Forgetting causality in the concurrent game semantics of probabilistic PCF

Simon Castellan[1], Pierre Clairambault[1], Hugo Paquet[2], and Glynn Winskel[2]

[1]Univ Lyon, CNRS, ENS de Lyon, UCB Lyon 1, LIP, Lyon, France
[2]Computer Laboratory, University of Cambridge, Cambridge, U.K.

***Abstract*—We enrich *thin concurrent games with symmetry*, recently introduced by Castellan *et al*, with probabilities, and build on top of it a cartesian closed category with an interpretation of Probabilistic PCF (PPCF). Exploiting that the strategies obtained from PPCF terms have a deadlock-free interaction, we deduce that there is a functor preserving the interpretation from our games to the *probabilistic relational model* recently proved fully abstract by Ehrhard *et al*. It follows that our model is intensionally fully abstract, without the need for a probabilistic notion of innocence. This holds both for a *sequential* and a *parallel* interpretation of PPCF in the style of Castellan *et al*.**

I. INTRODUCTION

One may separate denotational models of programming languages into two distinct families, *static* and *dynamic*. *Static* models represent a program in terms of the *states* it may reach. The states can take different forms and static models are many and various. Organised as a category, a model may be *well-pointed* (*i.e. extensional*, as in Scott domains [25]) or not (as in the relational model of linear logic [16]). States may come annotated with additional quantitative information (as in the probabilistic relational model [17]), and a model may even allow multiple witnesses for the same states (as in *spans* or *profunctors*). All those cases have in common that composition is an elaboration of relational composition; it collects information about the states reached, with no concern for the causal history of those states.

On the other hand, in a *dynamic* model one remembers some temporal or causal information from the execution of the source programs. Most dynamic semantics have been given within *game semantics* [18], [3], but one should not forget an early precursor, *sequential algorithms* [8]. The additional intensional information of game semantics makes it well-suited to accommodate *non-commutative effects*, dependent on the evaluation order; one may cite various games models of stateful languages [4], [2]. This intensionality also makes such models more modular, as powerfully illustrated by the so-called *Abramsky cube* [1], describing programming language features in terms of abstract conditions on strategies. Composition in such models is dramatically different from that in static models: for two strategies to synchronise they must not only agree on the state to be reached, but also have compatible requirements as to *how* it is going to be reached – otherwise the composition will deadlock.

Accordingly, *time-forgetting* operations from games to relations are naturally lax functorial [7], [31]; but become functorial for well-behaved programs [23], [9] – expressed by Melliès as, *innocent strategies are positional* [23].

In this paper, we are interested in both types of models for PPCF, *i.e.* PCF [25] extended with probabilistic choice. On the dynamic side, Danos and Harmer [15] gave a games model for *Probabilistic Idealized Algol*, which includes PPCF. On the static side, Danos and Ehrhard gave a model for PPCF in *probabilistic coherence spaces*, based on the probabilistic relational model [14]. Strikingly their model is fully abstract for PPCF [17], *i.e.* two programs have the same denotation if and only if they cannot be distinguished.

But again, the time-forgetting operation from Danos and Harmer's model to the probabilistic relational model is only lax functorial. This is inevitable since the games model supports references, which are time-sensitive. But how may we adapt it so that it still supports the interpretation of PPCF and yet has a functorial time-forgetting operation, *i.e.* how do we construct a deadlock-free games model of PPCF? To achieve an absence of deadlocks, we move to a games model that takes causality more seriously.

Our first contribution is the construction of a causal games model of PPCF. It draws on two recent developments in *concurrent game semantics*: the *thin concurrent games with symmetry* of [12], [13] used to build a parallel model for PCF [12], and the *probabilistic concurrent strategies* of [30]. Our second contribution is a collapse operation to the probabilistic relational model. We prove a deadlock-free property for *visible* [12] concurrent strategies, which entails that the collapse is functorial and preserves the interpretation of PPCF. From that it follows that the model is *intensionally fully abstract*: it is adequate, and strategies have the same distinguishing power as terms – just as in the original HO and AJM games models with respect to PCF. Interestingly, the result does not require any notion of probabilistic innocence or definability result. Moreover, it holds both for a *sequential* interpretation of PPCF (compatible with its interpretation in Danos and Harmer's model) and a *parallel* interpretation in the style of [12].

