circledast \mathcal{S} \arrow[dl, "\Pi_1"'] \arrow[dr, "\Pi_2"] \arrow[dd, phantom, "\vee"] & \\ \mathcal{S} \parallel \mathcal{C} \arrow[dr, "\sigma \parallel \mathcal{C}"'] & & \mathcal{A} \parallel \mathcal{T} \arrow[dl, "\mathcal{A} \parallel \tau"] \\ & \mathcal{A} \parallel \mathcal{B} \parallel \mathcal{C} & \end{tikzcd}$$

in the category of event structures with symmetry (and *without* polarity). The interaction is *probabilistic*: for any configuration $x \in \mathcal{C}(T \circledast S)$, we set $v_{T \circledast S}(x) = v_S((\Pi_1 x)_S) \times v_T((\Pi_2 x)_T)$, where $(\Pi_1 x)_S$ is the S-component of $\Pi_1 x \in \mathcal{C}(S \parallel C)$, and likewise for $(\Pi_2 x)_T$. The resulting map $\tau \circledast \sigma : \mathcal{T} \circledast \mathcal{S} \to \mathcal{A} \parallel \mathcal{B} \parallel \mathcal{C}$ is not quite a probabilistic $\sim$-strategy, because σ and τ play on dual versions of $\mathcal{B}$, making ambiguous the polarity of some events.

So as in [20, 6], the **composition** of $\mathcal{S}$ and $\mathcal{T}$ is obtained after *hiding* those moves of the interaction which act as *synchronisation events* — the moves $e \in T \circledast S$ such that $(\tau \circledast \sigma)e = (2, b)$ for some $b \in B$. The remaining set of events (so-called *visible*) induces an event structure $T \odot S$ with all structure inherited from $T \circledast S$, and polarity induced from $A^\perp \parallel C$. Any configuration $x \in \mathcal{C}(T \odot S)$ has a **unique witness** $[x] \in \mathcal{C}(T \circledast S)$. The isomorphism family $\widetilde{T \odot S}$ comprises bijections $\theta : x \cong y$ such that there is $\theta' : [x] \cong [y]$ in $\widetilde{T \circledast S}$ with $\theta \subseteq \theta'$. We get a map $\tau \odot \sigma : T \odot S \to A^\perp \parallel C$ which satisfies all the conditions for a probabilistic $\sim$-strategy, with $v_{T \odot S}(x) = v_{T \circledast S}([x])$ for every $x \in \mathcal{C}(T \odot S)$.

Copycat

As usual in game semantics, the identity morphism on a $\sim$-arena $\mathcal{A}$ will be a probabilistic $\sim$-strategy $\text{cc}_\mathcal{A} : \mathbb{C}_\mathcal{A} \to \mathcal{A}^\perp \parallel \mathcal{A}$ called **copycat**, in which Player deterministically copies the behaviour of Opponent — so any Opponent move immediately triggers the corresponding Player move in the dual game, with probability 1. Formally, $\mathbb{C}_A$ has the same events, polarity, and consistent subsets as $A^\perp \parallel A$ and the extra immediate causal dependencies $\{((1, a), (2, a)) \mid a \in A, \text{pol}_{A^\perp}(a) = -\}$ and $\{((2, a), (1, a)) \mid a \in A, \text{pol}_A(a) = -\}$ (from this $\leq_{\mathbb{C}_A}$ is obtained by transitive closure). Copycat has an isomorphism family $\mathbb{C}_{\widetilde{A}}$ which we do not define here for lack of space (it can be found *e.g.* in [8]). Together with the valuation $v_{\mathbb{C}_A}(x) = 1$ for all $x \in \mathcal{C}(\mathbb{C}_A)$, this turns copycat into a probabilistic $\sim$-strategy.

Recall that strategies are considered up to *weak isomorphism* (Definition 13). Doing so crucially relies on the thinness axiom on strategies, which implies [8] that weak isomorphism is stable under composition, so that we may perform a quotient and retain a well-defined notion of composition. Though identity and associativity laws for strategies only hold up to isomorphism, the quotient will turn them into strict equalities. So as in [5], we have:

▶ **Lemma 14.** *There is a category* **PG** *having*

- *objects: $\sim$-arenas*
- *morphisms $\mathcal{A} \nrightarrow \mathcal{B}$: weak isomorphism classes of probabilistic $\sim$-strategies on $\mathcal{A}^\perp \parallel \mathcal{B}$.*

Categorical structure

PG itself is a *compact closed category*, but we are interested in the subcategory $\mathbf{PG}^-$, where $\sim$-arenas and strategies are **negative** (that is, all initial moves are negative), and strategies are moreover **well-threaded** (meaning that events in $\mathcal{S}$ depend on a *unique* initial move).

