Probabilistic game semantics

Vincent Danos* & Russell Harmer
Université Paris VII

Abstract

A category of HO/N-style games and probabilistic strategies is developed where the possible choices of a strategy are quantified so as to give a measure of the likelihood of seeing a given play. A 2-sided die is shown to be universal in this category, in the sense that any strategy breaks down into a composition between some deterministic strategy and that die. The interpretative power of the category is then demonstrated by delineating a Cartesian closed subcategory which provides a fully abstract model of a probabilistic extension of Idealized Algol.

1 Introduction

Consider the fixpoint expression $\mathbf{Y}(\lambda x.(\mathtt{1}\ \mathtt{or}\ x))$. This can have two very different meanings attached. If, by `or`, it is meant that both sides can happen with, say, equal chances then, almost surely, we'll get the result `1`. If, on the other hand, it is meant only that either can happen—but we don't know which—then the only consistent answer is to say that we *may* get `1` but we also *may* diverge. This latter meaning was investigated in a recent games model for nondeterminism by Harmer & McCusker [12]. In this paper, we construct a games model of the former by developing a suitable notion of *probabilistic* strategy.

One motivation for doing this is to give a means of understanding quantitative aspects of uncertainty that might arise in the description of a process. Dually, inherently probabilistic processes—such as randomized optimization procedures, optimal strategies in classical game theory or more general adaptive behavior—can be accounted for in this framework.

Another motivation is purely mathematical. Extending games semantics to the probabilistic case is a natural step to take, just as probabilistic strategies are extremely natural in Von Neumann matrix games. Indeed, the category of probabilistic strategies that we construct smoothly extends the deterministic games world as charted out by Hyland & Ong [13] and Nickau [20] and further developed by Abramsky & McCusker [3, 4, 5, 19], Abramsky, Honda & McCusker [2] and Laird [15]. The category also fits in very nicely with the basic concepts of probability theory. The interaction of a strategy

*address: Équipe Preuves Programmes Systèmes, UFR de Mathématiques, Université de Paris 7, 2 place Jussieu, 75251 Paris Cedex 05, France. tel: +33 1 44 27 86 79. fax: +33 1 44 27 61 48. email: `danos@logique.jussieu.fr`

and a counter-strategy on a game gives rise to a Borelian subprobability on the infinite plays. The category has a tensor product which corresponds to the classical product measure construction, and the composition of strategies uses the image measure under the composition function (which in our case is always measurable) to define the composite. A factorization result is shown, saying that any probabilistic strategy can be broken down into the composition of a deterministic one and the 2-sided die. In that sense, the 2-sided die is shown to be universal. This is a discrete version of the basic fact that any probability on $[0,1]$ can be reconstructed as the law of some random variable defined on the Lebesgue probability triple. In short, we give what seems the fundamentals of a higher-order probability theory in "interactive" form.

A case study To illustrate the interpretative power of our structure, we eventually zero in on a Cartesian closed subcategory which we show provides a fully abstract model of a probabilistic extension of Idealized Algol. The exact choice of probabilistic primitive used to extend Algol is not overly important since, by our factorization result, any probabilistic process can be programmed using only a fair coin. The recent shifting of conditions from plays to strategies and other advances in the field allow, at this point, the use of a pretty standard machinery: we simply restrict to *single-threaded* strategies—those that are copyable and discardable with respect to the tensor product—and show that homsets admit a ω-continuous CPO structure of which the basis is easily seen to be Algol-definable thanks to the factorization theorem and a previous theorem of Abramsky & McCusker [3]. For general reasons, this suffices to show full abstraction.

Generality Our construction in no way depends on a particular games model: we give it in sufficient generality so that it can be adapted to whatever particular language we already know a deterministic games model of. And that is already quite a lot since, apart from the purely functional PCF, there are models for control, ground and general references. That games semantics concepts adapt in this way gives a sense of the homogeneity of this semantic universe, in due contrast with traditional denotational semantics.

Related work The most obvious precursor to this work is the study of *probabilistic powerdomains* and their application to denotational semantics as in the early work of Saheb-Djahromi [22, 23] and of Jones & Plotkin [14] amongst others. More recently, Edalat *et al.* [8, 9] have investigated bisimulation in a probabilistic context. It would be interesting to consider more closely the connection between that work and our own as this remains rather unclear for now.

Another possible connection is with *exact real arithmetic* [7]. Intriguingly, several of the issues raised in our work, particularly those surrounding the factorization result, have a similar flavour to those raised in real arithmetic [21]. It would be interesting to pursue this line of investigation further and, indeed, to consider using game semantics to model a language such as Real PCF [10].

Future work One challenging application could be in probabilistic program analysis. Hankin and Malacaria's static analysis techniques derived from games semantics [16, 17, 18] could get on well with probabilities, and give tools in code optimization for instance. Less an application maybe, but still sort of practical, is the tempting question of whether the probabilistic observational equivalence can be shown to be decidable for our extension of Idealized Algol since, in that particular case, the so-called intrinsic equivalence in the intensional model has a simple concrete characterization. If solved to the positive, that could relate to formal methods in programming.

A more theoretical question is that of relating our model to Harmer & McCusker's model [12]. There is an obvious embedding of theirs into ours, based on an equal chances interpretation of non-determinism, but it isn't fully abstract, as the opening example clearly shows. Another possibility is the investigation of a "timeless" version of our model in the spirit of the Baillot, Danos, Ehrhard & Regnier [6] construction starting from a symmetrized games model of Linear Logic and ending in a polarized variant of relational semantics. The concurrent games of Abramsky & Mellies [1] and their polarized version by Mellies could prove useful here, since they are a sort of intermediate step between games and hypercoherences. Current investigation with Ehrhard shows that there are probabilistic coherence spaces that could bridge the gap, at least in the intuitionistic case.

2 Probabilistic games

Game semantics represents computation with "execution traces" that describe the interaction between a System and its Environment. The System and Environment are modelled as the *protagonists*, respectively Player and Opponent, in a game; an execution trace by a sequence of moves played by the two protagonists. To make these ideas precise, we need a few bits of notation. The set of strings over an alphabet Σ is written $\Sigma^\star$. If s and t are in $\Sigma^\star$, we write $s \sqsubseteq t$ if s is a prefix of t, denote the length of s by $|s|$ and, for $1 \leq i \leq |s|$, denote the ith symbol of s by s_i. The unique string of length 0 is written ε. If $\Sigma' \subseteq \Sigma$, we write $s \upharpoonright \Sigma'$ for the restriction of s to Σ', *i.e.* the subsequence obtained by erasing all symbols of s that aren't in Σ'.