Related work: Our probabilistic game semantics is related to Tsukada and Ong's sheaf-based non-deterministic and probabilistic innocence [27], although precise connections have not been investigated. That innocent strategies have a deadlock-free composition is implicit in Melliès' work on game semantics for linear logic [5], [22], and exploited in Boudes' work on relating games with the relational model –

our deadlock-free property for visible strategies generalises that to a non-sequential and non-innocent setting.

Outline: In Section II we introduce PPCF, its relational semantics, and describe statically the probabilistic event structures used to represent it. In Section III we introduce *probabilistic thin concurrent games*, the setting on which the compositional interpretation of PPCF relies. Finally, in Section IV, we give our collapse operation to the probabilistic relational model and deduce intensional full abstraction.

II. Semantics for Probabilistic Programs

A. Probabilistic PCF

We present the language PPCF, the extension of Plotkin's PCF [25] with a probabilistic primitive **coin** : $\mathbf{Bool}$. Its **types** are those obtained from the basic types $\mathbf{Bool}$ and $\mathbf{Nat}$, and the arrow $\Rightarrow$. Its **terms** are the following

$$\begin{array}{rcl} M, N & ::= & \lambda x.\, M \mid M\, N \mid x \mid \mathbf{tt} \mid \mathbf{ff} \mid \mathbf{if}\, M\, N_1\, N_2 \mid Y \\ & & n \mid \mathbf{pred}\ M \mid \mathbf{succ}\ M \mid \mathbf{iszero}\ M \mid \mathbf{coin} \end{array}$$

The typing rules are standard and omitted – we assume that in $\mathbf{if}\, M\, N_1\, N_2$, N_1 and N_2 have ground type ($\mathbf{Bool}$ or $\mathbf{Nat}$), a general **if** can be defined as syntactic sugar.

The usual call-by-name operational semantics for PCF generalises to a probabilistic reduction relation $\xrightarrow{p}$, for $p \in [0,1]$. All rules are straightforward, with the primitive **coin** representing a fair coin: $\mathbf{coin} \to^{\frac{1}{2}} b$ for all $b \in \{\mathbf{tt}, \mathbf{ff}\}$. Because reduction is non-deterministic, there can be countably many **reduction paths** from M to N, *i.e.* sequences of the form $M = M_0 \xrightarrow{p_1} \dots \xrightarrow{p_n} M_n = N$. Given such a path π, its **weight** $w(\pi)$ is $\prod_{1\leq i \leq n} p_i$, and we define the coefficient $\mathrm{Pr}(M \to N)$ as $\sum \{ w(\pi) \mid \pi \text{ is a path from } M \text{ to } N\}$.

Definition 1. *Let M and N be PPCF terms such that $\Gamma \vdash M : A$ and $\Gamma \vdash N : A$. We write $M \lesssim_{ctx} N$ if for every context $C[\cdot]$ such that $\vdash C[P] : \mathbf{Bool}$ for every $\Gamma \vdash P : A$,*

$$\mathrm{Pr}(C[M] \to b) \leq \mathrm{Pr}(C[N] \to b)$$

for all $b \in \{\mathbf{tt}, \mathbf{ff}\}$. The equivalence induced by this preorder, ***contextual equivalence****, is denoted $\simeq_{ctx}$.*

B. The weighted relational model

In [17], Ehrhard *et al* proved that *probabilistic coherence spaces* ($\mathbf{PCoh}$) are **fully abstract** for PPCF – in other words, two PPCF terms are contextually equivalent iff they have the same denotation in $\mathbf{PCoh}$. In fact, $\mathbf{PCoh}$ is cut down (via *biorthogonality*) from a more liberal model $\mathbf{PRel}$, the *probabilistic relational model*, on which we will now focus.

1) The relational model of PCF: Ignoring probability for now, the *relational* model of PCF records the *input-output behaviour* of a term, along with the *multiplicity* of resources.