Let $\mathcal{A}$ and $\mathcal{B}$ be objects of $\mathbf{PG}^-$. Their **tensor product** $\mathcal{A} \otimes \mathcal{B}$ is simply defined as $\mathcal{A} \parallel \mathcal{B}$. The tensorial unit is the empty $\sim$-arena, and moreover the tensor is *closed*: the **function space** $\mathcal{A} \multimap \mathcal{B}$ has events those of $(\parallel_{\min(B)} A^\perp) \parallel B$ with same polarity. The causal dependency is induced, with extra causal links $\{((2, b), (1, (b, a))) \mid b \in \min(B), a \in A\}$. The function $\chi : (A \multimap B) \to A^\perp \parallel B$ defined as $(1, (b, a)) \mapsto (1, a)$ and $(2, b) \mapsto (2, b)$ allows us to characterise consistent sets and iso families concisely: $\text{Con}_{A \multimap B}$ is defined as the largest set making χ a map of esps, and an order-isomorphism θ between configurations of $A \multimap B$ is in $\widetilde{A \multimap B}$ iff $\chi\theta \in \widetilde{\mathcal{A}^\perp} \parallel \mathcal{B}$. $\mathbf{PG}^-$ also has **cartesian products**, with $\mathcal{A} \,\&\, \mathcal{B}$ defined as $\mathcal{A} \parallel \mathcal{B}$, only with consistent sets restricted to those of $\mathcal{A} \parallel \emptyset$ and $\emptyset \parallel \mathcal{B}$. The rest of the structure, including symmetry, is induced from $\mathcal{A} \parallel \mathcal{B}$ by restriction.

Finally there is a **linear exponential comonad** [18] ! on $\mathbf{PG}^-$. Given $\mathcal{A} \in \mathbf{PG}^-$, the $\sim$-arena $!\mathcal{A}$ is an expanded version of $\mathcal{A}$ with countably many copies of every move. Accordingly, the esp $!A$ is simply $\parallel_{i \in \omega} A$, and the bijections in $\widetilde{!A}$ are those $\theta : \parallel_{i \in I} x_i \cong \parallel_{j \in J} y_j$ such that there exists a permutation $\pi : I \cong J$ and bijections $\theta_i \in \widetilde{A}$ with $\theta((i, a)) = (\pi i, \theta_i a)$

for all $(i,a) \in \|_{i\in I}\, x_i$. Recall that $\mathcal{A}$ is negative, so the set $\widetilde{!A}_+$ of positive bijections (those in which only Player moves are reindexed) comprises those $\theta \in \widetilde{!A}$ for which $I = J$ and $\pi : I \to J$ is the identity function, and such that each $\theta_i \in \widetilde{A}_+$. On the other hand, bijections in $\widetilde{!A}_-$ can consist of any $\pi : I \cong J$, so long as $\theta_i \in \widetilde{A}_-$ for all i.

We leave out all further details of the categorical structure of $\mathbf{PG}^-$, including the various constructions on morphisms. It can be shown that $\mathbf{PG}^-$, together with the data above, is a model of Intuitionistic Linear Logic. From here it is standard that the Kleisli category for ! is a ccc:

▶ **Lemma 15.** *There is a cartesian closed category* $\mathbf{PG}^-_!$ *having*

- *objects: negative $\sim$-arenas*
- *morphisms $\mathcal{A} \nrightarrow \mathcal{B}$: (weak isomorphism classes of) negative and well-threaded probabilistic $\sim$-strategies on $!\mathcal{A}^\perp \parallel \mathcal{B}$.*

With a slight abuse of notation, we shall keep using $\odot$ for composition in the Kleisli category $\mathbf{PG}^-_!$. We use the following notations for the cartesian closed structure: $\mathcal{A} \Rightarrow \mathcal{B}$ is the function space $!\mathcal{A} \multimap \mathcal{B}$, cur is the bijection $\mathbf{PG}^-_!(\mathcal{A} \,\&\, \mathcal{B}, \mathcal{C}) \cong \mathbf{PG}^-_!(\mathcal{A}, \mathcal{B} \Rightarrow \mathcal{C})$, and $\mathrm{ev}_{\mathcal{A},\mathcal{B}} : (\mathcal{A} \Rightarrow \mathcal{B}) \,\&\, \mathcal{A} \nrightarrow \mathcal{B}$ is the evaluation morphism.