2.1 Arenas and legal plays

The basic ground rules of a game are determined by the *arena* in which it is played. We formally define an ***arena*** as a triple $\langle M_A, \lambda_A, \vdash_A \rangle$ where:

- M_A is a countable set of ***moves***.
- $\lambda_A : M_A \to \{\mathsf{O},\mathsf{P}\} \times \{\mathsf{Q},\mathsf{A}\} \times \{\mathsf{I},\mathsf{n}\}$ is a ***labelling*** function designating each move m as an Opponent or a Player move, a Question or an Answer and as an Initial or non-initial move.

  We define λ_A^{OP}, λ_A^{QA} and λ_A^{In} as the first, second and third projections respectively from λ_A.
- $\vdash_A$ is a binary relation, known as ***enabling***, on M_A such that:

  (e1) if $m \vdash_A n$ then $\lambda_A^{\mathsf{OP}}(m) \neq \lambda_A^{\mathsf{OP}}(n)$ and $\lambda_A^{\mathsf{In}}(n) \neq \mathsf{I}$.

  (e2) if $\lambda_A^{\mathsf{In}}(m) = \mathsf{I}$ then $\lambda_A(m) = \mathsf{OQI}$.

  (e3) if $m \vdash_A n$ and $\lambda_A^{\mathsf{QA}}(n) = \mathsf{A}$ then $\lambda_A^{\mathsf{QA}}(m) = \mathsf{Q}$.

These conditions say that the protagonists enable each other's moves, not their own; that initial moves can't be enabled; that only Opponent can have initial moves and, moreover, they must be Questions; and that Answers must be enabled by Questions.

Flat arenas If X is a countable set, we define the ***flat arena*** over X by setting M_X to be $X + \{\mathsf{q}\}$; defining $\lambda_X(\mathsf{q}) = \mathsf{OQI}$ and, for $x \in X$, $\lambda_X(x) = \mathsf{PAn}$; and making $\mathsf{q} \vdash_X x$ for every $x \in X$. We denote by $\mathbf{C}$, $\mathbf{B}$ and $\mathbf{N}$ respectively the flat arenas over the one-point set $\{\mathsf{a}\}$, the two-point set $\{\mathsf{tt},\mathsf{ff}\}$ and the countably infinite set $\{0, 1, 2, \ldots\}$.

Legal plays A string $s \in M_A^\star$ can be endowed with a "justification structure" by equipping each non-initial move in s with a "pointer" pointing back to an earlier enabling move. This can be formalized with a function $f_s : \{1, \ldots, |s|\} \to \{0, \ldots, |s| - 1\}$ satisfying:

- $f_s(i) < i$, for $1 \leq i \leq |s|$.
- if $f_s(i) = 0$ then $\lambda_A^{\mathsf{In}}(s_i) = \mathsf{I}$.
- if $f_s(i) = j \neq 0$ then $s_j \vdash_A s_i$.

If $f_s(i) = j \neq 0$ then we say that s_j ***justifies*** s_i; if $f_s(i) = 0$ then we call s_i an ***initial occurrence***. If $f_s(i) = j_1$, $f_s(j_1) = j_2$, $\ldots$, $f_s(j_n) = 0$ then we say that s_i is ***hereditarily justified*** by s_{j_n}. A string equipped with such an f_s is called a ***justified string***.

A ***legal play*** is a justified string that strictly alternates between Opponent and Player:

$$\text{if } s = s_1 a b s_2 \text{ then } \lambda_A^{\mathsf{OP}}(a) \neq \lambda_A^{\mathsf{OP}}(b).$$

We write $\mathcal{L}_A$ for the set of all legal plays in the arena A; $\mathcal{L}_A^{\mathsf{even}}$ and $\mathcal{L}_A^{\mathsf{odd}}$ denote the obvious subsets so that Player made the last move in $s \in \mathcal{L}_A^{\mathsf{even}}$, *etc.*

If $s \in \mathcal{L}_A$, we write $\mathsf{ie}(s)$ for the set of all ***immediate extensions*** of s, *i.e.* those legal plays t such that $s \sqsubseteq t$ and $|t| = |s| + 1$. The ***current thread*** of $sa \in \mathcal{L}_A^{\mathsf{odd}}$, written $\lceil sa \rceil$, is defined to be the subsequence of sa consisting of all moves hereditarily justified by the hereditary justifier of a.

Product & arrow If A and B are arenas, we define their product $A \times B$ by:

- $M_{A\times B} = M_A + M_B$, the disjoint union.
- $\lambda_{A\times B} = [\lambda_A, \lambda_B]$, the copairing.
- $m \vdash_{A\times B} n$ iff $m \vdash_A n$ or $m \vdash_B n$.

This places A and B "side by side" with no chance of any interaction between them. The ***empty arena*** $\mathbf{1} = \langle \varnothing, \varnothing, \varnothing \rangle$ with sole legal play ε is the unit for this constructor.

Our other constructor is the arrow, defined by:

- $M_{A\Rightarrow B} = M_A + M_B$.
- $\lambda_{A\Rightarrow B} = [\langle \overline{\lambda}_A^{\mathsf{OP}}, \lambda_A^{\mathsf{QA}}, \overline{\lambda}_A^{\mathsf{In}} \rangle, \lambda_B]$, where
  - $\overline{\lambda}_A^{\mathsf{OP}}(m) = \mathsf{O}$ iff $\lambda_A^{\mathsf{OP}}(m) = \mathsf{P}$.
  - $\overline{\lambda}_A^{\mathsf{In}}(m) = \mathsf{n}$, for all $m \in M_A$.
- $m \vdash_{A\Rightarrow B} n$ iff $m \vdash_A n$ or $m \vdash_B n$ or $\lambda_B^{\mathsf{In}}(m) = \mathsf{I} = \lambda_A^{\mathsf{In}}(n)$.