Write $\mathbb{B} = \{\mathbf{tt}, \mathbf{ff}\}$ and $\mathcal{M}_{\mathrm{f}}(X)$ for the set of **finite multisets** of elements of a set X. Objects of $\mathcal{M}_{\mathrm{f}}(X)$ are written with square brackets with elements annotated with their multiplicity; *e.g.* we have $[\mathbf{tt}^2, \mathbf{ff}] \in \mathcal{M}_{\mathrm{f}}(\mathbb{B})$, where $\mathbf{tt}$ has multiplicity 2 and $\mathbf{ff}$ has multiplicity 1. Using this notation, the term $b_1 : \mathbf{Bool}, b_2 : \mathbf{Bool} \vdash \mathbf{if}\, b_1\, b_1\, b_2 : \mathbf{Bool}$ will be represented as the subset of $\mathcal{M}_{\mathrm{f}}(\mathbb{B}) \times \mathcal{M}_{\mathrm{f}}(\mathbb{B}) \times \mathbb{B}$ containing:

$$\begin{array}{ccccc} \mathcal{M}_{\mathrm{f}}(\mathbb{B}) & \times & \mathcal{M}_{\mathrm{f}}(\mathbb{B}) & \times & \mathbb{B} \\ ([\mathbf{tt}^2], & & [\,], & & \mathbf{tt}) \\ ([\mathbf{tt}, \mathbf{ff}], & & [\,], & & \mathbf{ff}) \\ ([\mathbf{ff}], & & [\mathbf{tt}], & & \mathbf{tt}) \\ ([\mathbf{ff}], & & [\mathbf{ff}], & & \mathbf{ff}) \end{array}$$

The model is *non-uniform*: it shows how the term behaves if its argument ever *changes its mind*.

The interpretation of PCF in the relational model follows the usual methodology of denotational semantics, and in particular the interpretation of the simply-typed λ-calculus in a cartesian closed category, see *e.g.* [20] for an introduction. To construct the target cartesian closed category, we start with one of the simplest models of linear logic: the category $\mathbf{Rel}$ of sets and relations. In $\mathbf{Rel}$ the linear logic connectives are interpreted as follows: given X and Y, $X \otimes Y = X \multimap Y = X \times Y$, $X \& Y = X + Y$ (the tagged disjoint union) and $!X = \mathcal{M}_{\mathrm{f}}(X)$. The cartesian closed category $\mathbf{Rel}_!$ is then the Kleisli category for the comonad $!$, see *e.g.* [24]. We delay the details of the interpretation of PCF in $\mathbf{Rel}_!$, which we will cover in the presence of probabilities.

2) The weighted relational model: Because the model is non-uniform, it supports non-deterministic primitives. The idea behind the probabilistic relational model is to enrich this non-uniform model with quantitative information: each element comes with a *weight*, as shown for instance in the interpretation of $M_+ = b : \mathbf{Bool} \vdash \mathbf{if}\, b\, (\mathbf{if}\, \mathbf{coin}\, b\, \bot)\, (\mathbf{if}\, b\, \mathbf{ff}\, \mathbf{tt}) : \mathbf{Bool}$, where $\bot$ is a diverging term, *e.g.* $Y\, (\lambda x.\, x)$:

$$\begin{array}{ccc} \mathcal{M}_{\mathrm{f}}(\mathbb{B}) & \times & \mathbb{B} \\ ([\mathbf{tt}^2], & & \mathbf{tt})^{\frac{1}{2}} \\ ([\mathbf{tt}, \mathbf{ff}], & & \mathbf{ff})^{\frac{3}{2}} \\ ([\mathbf{ff}^2], & & \mathbf{tt})^{1} \end{array}$$

The weights can be greater than 1, because a multiset may correspond to several *execution traces*. In the example above the pair $([\mathbf{tt}, \mathbf{ff}], \mathbf{ff})$ has weight $\frac{3}{2} = \frac{1}{2} + 1$, summing over the different orders in which b can take its values from $[\mathbf{tt}, \mathbf{ff}]$.