3.3 Interpretation of Λ^+

We finally come to our interpretation of Λ^+ terms as probabilistic strategies. We start by imposing one key new condition on strategies: *sequential innocence*. The cut-down model will be closer to the language, allowing us to prove a correspondence result in Section 4. We assume from now on that all strategies are negative and well-threaded:

▶ **Definition 16.** A probabilistic $\sim$-strategy $\sigma : \mathcal{S} \to \mathcal{A}$ is **sequential innocent** if

- a subset $X \subseteq S$ is a configuration *if and only if* it is an Opponent-branching tree (that is, causality is tree-shaped and if $a \rightarrowtriangle b$ and $a \rightarrowtriangle c$ in X then $\mathrm{pol}(a) = +$) and $\sigma X \in \mathcal{C}(A)$;
- for every $x, y, z \in \mathcal{C}(S)$ such that $x = y \cap z$ and $y \cup z \in \mathcal{C}(S)$, either $v(x) = 0$ or

$$\frac{v(y \cup z)}{v(x)} = \frac{v(y)}{v(x)} \frac{v(z)}{v(x)}.$$

Less formally, innocence forces the independence (causal and probabilistic) of Opponent-forking branches of the strategy. Sequential innocent probabilistic $\sim$-strategies are closed under composition, stable under weak isomorphism, and copycat verifies all conditions, so we can consider the subcategory $\mathbf{PG}^{\mathrm{si}}_!$ of $\mathbf{PG}_!$ consisting of those strategies. It is easy to check that $\mathbf{PG}^{\mathrm{si}}_!$ is still a ccc; it is the category we will use to interpret Λ^+, and in what follows we refer to $\mathbf{PG}^{\mathrm{si}}_!$-strategies simply as Λ^+-strategies.

A reflexive object

Recall the $\sim$-arena $\mathcal{U}$ defined in 3.1. It is a **reflexive object**, meaning that there are maps $\lambda \in \mathbf{PG}^{\mathrm{si}}_!(\mathcal{U} \Rightarrow \mathcal{U}, \mathcal{U})$ and $\mathrm{app} \in \mathbf{PG}^{\mathrm{si}}_!(\mathcal{U}, \mathcal{U} \Rightarrow \mathcal{U})$ such that $\mathrm{app} \odot \lambda = \mathrm{id}_{\mathcal{U} \Rightarrow \mathcal{U}}$. It is easy to see that there is an isomorphism of essps $\rho : \mathcal{U} \cong \mathcal{U} \Rightarrow \mathcal{U}$. To turn this into a isomorphism is $\mathbf{PG}^{\mathrm{si}}_!$, we can lift it to a copycat-like strategy which "plays following ρ". Details of this lifting are omitted but can be found in [8].

Closed terms of the probabilistic λ-calculus are interpreted as probabilistic strategies on $\mathcal{U}$. Open terms M with free variables in Γ are interpreted as Λ^+-strategies $[\![M]\!]^\Gamma : \mathcal{U}^\Gamma \nrightarrow \mathcal{U}$, where $\mathcal{U}^\Gamma = \&_{x\in\Gamma}\mathcal{U}$. The interpretation of the λ-calculus constructions is standard, using that $\mathcal{U}$ is a reflexive object in a ccc:

$$\begin{aligned}
[\![x]\!]^\Gamma &= \pi_x, \text{ the } x^{\text{th}} \text{ projection}\\
[\![\lambda x.M]\!]^\Gamma &= \lambda \odot \mathrm{cur}([\![M]\!]^{\Gamma,x})\\
[\![MN]\!]^\Gamma &= \mathrm{ev}_{\mathcal{U},\mathcal{U}} \odot \langle \mathrm{app} \odot [\![M]\!]^\Gamma, [\![N]\!]^\Gamma\rangle
\end{aligned}$$

In order to give an interpretation to the probabilistic choice operator, we must define the sum of two strategies. Let $\sigma : \mathcal{S} \to (\mathcal{U}^\Gamma)^\perp \parallel \mathcal{U}$ and $\tau : \mathcal{T} \to (\mathcal{U}^\Gamma)^\perp \parallel \mathcal{U}$ be Λ^+-strategies, and let $p \in [0,1]$. The essp $\mathcal{S} +_p \mathcal{T}$ has a unique initial Opponent move (as do $\mathcal{S}$ and $\mathcal{T}$ — *wlog* call this move ε), and continues as either $\mathcal{S}$ or $\mathcal{T}$ non-deterministically. That is, it has events $\{\varepsilon\} \uplus (S \setminus \{\varepsilon\}) \uplus (T \setminus \{\varepsilon\})$, and all structure induced from $\mathcal{S}$ and $\mathcal{T}$, with $X \in \mathrm{Con}_{S+_pT}$ iff $X \in \mathrm{Con}_S$ or $X \in \mathrm{Con}_T$. We define $v_{S+_pT}(x)$ to be 1 if $x = \emptyset, \{\varepsilon\}$, $pv_S(x)$ if $x \in \mathcal{C}(S)$, and $(1-p)v_T(x)$ if $x \in \mathcal{C}(T)$. The obvious map $\sigma +_p \tau : \mathcal{S} +_p \mathcal{T} \to (\mathcal{U}^\Gamma)^\perp \parallel \mathcal{U}$ is a Λ^+-strategy, and the interpretation of the syntactic $+_p$ is simply $[\![M +_p N]\!]^\Gamma = [\![M]\!]^\Gamma +_p [\![N]\!]^\Gamma$. We have:

▶ **Theorem 17** (Adequacy). *For any $M \in \Lambda_0^+$, writing $\sigma : \mathcal{S} \to \mathcal{U}$ for $[\![M]\!]$, we have*

$$\mathrm{Pr}_\Downarrow(M) = \sum_{\substack{x\in\mathcal{C}(S)\\ |x^+|=1}} v_S(x),$$

where x^+ is the set of positive events of x.