In words, the initial moves of $A \Rightarrow B$ are the initial moves of B, the rôles of Opponent and Player are reversed in A and the (formerly) initial moves of A are now enabled by the (still) initial moves of B.

2.2 Strategies

A *strategy* is a kind of "rule book" saying which moves may be made by Player and with what probability. First of all, a ***prestrategy*** σ on an arena A is just a function $\sigma : \mathcal{L}_A^{\mathsf{even}} \to [0, \infty]$.

The ***traces*** of σ, which we denote by $\mathcal{T}_\sigma$, are those even-length legal plays assigned non-zero weight by σ, *i.e.* $\{s \in \mathcal{L}_A^{\mathsf{even}} \mid \sigma(s) > 0\}$. The ***domain*** of σ, written $\mathsf{dom}(\sigma)$, is those odd-length plays that are "reachable" by σ, *i.e.* $\{sa \in \mathcal{L}_A^{\mathsf{odd}} \mid s \in \mathcal{T}_\sigma\}$. Finally, given $s \in \mathsf{dom}(\sigma)$, we define the ***range*** of σ at s, written $\mathsf{rng}_\sigma(s)$, to be those immediate extensions of s that are in $\mathcal{T}_\sigma$, *i.e.* $\mathsf{ie}(s) \cap \mathcal{T}_\sigma$. This can be empty.

A prestrategy σ for A is a ***strategy*** iff

(p1) $\sigma(\varepsilon) = 1$;

(p2) if $sa \in \mathcal{L}_A^{\mathsf{odd}}$ then $\sigma(s) \geq \sum_{t \in \mathsf{ie}(sa)} \sigma(t)$.

Note that **(p2)** implies that σ is order reversing with respect to the prefix ordering on $\mathcal{T}_\sigma$ and the usual ordering on $[0, \infty]$. This means that $\mathcal{T}_\sigma$ is even-length prefix-closed. Furthermore, by **(p1)**, we get $\sigma(s) \leq 1$ for all $s \in \mathcal{L}_A$.

Basic constraints A strategy σ is ***deterministic*** iff for all $s \in \mathcal{T}_\sigma$, $\sigma(s) = 1$. Equivalently, σ takes values in $\{0, 1\}$. Note that **(p2)** asserts only an inequality. If σ is a strategy for which this is always in fact an equality, we say that it's a ***total*** strategy. A weaker form of determinism, known as ***subdeterminism***, just requires the trace set to be deterministic, *i.e.* if $sb \in \mathcal{T}_\sigma$ and $sb' \in \mathcal{T}_\sigma$ then $sb = sb'$. Any deterministic strategy is subdeterministic, but not vice versa. However, a total subdeterministic strategy is always deterministic.

Local probabilities Given $sa \in \mathsf{dom}(\sigma)$, define the ***conditional probability*** of sab given sa by:

$$\sigma(sab/sa) = \frac{\sigma(sab)}{\sigma(s)}$$

Since $sa \in \mathsf{dom}(\sigma)$, $\sigma(s) > 0$ so it's well-defined. By **(p2)**, this gives a subprobability on $\mathsf{rng}_\sigma(sa)$ for each $sa \in \mathsf{dom}(\sigma)$. Intuitively speaking, $\sigma(sab/sa)$ is the "die" that σ rolls for each $sa \in \mathsf{dom}(\sigma)$.

2.3 Examples of strategies

The fair coin The simplest example of a probabilistic strategy is probably the "fair coin" coin on the Boolean arena $\mathbf{B}$. Its trace set is just the set of even-length sequences of the form $\mathsf{q}b_1 \cdots \mathsf{q}b_n$ where each b_i is in $\{\mathsf{tt}, \mathsf{ff}\}$ and is justified by the immediately preceding move. We assign a probability to such an $s \in \mathcal{T}_{\mathsf{coin}}$ by $\mathsf{coin}(s) = (1/2)^{|s|/2}$, *e.g.* $\mathsf{coin}(\mathsf{q}\mathsf{tt}) = 1/2$, $\mathsf{coin}(\mathsf{q}\mathsf{ff}\mathsf{q}\mathsf{tt}) = 1/4$, $\mathsf{coin}(\mathsf{q}\mathsf{ff}\mathsf{q}\mathsf{ff}\mathsf{q}\mathsf{tt}) = 1/8$, *etc.* All local probabilities, such as $\sigma(\mathsf{q}\mathsf{ff}\mathsf{q}\mathsf{tt}/\mathsf{q}\mathsf{ff}\mathsf{q})$, are $1/2$.

The polydie Another useful probabilistic strategy is $\mathsf{polydie}$ on $\mathbf{N} \Rightarrow \mathbf{N}$. A typical play of $\mathsf{polydie}$ is:

$$\begin{array}{ccc}
\mathbf{N} & \Rightarrow & \mathbf{N} \\
 & & \mathsf{q} \\
\mathsf{q} & & \\
\mathsf{n} & & \\
 & & \mathsf{i}
\end{array}$$

where $0 \leq i < n$ and the probability assigned to this play is $1/n$. In other words, $\mathsf{polydie}$ asks for a number n and then rolls an n-sided die, returning something between 0 and $n-1$ with uniform probability.

Strategies on flat arenas as probability measures Plays in $\mathbf{B}$ of the form $\mathsf{qttqffqff}\ldots$, are in a 1-1 correspondence, denoted by $B(\cdot)$, with finite binary strings over $\{0,1\}$. Given such a binary string w, we set $V_w = \{x \in \{0,1\}^\omega | w \sqsubseteq x\}$. These V_ws are the basic open sets of the Cantor topology on infinite binary strings.

Given a total strategy σ on $\mathbf{B}$, we define $p_\sigma(V_{B(s)}) = \sigma(s)$, which standardly extends to a unique probability measure on the Borelian σ-algebra $\mathcal{B}$ generated by the V_ws. We can now see that **(p1)** says that the measure of the whole space is 1, while **(p2)** says that p is countably additive (or subadditive if σ isn't assumed total). Conversely, any probability measure on $\mathcal{B}$ defines a total stategy on $\mathbf{B}$.