The *pure* relational interpretation from before was based on the category $\mathbf{Rel}$ with objects sets and morphisms from X to Y relations $\varphi \subseteq X \times Y$, *i.e.* "matrices" $(\varphi_{x,y})_{x,y \in X \times Y} \in \{0,1\}^{(X\times Y)}$. Accordingly, the composition of relations can be regarded as matrix multiplication:

$$(\psi \circ \varphi)_{x,z} = \bigvee_{y \in Y} (\varphi_{x,y} \wedge \psi_{y,z})$$

So naively, one might try to construct a probabilistic variant of $\mathbf{Rel}$ by simply replacing the boolean semiring $(\{0,1\}, \vee, \wedge)$ above by the semiring $(\mathbb{R}_+, +, \times)$ where $\mathbb{R}_+$ denotes non-negative real numbers; except that the composition formula

$$(\psi \circ \varphi)_{x,z} = \sum_{y \in Y} (\varphi_{x,y} \times \psi_{y,z})$$

for $\varphi \in \mathbb{R}_+^{X\times Y}$, $\psi \in \mathbb{R}_+^{Y \times Z}$, will not in general satisfy that $(\psi \circ \varphi)_{x,z} \in \mathbb{R}_+$ as there is no reason for it to converge.

So instead, we will consider the interpretation of PPCF in a quantitative generalization of $\mathbf{Rel}$ weighted by elements of

$\overline{\mathbb{R}}_+ = \mathbb{R}_+ \uplus \{\infty\}$, the infinity added to ensure convergence of the (potentially) infinite sum above. There is a category $\mathbf{PRel}$ with sets as objects, and as morphisms from X to Y the (potentially) infinite matrices $\varphi \in \overline{\mathbb{R}}_+^{X \times Y}$, composed as above. The *identity* on X is the diagonal matrix $(\delta_{x_1,x_2})_{x_1,x_2 \in X}$ where δ_{x_1,x_2} is 1 whenever $x_1 = x_2$, and 0 otherwise.

Now, just like $\mathbf{Rel}$, $\mathbf{PRel}$ supports the structure of a model of linear logic with the constructions on objects the same as in $\mathbf{Rel}$ and analogous constructions on morphisms. We proceed to define the interpretation of PPCF in $\mathbf{PRel}_!$. As for $\mathbf{Rel}$ the interpretation of the λ-calculus combinators follows from the cartesian closed structure of the Kleisli category $\mathbf{PRel}_!$, which we do not detail further [20]. The interpretation of Y is also obtained in a standard way as a least upper bound of finite approximations, using that homsets of $\mathbf{PRel}$ are dcpos when ordered pointwise. We now focus on the interpretation of ground types and associated combinators.

The types $\mathbf{Bool}$ and $\mathbf{Nat}$ are interpreted by the sets $[\![\mathbf{Bool}]\!] = \mathbb{B}$ and $[\![\mathbf{Nat}]\!] = \mathbb{N}$, respectively. For $n \in \mathbb{N}$, the constant $\underline{n}$ has semantics given by $([\![\underline{n}]\!])_k = \delta_{k,n}$ for $k \in \mathbb{N}$. The boolean constants $\mathbf{tt}$ and $\mathbf{ff}$ are interpreted in the same way. The semantics of **succ** and **pred** are defined by

$$
\begin{array}{rcccccl}
[\![\textbf{succ}]\!] & : & \mathcal{M}_{\mathrm{f}}(\mathbb{N}) & \times & \mathbb{N} & \to & \overline{\mathbb{R}}_+ \\
& & ([n] & , & n+1) & \mapsto & 1 \\
& & (\_ & , & \_) & \mapsto & 0 \\
[\![\textbf{pred}]\!] & : & \mathcal{M}_{\mathrm{f}}(\mathbb{N}) & \times & \mathbb{N} & \to & \overline{\mathbb{R}}_+ \\
& & ([n+1] & , & n) & \mapsto & 1 \\
& & ([0] & , & 0) & \mapsto & 1 \\
& & (\_ & , & \_) & \mapsto & 0
\end{array}
$$