We only state the result at this point; it will follow directly from the interpretation-preserving functor of Section 5 and the adequacy of the weighted relational model for Λ^+. A direct corollary of Theorem 17 is the following soundness result:

▶ **Lemma 18** (Soundness). *For any $M, N \in \Lambda^+$ with free variables in Γ, if $[\![M]\!]^\Gamma = [\![N]\!]^\Gamma$ then $M =_{obs} N$.*

In fact we will prove in Section 5 that the converse, *full abstraction*, also holds modulo a mild (effective) quotient. It will also follow that the weighted relational model itself is also fully abstract, which was open. These facts rely on Leventis' result [16] along with the formal correspondence between strategies and Böhm trees, to which we now move on.

4 The Correspondence Theorem

In [13], the authors prove an *exact correspondence theorem* for the pure λ-calculus: infinitely extensional Böhm trees precisely correspond to deterministic innocent strategies on a universal arena. They work in a different games framework, but the analogous phenomenon occurs in ours (the main technical difference, if we were to conduct the proof in the deterministic case, would be the explicit duplication of moves: our strategies are *expanded*, in order to accommodate Opponent's choice of copy index for every move).

For Λ^+ however, the correspondence is not so exact: although terms M and $M +_p M$ have the same probabilistic Böhm tree, they have different interpretations in $\mathbf{PG}_!^{\mathrm{si}}$, where each probabilistic choice is recorded as an explicit branching point.[2] In what follows, we identify a class of *Böhm tree-like* probabilistic strategies for which the exact correspondence

[2] In particular, $\mathbf{PG}_!^{\mathrm{si}}$ does not yield a *probabilistic λ-theory* in the sense of Leventis [16].

does hold, and we show that any strategy can be reduced to a Böhm tree-like one. Two strategies can then be considered equivalent if they reduce to the same.

First, given a Λ^+-strategy $\sigma : \mathcal{S} \to \mathcal{U}$, define a relation $\approx$ on the events of S as the smallest equivalence relation such that if $s_1 \approx s_1'$, $s_1 \twoheadrightarrow s_2$, $s_1' \twoheadrightarrow s_2'$ and there is an order-isomorphism $\varphi : \{s \in S \mid s_2 \leq s\} \cong \{s' \in S \mid s_2' \leq s'\}$ such that $\sigma\, s \sim^+ (\sigma \circ \varphi)\, s$ for all $s \geq s_2$, then $s_2 \approx s_2'$. Informally, $\approx$ identifies events coming from the same syntactic construct in two copies of a term in an idempotent probabilistic sum, as in $M +_p M$ (where Opponent has played the same copy indices).

▶ **Definition 19.** We say σ is **Böhm tree-like** if it satisfies

(1) for every $x \in \mathcal{C}(S)$, $v_S(x) > 0$; and
(2) for every $s, s' \in S$, if $s \approx s'$ then $s = s'$.

In other words, a Böhm tree-like strategy is one with no redundant branches. Many Λ^+-strategies do not satisfy this property, but all can be reduced to one that does:

▶ **Definition 20.** Given a Λ^+-strategy $\sigma : \mathcal{S} \to \mathcal{U}$, let S_{bt} be the set of $\approx$-equivalence classes containing at least one event s such that $v_S([s]) > 0$ (where $[s]$ is the down-closure of s).

It is direct to turn S_{bt} into an essp $\mathcal{S}_{\mathrm{bt}}$ with structure induced by $\mathcal{S}$. The (partial) quotient map $f : \mathcal{S} \to \mathcal{S}_{\mathrm{bt}}$ is then used to *push-forward* the valuation, *i.e.*

$$v_{S_{\mathrm{bt}}}(x) = \sum_{\substack{y \in \mathcal{C}(S) \\ f y = x}} v_S(y).$$

Then, $\sigma_{\mathrm{bt}} : \mathcal{S}_{\mathrm{bt}} \to \mathcal{U}$ is a Böhm tree-like Λ^+-strategy. Write $\sigma =_{\mathrm{bt}} \tau$ when $\sigma_{\mathrm{bt}} = \tau_{\mathrm{bt}}$.