Such a probability measure could be any discrete-time stochastic process with values in $\{0,1\}$ and, in particular, it needn't be a Markov process. An example of this is the *damped deviation die*, denoted $\mathsf{3d}$, a strategy where the larger the deviation between trues and falses so far, the likelier it is to return that value which diminishes the deviation. For any play such as $s = \mathsf{qttqffqff}\ldots$, let $d(s)$ be the difference between the number of trues and falses in s, suppose $d(s) \geq 0$ and define $\mathsf{3d}(s\mathsf{qtt}/s\mathsf{q}) = \frac{1}{d(s)+2}$. Hence, $\mathsf{3d}$ must be able to look arbitrarily far back in time.

2.4 Composing strategies

Let u be a finite string of moves from arenas A, B and C equipped with candidate "justification pointers", *i.e.* a strictly decreasing function $f_u : \{1,\cdots,|u|\} \to \{0,\cdots,|u|-1\}$. We define $u \upharpoonright B, C$ to be the subsequence of u where we delete all moves from A and suitably "renormalize" f_u to preserve the "pointers". We define $u \upharpoonright A, B$ similarly. We define $u \upharpoonright A, C$ by removing all moves from and "pointers" to B except in the case where $a \in M_A$ points to $b \in M_B$ which, in turn, points to $c \in M_C$, whence we make a point directly to c in $u \upharpoonright A, C$.

A ***legal interaction*** of arenas A, B and C is such a u satisfying $u \upharpoonright A, B \in \mathcal{L}_{A \Rightarrow B}$, $u \upharpoonright B, C \in \mathcal{L}_{B \Rightarrow C}$ and $u \upharpoonright A, C \in \mathcal{L}_{A \Rightarrow C}$. We write $\mathsf{int}(A, B, C)$ for the set of all legal interactions of A, B and C.

If $s \in \mathcal{L}^{\mathsf{even}}_{A \Rightarrow C}$, we define its set of ***witnesses*** by:

$$\mathsf{wit}(s) = \{u \in \mathsf{int}(A, B, C) \mid u \upharpoonright A, C = s\},$$

i.e. all legal interactions playing s "on the outside". Given strategies σ and τ for $A \Rightarrow B$ and $B \Rightarrow C$ respectively, we define the ***composite strategy***

$\sigma\,;\tau$ for $A \Rightarrow C$ by assigning a probability to $s \in \mathcal{L}_{A\Rightarrow C}$ with the following ***composition formula***.

$$\sigma\,;\tau(s) = \sum_{u \in \mathsf{wit}(s)} \sigma(u \upharpoonright A, B) \cdot \tau(u \upharpoonright B, C)$$

We sometimes write $(\sigma \,\|\, \tau)(u)$ as an abbreviation for $\sigma(u \upharpoonright A, B) \cdot \tau(u \upharpoonright B, C)$. The above sum is always countable because $\mathsf{int}(A, B, C)$ is.

This composition formula expresses exactly the same interaction mechanism as in the deterministic case with, in addition, σ and τ independently rolling their dice.

Well-defined composition It's clear that the composite of two prestrategies is itself a prestrategy (with the convention that $0 \cdot \infty = 0$). It's also easy to see that (**p1**) is preserved by composition. The lemma below does most of the rest of the work.

Lemma 2.1 If $sa \in \mathcal{L}^{\mathsf{odd}}_{A\Rightarrow C}$ and $ua \in \mathsf{int}(A, B, C)$ such that $u \in \mathsf{wit}(s)$, let W_{ua} be the set of all $u' \in \mathsf{int}(A, B, C)$ satisfying $ua \sqsubseteq u'$ and $u' \upharpoonright A, C \in \mathsf{ie}(sa)$, *i.e.* all those legal interactions starting like ua and witnessing an immediate extension of sa.

If $\sigma : A \Rightarrow B$ and $\tau : B \Rightarrow C$ then

$$(\sigma \,\|\, \tau)(u) \geq \sum_{u' \in W_{ua}} (\sigma \,\|\, \tau)(u').$$

We prove by induction that the interactions obtained by truncating W_{ua} at any depth satisfy the inequality. The idea is that, at any depth, some of the witnesses have already reached the "outside" while others are still engaged in interaction. But applying (**p2**) to the latter witnesses shows that the eventual contribution of these unfinished witnesses cannot exceed what we've already got. $\square$

Proposition 2.2 If σ and τ are strategies for $A \Rightarrow B$ and $B \Rightarrow C$ respectively then $\sigma\,;\tau$ is a strategy for $A \Rightarrow C$.

Remark If both strategies are deterministic, *i.e.* with values in $\{0, 1\}$, then $\sigma\,;\tau(s)$ is an integer which, by **(p1)** and **(p2)**, can only be 0 or 1, so there can be at most one witness of s. In the general case, a given play s can have (infinitely) many witnesses.

A category of arenas & strategies The above result tells us that we have a good notion of composition and it's routine to verify that this is associative. The usual ***copycat*** strategies are our identities:

$$\mathcal{T}_{\mathsf{id}_A} = \{s \in \mathcal{L}_{A_\ell \Rightarrow A_r} \mid \forall s' \sqsubseteq^{\mathsf{even}} s.\ s \restriction A_\ell = s \restriction A_r\}$$

where $\mathsf{id}_A(s) = 1$ for all $s \in \mathcal{T}_{\mathsf{id}_A}$.

In summary, we have a category where objects are arenas and an arrow from A to B is a strategy for $A \Rightarrow B$. We also have a subcategory of deterministic strategies since this class of strategies is closed under composition.

2.5 Examples of composing strategies

Consider composing coin with the "wait for true" strategy on $\mathbf{B} \Rightarrow \mathbf{C}$ defined as the semantics of

$$\mathbf{Y}(\lambda f.\ \lambda x.\ \texttt{if}\ x\ \texttt{then skip else}\ f(x)).$$

Some typical interactions are:

$$\begin{array}{ccccc}
\mathbf{1} & \Rightarrow & \mathbf{B} & \Rightarrow & \mathbf{C} \\
&&&& \mathsf{q} \\
&& \mathsf{q} && \\
&& \mathsf{tt} && \\
&&&& \mathsf{a} \\
\hline
&&&& \mathsf{q} \\
&& \mathsf{q} && \\
&& \mathsf{ff} && \\
&& \mathsf{q} && \\
&& \mathsf{tt} && \\
&&&& \mathsf{a} \\
\hline
&&&& \mathsf{q} \\
&& \mathsf{q} && \\
&& \mathsf{ff} && \\
&& \mathsf{q} && \\
&& \mathsf{ff} && \\
&& \mathsf{q} && \\
&& \mathsf{tt} && \\
&&&& \mathsf{a} \\
&& \vdots &&
\end{array}$$

The first witness of qa has probability $1/2$, the second has probability $1/4$, *etc.*, so the overall probability of qa is $1/2 + 1/4 + 1/8 + \cdots = 1$.