The morphism $[\![\textbf{iszero}]\!] \in \mathbf{PRel}_!(\mathbb{N}, \mathbb{B})$ is defined similarly. Finally given terms $M : \mathbf{Bool}, N : \mathbb{X}, P : \mathbb{X}$ (where $\mathbb{X}$ denotes any ground type, *i.e.* $\mathbf{Bool}$ or $\mathbf{Nat}$), the term $\textbf{if}\, M\ N\ P$ has semantics $\langle [\![M]\!], \langle [\![N]\!], [\![P]\!] \rangle \rangle \circ \mathit{if}$, where $\mathit{if} \in \mathbf{PRel}_!(\mathbb{B} \& ([\![\mathbb{X}]\!] \& [\![\mathbb{X}]\!]), [\![\mathbb{X}]\!]) \cong \mathbf{PRel}(!\mathbb{B} \otimes ![\![\mathbb{X}]\!] \otimes ![\![\mathbb{X}]\!], [\![\mathbb{X}]\!])$ is defined by

$$
\begin{array}{rcccccccccl}
\mathit{if} & : & \mathcal{M}_{\mathrm{f}}(\mathbb{B}) & \times & \mathcal{M}_{\mathrm{f}}([\![\mathbb{X}]\!]) & \times & \mathcal{M}_{\mathrm{f}}([\![\mathbb{X}]\!]) & \times & [\![\mathbb{X}]\!] & \to & \overline{\mathbb{R}}_+ \\
& & ([\mathbf{tt}] & , & [x] & , & [] & , & x) & \mapsto & 1 \\
& & ([\mathbf{ff}] & , & [] & , & [x] & , & x) & \mapsto & 1 \\
& & (\_ & , & \_ & , & \_ & , & \_) & \mapsto & 0
\end{array}
$$

Finally, the probabilistic primitive $\mathbf{coin}$ is interpreted as expected as having $[\![\mathbf{coin}]\!]_{\mathbf{tt}} = \frac{1}{2}$ and $[\![\mathbf{coin}]\!]_{\mathbf{ff}} = \frac{1}{2}$, completing the interpretation of PPCF. One may however question how satisfactory this model is – it is quite obviously very far from full completeness, as witnessed by the presence of infinite weights. And indeed, the authors of [17] do not stop with $\mathbf{PRel}$. Instead they cut it down by a biorthogonality construction to obtain another weighted model of linear logic, $\mathbf{PCoh}$. In $\mathbf{PCoh}$ weights remain finite, and the interpretation of $M : \mathbb{X}$ yields a sub-probability distribution on $[\![\mathbb{X}]\!]$. In fact, the main result of [17] is that $\mathbf{PCoh}$ is *fully abstract*, *i.e.* for any M, N we have that $M \simeq_{\mathrm{ctx}} N$ iff $[\![M]\!]_{\mathbf{PCoh}} = [\![N]\!]_{\mathbf{PCoh}}$.

But this has the interesting immediate corollary that despite its drawbacks, $\mathbf{PRel}$ is itself already fully abstract! Indeed there is an obvious faithful forgetful functor $\mathbf{PCoh} \hookrightarrow \mathbf{PRel}$ preserving all the structure on the nose – in fact a term M has exactly the same interpretation in $\mathbf{PRel}$ and $\mathbf{PCoh}$, the only difference being that the interpretation is more informative as it carries correctness information *w.r.t.* biorthogonality.

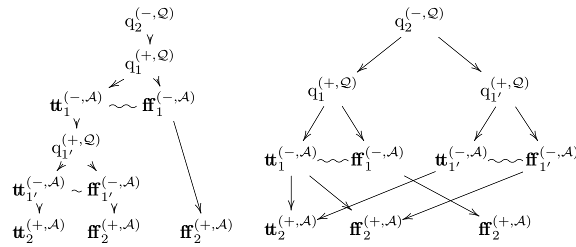

Fig. 1. Two strategies for $b : \mathbf{Bool}_1 \vdash M : \mathbf{Bool}_2$.

So we state the main theorem of [17] as:

Theorem 2. *For any terms $\Gamma \vdash M : A$ and $\Gamma \vdash N : A$ of PPCF, $M \simeq_{ctx} N$ iff $[\![M]\!]_{\mathbf{PRel}} = [\![N]\!]_{\mathbf{PRel}}$.*

Although proving this requires $\mathbf{PCoh}$, it can be forgotten when *using* the result. Accordingly, in the rest of this paper, we will only work with $\mathbf{PRel}$ and ignore biorthogonality.