We can now make formal the connection between Λ^+-strategies and probabilistic Böhm trees. To do so we define a bijective map from the set of Böhm tree-like Λ^+-strategies of depth d on $(\mathcal{U}^\Gamma)^\perp \parallel \mathcal{U}$, to the set $\mathcal{PT}_d^\Gamma$ of probabilistic Böhm trees of depth d with free variables in Γ. Let us say first what we mean by the *depth* of a strategy:

▶ **Definition 21.** The **depth** of a Λ^+-strategy $\sigma : \mathcal{S} \to \mathcal{U}$, $\mathrm{depth}(\sigma)$, is the maximum number of Player moves in a chain $s_0 \twoheadrightarrow \cdots \twoheadrightarrow s_n$ in S, and ∞ if such chains have unbounded length.

We can show by induction on d:

▶ **Lemma 22.** *For every $d \in \mathbb{N}$ and every $\Gamma \subseteq_{fin}$ Var there is a bijection*

$$\Psi_\Gamma^d : \{\sigma_{bt} \mid \sigma \in \mathbf{PG}_!^{si}(\mathcal{U}^\Gamma, \mathcal{U}) \text{ and } \mathrm{depth}\,\sigma \leq d\} \xrightarrow{\cong} \mathcal{PT}_\Gamma^d.$$

Proof (sketch). In Section 2.3, we motivated the definition of probabilistic strategies via a geometric correspondence with probabilistic Böhm trees, to be expected in the light of standard definability results in game semantics.

However, probabilistic strategies differ from the picture of Section 2.3 due to the necessity for Player to acknowledge Opponent's replications, spawning countably many symmetric copies of branches starting with an Opponent move. It follows however from the axioms of symmetry that events differing only by Opponent's choice of copy indices have isomorphic futures. One can, with no loss of information, focus on a sub-strategy where Opponent performs no duplication, and apply the correspondence explained in Section 2.3. ◀

We now show that this bijection preserves the interpretation of Λ^+.

▶ **Theorem 23** (Correspondence theorem). *For any $M \in \Lambda^+$ and $d \in \mathbb{N}$, $\Psi^d_\Gamma((\llbracket M \rrbracket^d)_{bt}) = PT^d(M)$, where $\llbracket M \rrbracket^d$ is the maximal sub-strategy of $\llbracket M \rrbracket$ with depth $\leq d$.*

Proof (sketch). The proof is by induction on d, and follows a similar argument as in the non-probabilistic case [13], with the additional difficulty of dealing with *infinite width*: a probabilistic Böhm tree may be a probability distribution with infinite support, and the first level of Player moves in a probabilistic strategy may be infinite. One must therefore consider finite-width approximations.

Probabilistic strategies are traditionally ordered using a probabilistic version of the prefix order: given $\sigma : \mathcal{S} \to \mathcal{A}$ and $\tau : \mathcal{T} \to \mathcal{A}$ we say $\sigma \sqsubseteq \tau$ if $\mathcal{S} \subseteq \mathcal{T}$ (*i.e.* $S \subseteq T$ and all data is inherited), and for all $x \in \mathcal{C}(S)$, $v_S(x) \leq v_T(x)$. However the naive restriction of this order to the set of Böhm tree-like strategies is not sensible, because $\sigma \sqsubseteq \tau$ does not imply $\sigma_{\mathrm{bt}} \sqsubseteq \tau_{\mathrm{bt}}$. An alternative is given by Leventis [16, p. 111], who defines an order $\preccurlyeq$ on the set $\mathcal{PT}^d_\Gamma$, characterised in this setting as follows: $t \preccurlyeq t'$ iff there exists a strategy σ such that $\left(\Psi^d_\Gamma\right)^{-1}(t) =_{\mathrm{bt}} \sigma$ and $\sigma \sqsubseteq \left(\Psi^d_\Gamma\right)^{-1}(t')$. Intuitively, the branches of σ are those of $\left(\Psi^d_\Gamma\right)^{-1}(t)$, duplicated and assigned probability in such a way that they can be extended to those of $\left(\Psi^d_\Gamma\right)^{-1}(t')$ using the prefix order $\sqsubseteq$.

Under $\preccurlyeq$ the set $\mathcal{PT}^d_\Gamma$ is a cpo, and we also call $\preccurlyeq$ the corresponding order on the set of Böhm tree-like strategies (this automatically makes Ψ^d_Γ a continuous bijection).