If we use a *biased* coin instead—where the more time that's passed, the less likely it is to return true—the probability of becoming "trapped" in the loop can be non-zero. For instance, set

$$\mathsf{coin}'(s\mathsf{qtt}) = \frac{\mu}{\lambda^{1+|s|/2}}$$

where $0 < \mu < 1-\lambda < 1$. The overall probability of qa is now $\sum_0^\infty \mu\lambda^n = \frac{\mu}{1-\lambda}$ which can be anything between 0 and 1. Of course, for coin' to satisfy (**p2**) we must have, for all k, $1 - \sum_0^k \mu\lambda^n \geq \mu\lambda^{k+1}$, but this is the case since, by assumption, $\mu < 1 - \lambda$.

We can easily make coin' into a total strategy so that the class of total strategies can't be closed under composition: pieces of chance can get lost in infinite chatters. This is the reason why (**p2**) is an *in*equality.

Another variation on this coin tossing example uses the "count for true" strategy on $\mathbf{B} \Rightarrow \mathbf{N}$ defined as the semantics of

$$\mathbf{Y}(\lambda f.\ \lambda x.\ \mathtt{if}\ x\ \mathtt{then}\ \mathtt{0}\ \mathtt{else}\ \mathtt{1}{+}f(x)).$$

In the composite with (the usual) coin, we get $\mathsf{q0}$ with probability $1/2$, $\mathsf{q1}$ with probability $1/4$, $\mathsf{q2}$ with probability $1/8$ and so on. In other words, the result is a geometric distribution on the natural numbers: we output n with probability $2^{-(n+1)}$.

Variations on a μ-scheme Consider the strategy mu_2 defined as the semantics of

$$\lambda f.\ [\mathbf{Y}(\lambda F.\lambda f.\lambda x.\ \mathtt{if0}\ f(x)\ \mathtt{skip}\ (F f\ (x{+}\mathtt{1})))]\ f\ \mathtt{2}$$

of type $(\mathtt{Nat} \to \mathtt{Nat}) \to \mathtt{Com}$. This term takes an input f and feeds it successive arguments, beginning with 2, until it gives an output 0, at which point it just terminates.

If we compose $\mathsf{polydie}$ with mu_2, the behaviour is as follows. Initially, a 2-sided die is rolled. If it turns up a 0, we're done; otherwise, a 3-sided die is rolled. If a 0 turns up this time, we're done; otherwise a 4-sided die is rolled, *etc.* The trace qa appears in the composite strategy with probability $1/2 + (1/2)\cdot(1/3) + (1/2)\cdot(2/3)\cdot(1/4) + (1/2)\cdot(2/3)\cdot(3/4)\cdot(1/5) + \cdots = \sum_{n\geq 1}(1/n)(1/n+1) = 1$.

An interesting variation on this is to stop once a 0 has been rolled, but only continue with the larger die if a 1 was rolled; otherwise, we just give up. This can be done by replacing the '$\mathtt{if0}$' subterm with

$$\mathtt{case}\ f(x)\ \mathtt{of}\ \mathtt{0} \mapsto \mathtt{skip} \mid \mathtt{1} \mapsto (F f\ (x{+}\mathtt{1})).$$

If we compose $\mathsf{polydie}$ with the strategy rr associated with this modified μ-scheme term, the probability of qa in the composite is $1/2 + (1/2)\cdot(1/3) + (1/2)\cdot(1/3)\cdot(1/4) + (1/2)\cdot(1/3)\cdot(1/4)\cdot(1/5) + \cdots = \sum_{n\geq 2} 1/n! = e - 2 \sim 0.718$. This shows that the class of strategies where all probabilities are rational isn't closed under composition.

2.6 SMCC structure

If σ and τ are strategies for $A \Rightarrow C$ and $B \Rightarrow D$ respectively then we define $\sigma \times \tau$ for $(A \times B) \Rightarrow (C \times D)$ in the following manner. Given $s \in \mathcal{T}_{\sigma\times\tau}$, we define its probability simply as the product of its probabilities in the two components:

$$(\sigma \times \tau)(s) = \sigma(s \upharpoonright A, C) \cdot \tau(s \upharpoonright B, D).$$

Note that there is no interaction between σ and τ: they each have their own die. This makes $\times$ into a bifunctor on our category. In fact, the evident copycat strategies provide us with a symmetric monoidal structure. In the usual way, a simple relabelling of moves gives rise to a natural isomorphism between σ on $(A \times B) \Rightarrow C$ and $\Lambda(\sigma)$ on $A \Rightarrow (B \Rightarrow C)$ so we have a symmetric monoidal closed category.

Finally, for any arena A, the ***diagonal*** strategy Δ_A on $A \Rightarrow (A \times A)$ equips A with a comonoid structure: for $s \in \mathcal{L}^{\mathsf{even}}_{A\Rightarrow(A_\ell \times A_r)}$, we have $s \in \mathcal{T}_{\Delta_A}$ (with probability 1) iff

$$\forall s' \sqsubseteq^{\mathsf{even}} s.\ s' \upharpoonright \ell \in \mathcal{T}_{\mathsf{id}_A} \wedge s' \upharpoonright r \in \mathcal{T}_{\mathsf{id}_A}$$

where $s' \upharpoonright \ell$ means the subsequence of s' consisting of all moves hereditarily justified by an initial occurrence in A_ℓ, and $s' \upharpoonright r$ is defined similarly.

2.7 Ordering strategies

Given strategies σ and τ for some arena A, we set $\sigma \leqslant \tau$ iff, for all $s \in \mathcal{L}_A$, we have $\sigma(s) \leq \tau(s)$. This is clearly a partial order. It's straightforward to show that composition is monotone with respect to this.

CPO-enrichment In fact, we can show that each homset of our category is a CPO and that composition is continuous, *i.e.* our category "is" CPO-enriched.