C. Game semantics and event structures

The interpretation of a term in $\mathbf{PRel}$ "flattens out" its behaviour: it only displays the *multiplicity* of its use of resources, but forgets in what *order* these resources are being evaluated. This is as opposed to *game semantics*, which also records the order in which computational events are performed, or at least the causal dependencies between them. In the concurrent game semantics presented here (very close to [12]), the term $b : \mathbf{Bool} \vdash M = \textbf{if}\, b\ b\ \mathbf{ff} : \mathbf{Bool}$ can be represented by either of the two diagrams in Figure 1 (*i.e.* there will be two interpretation functions, sending M to one or the other).

These diagrams, read from top to bottom, represent dialogues (or collections of dialogues) between two players **Player** and **Opponent**, respectively playing for a program and its execution environment. Nodes, called **moves**, are computational events. Each one is due to either Player ($+$) or Opponent ($-$), as indicated by their polarity. Moves are annotated by a **Question/Answer** labelling ($\mathcal{Q}/\mathcal{A}$): **questions** correspond to variable calls, whereas **answers** correspond to calls returning. The wiggly lines denote *incompatible branchings*: moves related by them cannot occur together in an execution.

The diagram on the left is a tree, and each of its branches denotes a dialogue between Player (playing for M) and Opponent (playing for the environment) tracing one possible execution path of M. For instance, the leftmost path reads:

$\mathrm{q}_2^{(-,\mathcal{Q})}$	Opponent:	"What is the output of M (on $\mathbf{Bool}_2$)?"
$\mathrm{q}_1^{(+,\mathcal{Q})}$	Player:	"What is the value of b (on $\mathbf{Bool}_1$)?"
$\mathbf{tt}_1^{(-,\mathcal{A})}$	Opponent:	"The value of b is $\mathbf{tt}$."
$\mathrm{q}_{1'}^{(+,\mathcal{Q})}$	Player:	"Then, what is, again, the value of b?"
$\mathbf{tt}_{1'}^{(-,\mathcal{A})}$	Opponent:	"The value of b is $\mathbf{tt}$."
$\mathbf{tt}_2^{(+,\mathcal{A})}$	Player:	"Then, the output of M is $\mathbf{tt}$."

In particular, this dialogue explicitly displays the several consecutive calls to b, leaving Opponent the opportunity to change his mind. The full diagram on the left-hand side of Figure 1 appends all such dialogues together in a single picture, the wiggly lines separating incompatible branches.

But beyond simple sequential execution, our framework for game semantics, as it is based on an independence model of concurrency, also supports a partial order-based representation of parallel executions. The diagram on the right-hand side of Figure 1 represents another implementation strategy for M. Taking advantage that the order of evaluation is irrelevant in PPCF, the diagram expresses that one can evaluate the two occurrences of b *in parallel*. For each pair of results for the two independent calls to b, there is a Player answer to the original Opponent question $\mathrm{q}_2^{(-,\mathcal{Q})}$. Rather than just chronological contiguity, the arrows there describe the causal dependency of a move, *i.e.* the events that must have occurred before. We will see later that both diagrams denote (up to minor details, explained later) objects called *strategies*, representing terms. We will describe later two interpretations of PPCF as strategies: one sequential, one parallel, respectively computing the two strategies of Figure 1 from M.

Diagrams such as in Figure 1, that convey information about both *causal dependency* and *incompatibility*, are naturally formalised as *event structures*, a concurrent analogue of trees.

Definition 3. *An* ***event structure*** *is* $(E, \leq_E, Con_E)$ *with a set* E *of* ***events****,* $\leq_E$ *a partial order stipulating* causal dependency*, and Con a non-empty set of* ***consistent*** *subsets of* E*,* s.t.

$$[e] = \{e' \mid e' \leq e\} \text{ is finite for all } e \in E$$
$$\{e\} \in Con_E \text{ for all } e \in E$$
$$Y \subseteq X \in Con_E \implies Y \in Con_E$$
$$X \in Con_E \text{ and } e \leq e' \in X \implies X \cup \{e\} \in Con_E.$$

With an eye to game semantics, we equip an event structure with a function pol $: E \to \{-,+\}$*. Then* $(E, \leq_E, Con_E, pol_E)$ *is called an* ***event structure with polarity (esp)****.*

Let us 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 $x \in$ Con and x is down-closed, *i.e.* if $e \leq e' \in x$ then $e \in x$. Given an event e, write $[e]$ for $\{e' \in E \mid e' \leq e\}$ – such configurations, those with a top element, are called **prime configurations**. If E has polarity, we might give information about the polarity of events by simply annotating them 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.