Leventis proves the crucial property that for every term M there is a chain $t_0, t_1, \ldots$ of finite-width trees satisfying $\mathrm{PT}^d(M) = \bigvee t_i$. Replaying his argument in our game semantics, we show that the chain $\left(\Psi^d_\Gamma\right)^{-1}(t_i), i \in \mathbb{N}$ has lub $(\llbracket M \rrbracket^d)_{\mathrm{bt}}$. We conclude, because $\left(\Psi^d_\Gamma\right)^{-1}\left(\mathrm{PT}^d_\Gamma(M)\right) = \left(\Psi^d_\Gamma\right)^{-1}\left(\bigvee_{i\in\mathbb{N}} t_i\right) = \bigvee_{i \in I}\left(\Psi^d_\Gamma\right)^{-1}(t_i) = (\llbracket M \rrbracket^d)_{\mathrm{bt}}$. ◀

Using the correspondence it follows easily that:

▶ **Lemma 24.** *For any $M, N \in \Lambda^+$, $M =_{PT} N$ if and only if $\llbracket M \rrbracket =_{bt} \llbracket N \rrbracket$.*

▶ **Theorem 25** (Full abstraction). *The model $\mathbf{PG}^{si}_!/=_{bt}$ is fully abstract,* i.e. *$M =_{obs} N$ if and only if $\llbracket M \rrbracket =_{bt} \llbracket N \rrbracket$.*

5 Weighted Relational Semantics

In this final section, we consider the weighted relational model of Λ^+. It lives in the category $\mathbf{PRel}_!$ whose objects are sets and whose morphisms are certain matrices with coefficients in the set $\overline{\mathbb{R}}_+ = \mathbb{R}_+ \cup \{\infty\}$. This interpretation of probabilistic λ-terms was first suggested in [11], where authors consider the category $\mathbf{PCoh}_!$ of **probabilistic coherence spaces**, a refinement (using *biorthogonality*) of the model $\mathbf{PRel}_!$ presented here. $\mathbf{PCoh}_!$ has desirable properties (notably, all coefficients are finite) but because there is a faithful functor $\mathbf{PCoh}_! \to \mathbf{PRel}_!$ preserving the interpretation of Λ^+, all the results of [11] hold for the simpler model $\mathbf{PRel}_!$, which we focus on in this paper and proceed to define.

5.1 The weighted relational model of Λ^+

We use the notation $\mathbf{PRel}_!$ to indicate that the model is obtained as the Kleisli category for a comonad !, much like $\mathbf{PG}_!$. The underlying category $\mathbf{PRel}$ is a well-known model of intuitionistic linear logic (see *e.g.* [15]), but we skip its construction and give a direct presentation of $\mathbf{PRel}_!$:

▶ **Definition 26.** The category $\mathbf{PRel}_!$ is defined as follows:

- *objects*: sets;
- *morphisms from X to Y*: maps $\varphi : \mathcal{M}_{\mathrm{f}}(X) \times Y \to \overline{\mathbb{R}}_+$, where $\mathcal{M}_{\mathrm{f}}(X)$ is the set of **finite multisets** of elements of X;
- *composition*: for $\varphi \in \mathbf{PRel}_!(X, Y)$, $\psi \in \mathbf{PRel}_!(Y, Z)$, define $\psi \circ \varphi : \mathcal{M}_{\mathrm{f}}(X) \times Z \to \overline{\mathbb{R}}_+$ as

$$(\psi \circ \varphi)(m, c) = \sum_{p \in \mathcal{M}_{\mathrm{f}}(Y)} \psi_{p,c} \sum_{\substack{(m_b)_{b \in p} \\ \text{s.t. } m = \uplus m_b}} \prod_{b \in p} \varphi_{(m_b, b)}$$

  for every $m \in \mathcal{M}_{\mathrm{f}}(X)$ and $c \in Z$.
- *identity*: for any set X, and for any $m \in \mathcal{M}_{\mathrm{f}}(X)$ and $a \in X$, define

$$\mathrm{id}_X(m, a) = \begin{cases} 1 & \text{if } m = [a] \\ 0 & \text{otherwise.} \end{cases}$$

$\mathbf{PRel}_!$ is cartesian closed, with $X \,\&\, Y = X \uplus Y$ and $X \Rightarrow Y = \mathcal{M}_{\mathrm{f}}(X) \times Y$. There is a reflexive object $\mathcal{D}$ in $\mathbf{PRel}_!$, supporting the interpretation of Λ^+, and defined as the least fixed point of the operation mapping X to the set $\mathcal{M}_{\mathrm{f}}(X)^{(\omega)}$ of **quasi-finite** sequences of finite multisets over X, *i.e.* with all but finitely many elements equal to $[\,]$. Concretely, $\mathcal{D}$ is the lub of the chain $D_0, D_1, \ldots$ where $D_0 = \emptyset$ and $D_{i+1} = \mathcal{M}_{\mathrm{f}}(D_i)^{(\omega)}$ for all i. It is the case that $\mathcal{D} \cong \mathcal{D} \Rightarrow \mathcal{D}$; the set-theoretical bijection and its lifting to a $\mathbf{PRel}_!$ isomorphism can be found in [11].