Given a $\leqslant$-directed set of strategies Δ, we define $\bigsqcup \Delta$ as follows.

$$(\bigsqcup \Delta)(s) = \bigvee \{\sigma(s) \mid \sigma \in \Delta\}$$

It can be shown that this defines a good strategy and that $\bigsqcup \Delta$ is indeed the least upper bound of Δ. A routine but rather involved argument establishes that composition is continuous.

A countable basis For any arena A, the collection of strategies with *finite trace set* and where all probabilities are *rational* forms a basis for the CPO of all strategies, *i.e.* for any σ on A, the set Δ_σ of such strategies approximating σ is directed and has σ as its least upper bound.

This basis is always countable so that, for any arena A, the set of all strategies for A forms an ω-continuous CPO. This contrasts with *all* previous (HO-style) games models, including [12], where strategies form ω-*algebraic* CPOs—the ***compact*** strategies being those with a finite number of traces.

2.8 Deterministic factorization

The technique of *factorization* has become standard in game semantics. The underlying idea is similar to that of [12] where a nondeterministic strategy is reduced to a deterministic strategy that has access to a nondeterministic "oracle". In our case, the "oracle" strategy is a die. It suffices to use a 2-sided die but, for convenience, we use the $\mathsf{polydie}$ defined above.

We first present a simplified factorization that works for basis strategies. Afterwards, we sketch a more general factorization that allows us to simulate *any* probabilistic strategy with a deterministic strategy and the $\mathsf{polydie}$.

Basis factorization Let σ be a basis strategy for arena A and consider $sa \in \mathsf{dom}(\sigma)$. Since $\mathcal{T}_\sigma$ is finite, we know that $\mathsf{rng}_\sigma(sa)$ must be finite; we'll write it as $\mathsf{rng}_\sigma(sa) = \{sab_1, \ldots, sab_n\}$. Since all probabilities in σ are rational, the *local* probability $\sigma(sab_i/sa)$, for each $sab_i \in \mathsf{rng}_\sigma(sa)$, is too.

Theorem 2.3 (deterministic factorization) If σ is a basis strategy for the arena A then there exists a *deterministic* strategy $\mathsf{Det}(\sigma)$ for $(\mathbf{N} \Rightarrow \mathbf{N}) \Rightarrow A$ such that $\mathsf{polydie}\,;\mathsf{Det}(\sigma) = \sigma$ and $\mathsf{Det}(\sigma)$ is compact.

The above remarks imply that we can find a common denominator d for the local probabilities $\sigma(sab_i/sa)$. Let $\{p_1, \ldots, p_n\}$ be the numerators such that $p_i/d = \sigma(sab_i/sa)$. The strategy $\mathsf{Det}(\sigma)$ proceeds in the following fashion...

$$
\begin{array}{ccccc}
(\mathbf{N} & \Rightarrow & \mathbf{N}) & \Rightarrow & A \\
 & & & \vdots & \\
 & & & & a \\
 & & \mathsf{q} & & \\
\mathsf{q} & & & & \\
\mathsf{d} & & & & \\
 & & \mathsf{i} & & \\
 & & & & b_j
\end{array}
$$

where j is 1 if $0 \le i < p_1$, j is 2 if $p_1 \le i < p_1 + p_2$, j is 3 if $p_1 + p_2 \le i < p_1 + p_2 + p_3$ and so on. In other words, we "slice up" the interval from 0 to $d - 1$ according to the numerators p_i.

This clearly defines a deterministic strategy and it's easy to check that supplying the $\mathsf{polydie}$ as input achieves a correct simulation of the original strategy so that $\mathsf{polydie}\,;\mathsf{Det}(\sigma) = \sigma$. It's also obvious that, since σ has a finite trace set, $\mathsf{Det}(\sigma)$ must have too. $\Box$ Conversely, a small calculation with the composition formula shows that any strategy that can be factorized as $\mathsf{polydie};\sigma$ where σ is deterministic and compact must belong to the basis.

General factorization To extend this idea to the general case, consider first a strategy σ all of whose probabilities are rational (but not necessarily having a finite trace set). If $sa \in \mathsf{dom}(\sigma)$ with range $sab_1, sab_2, \ldots$, we again consider the local probabilities. If $\sigma(sab_1/sa) = n_1/d_1$, we roll a d_1-sided die. If the result is less than n_1, we play b_1; otherwise, we *renormalize* the remaining probabilities by dividing through by $1 - (n_1/d_1)$ and then repeat this process for $\sigma(sab_2/sa)$...

e.g. if the local probabilities of b_1 and b_2 are 2/3 and 1/6 respectively, the following play produces b_2.

$$
\begin{array}{ccccc}
(\mathbf{N} & \Rightarrow & \mathbf{N}) & \Rightarrow & A \\
 & & & \vdots & \\
 & & & & a \\
 & & \mathsf{q} & & \\
\mathsf{q} & & & & \\
\mathsf{3} & & & & \\
 & & \mathsf{2} & & \\
 & & \mathsf{q} & & \\
\mathsf{q} & & & & \\
\mathsf{2} & & & & \\
 & & \mathsf{0} & & \\
 & & & & b_2
\end{array}
$$

We roll a 2-sided die the second time since the renormalized probability of b_2 is $(1/6)/(1/3) = 1/2$.

Finally, if a play in σ has irrational probability p, we can always find a sequence of rationals whose series of partial sums converges to p. So, for each point sab_i in the range, we either have a single rational or a sequence of rationals. Using a standard trick (such as the one used to enumerate $\mathbf{N} \times \mathbf{N}$) we can reduce this to the previous case. The essential point of this is that a play with irrational probability is *not* factorized in "one go". In fact, it will have an infinite number of witnesses: the desired probability gradually builds up as we progress further and further down its sequence of rationals.

3 A fully abstract model of PA

3.1 Probabilistic Algol

Our starting point is the language IA as defined in [3] for example. All we need to do is add in some probabilistic primitive. But should we add a coin or should we rather choose the polydie? Well, the polydie might be more handy to program with—indeed, the programming language CAML has a "probabilistic function" `Random.int` which exactly implements the polydie—but, for the contextual equivalence, it makes no difference whatsoever. In fact, the polydie can even be programmed inside Algol using only a coin. Informally it goes this way: given an n, choose a k such that $2^k \geq n$, then roll for a first round the 2-sided die k times. If the answer, which is a binary string of length k, codes for some integer in $\{0, \ldots, n-1\}$, return it; if not, go for another round, and so on, until you succeed.

So, for the sake of simplicity, we just add a single term `coin` of type `Bool`. We call this extended language PA for *probabilistic Algol*.

To complete the definition of the language, we need to give its operational interpretation. We do this in the usual "big step" style except that derivations must be decorated with probabilities.

We have the obvious rules for the `coin`:

$$\frac{}{\texttt{coin} \Downarrow^{1/2} \texttt{tt}} \qquad \frac{}{\texttt{coin} \Downarrow^{1/2} \texttt{ff}}$$

The other rules follow naturally; we just give a few examples. (For the sake of readability, we're eliding the use of *stores*; adding them poses no problem.)

$$\frac{M \Downarrow^p \lambda x.\ M' \quad M'[N/x] \Downarrow^q V}{MN \Downarrow^{pq} V}$$

$$\frac{M \Downarrow^p \texttt{n}}{\texttt{succ}\ M \Downarrow^p \texttt{n}+\texttt{1}} \qquad \frac{M \Downarrow^p \texttt{tt} \quad N \Downarrow^q V}{\texttt{if}\ M\ N\ L \Downarrow^{pq} V}$$

Finally, let M and N be closed terms of the same type. They are ***contextually equivalent*** iff, for all contexts $C[-]$ of ground type, we have $C[M] \Downarrow^p V$ if, and only if, $C[N] \Downarrow^p$. We denote this by $M \simeq N$.

As an example, let M_n be the "nth Church numeral" '$\lambda h.h^n(\texttt{tt})$' and N be '$\lambda x.$ `if coin ff` x' of types $(\texttt{Bool} \to \texttt{Bool}) \to \texttt{Bool}$ and $\texttt{Bool} \to \texttt{Bool}$. We observe that $M_n N \Downarrow^{2^{-n}} \texttt{tt}$ so that the context $(-)N$ allows us to distinguish between all the M_ns.

This demonstrates the discriminating power, even in the purely functional world, of probabilistic contexts, for it is easily seen that M_n and M_{n+2} are contextually equivalent in PCF.

3.2 The model

We now begin the process of building a model of PA by introducing the class of single-threaded strategies.

Recall that the current thread of $sa \in \mathcal{L}_A^{\mathsf{odd}}$ is written as $\lceil sa \rceil$. If we have some $tab \in \mathcal{L}_A^{\mathsf{even}}$ such that $\lceil sa \rceil = \lceil ta \rceil$ and the justifier of b occurs in $\lceil ta \rceil$, we denote by $\mathsf{Match}(tab, sa)$ the play sab where b is justified in the same way as in tab.

We say that a strategy σ is ***single-threaded*** iff

- for all $sab \in \mathcal{T}_\sigma$, the justifier of b occurs in $\lceil sa \rceil$.
- if $sab \in \mathcal{T}_\sigma$, $ta \in \mathsf{dom}(\sigma)$ and $\lceil sa \rceil = \lceil ta \rceil$ then:

$$\sigma(sab/sa) = \sigma(\mathsf{Match}(sab, ta)/ta).$$

Note that this condition is expressed in terms of the *local* probabilities of σ.

Proposition 3.1 A strategy σ for $A \Rightarrow B$ is single-threaded if, and only if, it's a comonoid homomorphism, *i.e.* $\sigma\,;\Delta_B = \Delta_A\,;(\sigma \times \sigma)$.

For general reasons [11], the single-threaded strategies form a Cartesian closed category $\mathcal{C}$. Furthermore, the CPO-enriched structure all restricts down to $\mathcal{C}$, *i.e.* we have a CPO-enriched CCC.

Visibility & bracketing The remaining constraints required to cut the model down to match well with Algol both rely on the following definition of the ***view***, written $\ulcorner s \urcorner$, of a non-empty legal play s.

$$\begin{aligned} \ulcorner sm \urcorner &= m, & &\text{if } m \text{ is initial;} \\ \ulcorner sntm \urcorner &= \ulcorner s \urcorner nm, & &\text{if } m \text{ is an } \mathsf{O}\text{-move and} \\ & & &n \text{ justifies } m; \\ \ulcorner sm \urcorner &= \ulcorner s \urcorner m, & &\text{if } m \text{ is a } \mathsf{P}\text{-move.} \end{aligned}$$

A strategy σ satisfies the ***visibility*** condition iff, for all $sab \in \mathcal{T}_\sigma$, the justifier of b occurs in $\ulcorner sa \urcorner$; and σ satisfies the ***bracketing*** condition iff, for all $sab \in \mathcal{T}_\sigma$ where b is an Answer, b is justified by the pending Question of $\ulcorner sa \urcorner$, *i.e.* the most recently asked but, as yet, unanswered Question in $\ulcorner sa \urcorner$. Both these constraints are preserved by composition.

Our interpretation of PA is in the sub-CCC of $\mathcal{C}$, denoted by $\mathcal{C}_{vb}$, of (single-threaded) strategies subject to both these constraints. The interpretation of IA, *i.e.* the deterministic fragment of the language, follows that of Abramsky & McCusker [3]; our probabilistic primitive `coin` is interpreted by the strategy coin.

We have the following *soundness* result for our model of PA in $\mathcal{C}_{vb}$.

Theorem 3.2 If M and N are closed terms of PA of type T such that $[\![M]\!] = [\![N]\!]$ then $M \simeq N$.

3.3 Full abstraction

With the soundness result in place, the path to full abstraction is relatively standard. First of all, we prove a *definability* result for basis strategies in the model.

Theorem 3.3 If T is some type of PA and σ is a basis strategy on $[\![T]\!]$ then $\sigma = [\![M]\!]$ for a term M of PA.