Terms of Λ^+ are interpreted in the standard way, with $[\![M +_p N]\!]^{\Gamma}_{\mathbf{PRel}_!}(d) = p[\![M]\!]^{\Gamma}_{\mathbf{PRel}_!}(d) + (1-p)[\![N]\!]^{\Gamma}_{\mathbf{PRel}_!}(d)$ for every $d \in \mathcal{D}$. We have:

▶ **Theorem 27** (Adequacy [11]). *For any $M \in \Lambda^+_0$, the map $[\![M]\!]_{\mathbf{PRel}_!} : \mathcal{D} \to \overline{\mathbb{R}}_+$ satisfies*

$$\Pr_{\Downarrow}(M) = \sum_{d \in D_2} [\![M]\!]_{\mathbf{PRel}_!}(d).$$

5.2 Relational collapse

We now connect the two models *via* a functor $\downarrow : \mathbf{PG}^{\mathrm{si}}_! \to \mathbf{PRel}_!$, which intuitively *forgets* the causal information in a strategy, only remembering the states reached during the execution.

If $(E, \widetilde{E})$ is an event structure with symmetry, write $\cong$ for the equivalence relation on $\mathcal{C}(E)$ defined as $x \cong y$ if and only if there is $\theta : x \cong y$ in $\widetilde{E}$. For $\mathcal{A}$ an arbitrary negative $\sim$-arena, the set $\downarrow \mathcal{A}$ is then defined as the quotient $\{x \in \mathcal{C}(A) \mid x \text{ non-empty}\}/\cong$.

For any $\mathcal{A}, \mathcal{B}$, there is a bijection $\downarrow(\mathcal{A} \Rightarrow \mathcal{B}) \simeq \mathcal{M}_{\mathrm{f}}(\downarrow \mathcal{A}) \times \downarrow \mathcal{B}$, enabling morphisms of $\mathbf{PG}^{\mathrm{si}}_!$ to be mapped to those of $\mathbf{PRel}_!$: if $\sigma : \mathcal{S} \to\, !\mathcal{A} \Rightarrow \mathcal{B}$ is a Λ^+-strategy and $\mathbf{x} \in \downarrow(\mathcal{A} \Rightarrow \mathcal{B})$ (so $\mathbf{x}$ is an equivalence class of configurations), the set of **witnesses** of $\mathbf{x}$ is defined as $\mathrm{wit}_S(\mathbf{x}) = \{z \in \mathcal{C}(S) \mid \sigma z \in \mathbf{x} \text{ and the maximal moves of } z \text{ have polarity } +\}/\cong$. Because v_S is invariant under symmetry, we can transport σ to $\downarrow \sigma : \downarrow(\mathcal{A} \Rightarrow \mathcal{B}) \to \overline{\mathbb{R}}_+$ *via*

$$\downarrow \sigma(\mathbf{x}) = \sum_{\mathbf{z} \in \mathrm{wit}_S(\mathbf{x})} v_S(\mathbf{z})$$

for each $\mathbf{x} \in \downarrow(\mathcal{A} \Rightarrow \mathcal{B})$. One can then easily deduce from the *deadlock-free lemma* of [5]:

▶ **Lemma 28.** $\downarrow$ *is a functor* $\mathbf{PG}^{si}_! \to \mathbf{PRel}_!$.

Furthermore, $\downarrow$ preserves the interpretation of Λ^+ terms and is well-defined on the quotiented model $\mathbf{PG}^{\mathrm{si}}_! / =_{\mathrm{bt}}$:

▶ **Lemma 29.** *$\downarrow\mathcal{U} \cong \mathcal{D}$ and up to this iso, for any $M \in \Lambda^+$ we have $\downarrow[\![M]\!]_{\mathbf{PG}_!^{si}} = [\![M]\!]_{\mathbf{PRel}_!}$.*

▶ **Lemma 30.** *If $\sigma =_{bt} \tau$ then $\downarrow\sigma = \downarrow\tau$.*

Combining the previous two lemmas and the soundness theorem, we finally get:

▶ **Theorem 31** (Full abstraction)**.** *For any $M, N \in \Lambda^+$ with free variables in Γ , $M =_{obs} N$ if and only if $[\![M]\!]_{\mathbf{PRel}_!} = [\![N]\!]_{\mathbf{PRel}_!}$.*

6 Conclusion

An immediate corollary of Theorem 31 is that the probabilistic coherence space model of [11] is fully abstract, since $\mathbf{PCoh}_!$ and $\mathbf{PRel}_!$ induce the same equational theory on Λ^+ terms.

Interestingly, the results of this paper should further entail that the interpretation of Λ^+ in the simpler model of Danos and Harmer [10] is also fully abstract, since one can in principle map our strategies functorially to theirs. Note however that since it is not known how to state a notion of probabilistic innocence in Danos and Harmer's model, definability fails and the present work could not have been carried out there.