By the factorization theorem, $\sigma = \mathsf{polydie} \,;\, \mathsf{Det}(\sigma)$ where $\mathsf{Det}(\sigma)$ is a compact deterministic strategy. By the definability theorem of Abramsky & McCusker [3], this strategy is definable in IA. But, since $\mathsf{polydie}$ is definable in PA, we're done. $\square$

We define an equivalence relation on each homset of $\mathcal{C}_{vb}$ as follows. Given $f, g : A \Rightarrow B$, write $f \sim g$ iff for all "test morphisms" $\alpha : [A \Rightarrow B] \Rightarrow \mathbf{C}$,

$$`f\text{'} \,;\, \alpha = `g\text{'} \,;\, \alpha$$

where $`f\text{'} : \mathbf{1} \Rightarrow [A \Rightarrow B]$ is the *name* of f, defined by currying. The so-called ***intrinsic quotient*** category $\mathcal{E}_{vb}$ obtained by quotienting $\mathcal{C}_{vb}$ with this relation is also a CCC and, moreover, the above soundness result survives the quotient. The converse to soundness also holds there.

Theorem 3.4 (full abstraction) If M and N are closed terms of type T then $\mathcal{E}[\![M]\!] = \mathcal{E}[\![N]\!]$ if, and only if, $M \simeq N$.

We only need to prove the right-to-left direction. Suppose that $\mathcal{E}[\![M]\!] \neq \mathcal{E}[\![N]\!]$ so that, for some test α, $[\![M]\!] \,;\, \alpha \neq [\![N]\!] \,;\, \alpha$. Since homsets in $\mathcal{C}_{vb}$ are continuous CPOs, we can WLOG assume α to be a basis element. Hence, by definability, $\alpha = [\![x : T \vdash P : \mathtt{Com}]\!]$ for some context P and it follows that this context distinguishes M and N. $\square$

References

[1] S. Abramsky and P.-A. Mellies. Concurrent games and full completeness. In *Proceedings, Fourteenth Annual IEEE Symposium on Logic in Computer Science*. IEEE Computer Society Press, 1999.

[2] S. Abramsky, K. Honda, and G. McCusker. A fully abstract game semantics for general references. In *Proceedings, Thirteenth Annual IEEE Symposium on Logic in Computer Science*, pages 334–344, 1998.

[3] S. Abramsky and G. McCusker. Linearity, sharing and state: a fully abstract game semantics for Idealized Algol with active expressions. In P. W. O'Hearn and R. D. Tennent, editors, *Algol-like languages*. Birkhaüser, 1997.

[4] S. Abramsky and G. McCusker. Call-by-value games. In M. Nielsen and W. Thomas, editors, *Computer Science Logic: 11th International Workshop Proceedings*, Lecture Notes in Computer Science. Springer-Verlag, 1998.

[5] S. Abramsky and G. McCusker. Full abstraction for Idealized Algol with passive expressions. To appear in Theoretical Computer Science, 1998.

[6] P. Baillot, V. Danos, T. Ehrhard, and L. Regnier. Timeless games. In W. T. M. Nielsen, editor, *Computer Science Logic: 11th International Workshop Proceedings*, volume 1414 of *Lecture Notes in Computer Science*. EACSL, Springer Verlag, 1997.

[7] A. Edalat. Domains for computation in mathematics, physics and exact real arithmetic. *Bulletin of Symbolic Logic*, 3(4), 1997.

[8] A. Edalat, R. Blute, J. Desharnais, and P. Panangaden. Bisimulation for labeled Markov processes. In *Proceedings, twelfth Annual IEEE Symposium on Logic In Computer Science*, 1997.

[9] A. Edalat, J. Desharnais, and P. Panangaden. A logical characterization of bisimulation for labeled Markov processes. In *Proceedings, thirteenth Annual IEEE Symposium on Logic In Computer Science*, 1998.

[10] M. Escardo. PCF extended with real numbers. *Theoretical Computer Science*, 162(1):79–115, 1996.

[11] P. J. Freyd, P. W. Hearn, A. J. Power, M. Takeyama, and R. D. Tennent. Bireflectivity. In *Mathematical Foundations of Programming Semantics, Eleventh Annual Conference, Tulane University, New Orleans, LA, March 29–April 1, 1995*. Elsevier, 1995.

[12] R. S. Harmer and G. A. McCusker. A fully abstract game semantics for finite nondeterminism. In *Proceedings, Fourteenth Annual IEEE Symposium on Logic in Computer Science*. IEEE Computer Society Press, 1999.

[13] J. M. E. Hyland and C.-H. L. Ong. On full abstraction for PCF: I, II and III. To appear in Information and Computation, 200?, 1994.

[14] C. Jones and G. D. Plotkin. A probabilistic powerdomain of evaluations. In *Proceedings, fourth Annual IEEE Symposium on Logic In Computer Science*, 1989.

[15] J. Laird. Full abstraction for functional languages with control. In *Proceedings, Twelfth Annual IEEE Symposium on Logic in Computer Science*. IEEE Computer Society Press, 1997.

[16] P. Malacaria and C. Hankin. Generalised flowcharts and games. In *Proceedings, 25th International Colloquium on Automata, Languages and Programming*, volume 1443 of *Lecture Notes in Computer Science*. Springer-Verlag, 1998.

[17] P. Malacaria and C. Hankin. A new approach to control flow analysis. In *Proceedings, CC'98*, volume 1383 of *Lecture Notes in Computer Science*. Springer-Verlag, 1998.

[18] P. Malacaria and C. Hankin. Non-deterministic games and program analysis: An application to security. In *Proceedings, Fourteenth Annual IEEE Symposium on Logic in Computer Science*. IEEE Computer Society Press, 1999.

[19] G. McCusker. Games and full abstraction for FPC. In *Proceedings, Eleventh Annual IEEE Symposium on Logic in Computer Science*, pages 174–183. IEEE Computer Society Press, 1996.

[20] H. Nickau. Hereditarily sequential functionals. In *Proceedings, Logical Foundations of Computer Science*, volume 813 of *Lecture Notes in Computer Science*. Springer-Verlag, 1994.

[21] P. J. Potts, A. Edalat, and M. H. Escardó. Semantics of exact real number arithmetic. In *Proceedings of the 12th Annual IEEE Symposium on Logic In Computer Science*, 1997.

[22] N. Saheb-Djahromi. Probabilistic LCF. In *Proceedings, Mathematical Foundations of Computer Science*, volume 64 of *Lecture Notes in Computer Science*. Springer-Verlag, 1978.

[23] N. Saheb-Djahromi. CPOs of measures for nondeterminism. *Theoretical Computer Science*, 12(1), 1980.