So using probabilistic concurrent games, we obtain probabilistic analogues of well-established results from the theory of the pure λ-calculus: the correspondence between Böhm trees and innocent strategies [13], and the full abstraction property of the relational model [4].

References

1. Patrick Baillot, Vincent Danos, Thomas Ehrhard, and Laurent Regnier. Timeless games. In *International Workshop on Computer Science Logic*, pages 56–77. Springer, 1997.
2. Hendrik Pieter Barendregt. *The lambda calculus*, volume 2. North-Holland Amsterdam, 1984.
3. Johannes Borgström, Ugo Dal Lago, Andrew D. Gordon, and Marcin Szymczak. A Lambda-Calculus Foundation for Universal Probabilistic Programming. *CoRR*, abs/1512.08990, 2015.
4. Antonio Bucciarelli, Thomas Ehrhard, and Giulio Manzonetto. Not enough points is enough. In *International Workshop on Computer Science Logic*, pages 298–312. Springer, 2007.
5. Simon Castellan, Pierre Clairambault, Hugo Paquet, and Glynn Winskel. The Concurrent Game Semantics of Probabilistic PCF. In *Logic in Computer Science (LICS), 2018 33rd Annual ACM/IEEE Symposium on*, page to appear. ACM/IEEE, 2018.
6. Simon Castellan, Pierre Clairambault, and Glynn Winskel. Symmetry in Concurrent Games. In *Proceedings of the Joint Meeting of the Twenty-Third EACSL Annual Conference on Computer Science Logic (CSL) and the Twenty-Ninth Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, page 28. ACM, 2014.
7. Simon Castellan, Pierre Clairambault, and Glynn Winskel. The Parallel Intensionally Fully Abstract Games Model of PCF. In *Logic in Computer Science (LICS), 2015 30th Annual ACM/IEEE Symposium on*, pages 232–243. IEEE, 2015.
8. Simon Castellan, Pierre Clairambault, and Glynn Winskel. Concurrent hyland-ong games. 2016. https://arxiv.org/abs/1409.7542.
9. Vincent Danos and Thomas Ehrhard. Probabilistic coherence spaces as a model of higher-order probabilistic computation. *Information and Computation*, 209(6):966–991, 2011.
10. Vincent Danos and Russell S Harmer. Probabilistic game semantics. *ACM Transactions on Computational Logic (TOCL)*, 3(3):359–382, 2002.

11 Thomas Ehrhard, Michele Pagani, and Christine Tasson. The computational meaning of probabilistic coherence spaces. In *Logic in Computer Science (LICS), 2011 26th Annual IEEE Symposium on*, pages 87–96. IEEE, 2011.

12 J Martin E Hyland and C-HL Ong. On full abstraction for PCF: I, II, and III. *Information and computation*, 163(2):285–408, 2000.

13 Andrew D Ker, Hanno Nickau, and C-H Luke Ong. Innocent game models of untyped λ-calculus. *Theoretical Computer Science*, 272(1-2):247–292, 2002.

14 Dexter Kozen. Semantics of probabilistic programs. In *Foundations of Computer Science, 1979., 20th Annual Symposium on*, pages 101–114. IEEE, 1979.

15 Jim Laird, Giulio Manzonetto, Guy McCusker, and Michele Pagani. Weighted relational models of typed lambda-calculi. In *Proceedings of the 2013 28th Annual ACM/IEEE Symposium on Logic in Computer Science*, pages 301–310. IEEE Computer Society, 2013.

16 Thomas Leventis. *Probabilistic lambda-theories*. PhD thesis, Aix-Marseille Université, 2016.

17 Thomas Leventis. Probabilistic Böhm Trees and Probabilistic Separation. In *Logic in Computer Science (LICS), 2018 33rd Annual ACM/IEEE Symposium on*. ACM/IEEE, 2018.

18 Paul-André Mellies. Categorical semantics of linear logic. *Panoramas et syntheses*, 27:15–215, 2009.

19 Paul-André Mellies and Samuel Mimram. Asynchronous games: innocence without alternation. In *International Conference on Concurrency Theory*, pages 395–411. Springer, 2007.

20 Silvain Rideau and Glynn Winskel. Concurrent strategies. In *Logic in Computer Science (LICS), 2011 26th Annual IEEE Symposium on*, pages 409–418. IEEE, 2011.

21 Nasser Saheb-Djahromi. Cpo's of measures for nondeterminism. *Theoretical Computer Science*, 12(1):19–37, 1980.

22 Glynn Winskel. *Events in Computation*. PhD thesis, 1980.

23 Glynn Winskel. Event structures with symmetry. *Electronic Notes in Theoretical Computer Science*, 172:611–652, 2007.

24 Glynn Winskel. Distributed probabilistic and quantum strategies. *Electronic Notes in Theoretical Computer Science*, 298:403–425, 2013.