If for an event structure E there is a binary relation $\#_E$ such that for all $X \subseteq E$ finite, $X \in$ Con iff $\forall e \neq e' \in X, \neg(e \#_E e')$, we say that E has **binary conflict**. In that case we automatically have that if $e \# e'$ and $e' \leq e''$ then $e \# e''$ as well (the conflict is *inherited*). If $e \# e'$ and the conflict is not inherited (meaning that for all $e_0 < e$ and $e'_0 < e'$ we have $\neg(e_0 \# e'_0)$), we say that $e \# e'$ is a **minimal conflict**, written $e \sim\sim e'$. With all that in place, it should now be clear how the diagrams of Figure 1 denote event structures (with binary conflict) where rather than $\leq_E$ and $\#_E$, we draw immediate causality $\rightarrowtriangle$ and minimal conflict $\sim\sim$.

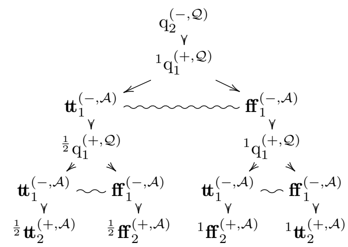

Fig. 2. A probabilistic strategy for $b : \mathbf{Bool}_1 \vdash M_+ : \mathbf{Bool}_2$

As *strategies*, we will see later that the esps of Figure 1 also come with a *labelling function* to a *game* representing the typing judgment $\mathbf{Bool} \vdash \mathbf{Bool}$, labelling from which the *annotations* $\mathrm{q}_2^{(-,\mathcal{Q})}, \mathbf{tt}_1^{(-,\mathcal{A})}, \ldots$ follow. But let us first discuss how *probability* is adjoined to event structures.

D. Event structures with probability

1) Probabilistic sequential esps: **Sequential** esps (such as that on the left of Figure 1) are those for which the causal dependency is forest-shaped, and for every configuration $x \in \mathcal{C}(E)$, if x has several distinct extensions $x \cup \{e_1^+\}, x \cup \{e_2^+\} \in \mathcal{C}(E)$ with positive events, then $x \cup \{e_1, e_2\} \notin \mathcal{C}(E)$. This means that for every $x \in \mathcal{C}(E)$, there is a set of positive extensions $\mathrm{ext}_E^+(x) = \{e_i^+ \mid i \in I\}$, all pairwise incompatible.

Sequential esps are easily enriched with probabilities, in the spirit of the game semantics of probabilistic Idealized Algol of Danos and Harmer [15]. The basic idea is that for each $x \in \mathcal{C}(E)$, Player adjoins to his set of extensions $\mathrm{ext}_E^+(x)$ a *sub-probability distribution* on $\mathrm{ext}_E^+(x)$. But rather than having a sub-distribution for each probabilistic branching in an esp, it is more convenient to carry a single **valuation**

$$v : \mathcal{C}(E) \to [0, 1]$$

putting together all the local probabilistic choices: the valuation assigned to x records all the Player probabilistic choices performed in order to reach x. Because v only records Player's probabilistic choices, it is then natural to require that *(1)* $v(\emptyset) = 1$ and *(2)* $v(x \cup \{e^-\}) = v(x)$ for any negative extension e^- of x. So as to enforce that local choices give sub-probability distributions, we also have *(3)* for all $x \in \mathcal{C}(E)$,

$$v(x) - \sum_{e \in \mathrm{ext}^+(x)} v(x \cup \{e\}) \geq 0$$

Furthermore, v is then entirely determined by the data of $v([e^+])$ for all positive $e \in E$, hence a probabilistic sequential esp can be represented by annotating positive events with the valuation of their prime configuration. Figure 2 displays the esp to be later obtained as the interpretation of the term M_+ (given in II-B2), with the probabilistic valuation written on the left of events.

2) General probabilistic esps: For non-sequential esps the axioms *(1)* and *(2)* still make sense, but finding the analogue of *(3)* is trickier, as there may be overlap between all positive extensions. This overlap leads to a redundancy in the valuation, that has to be corrected following the inclusion-exclusion principle. Following Winskel [30], we define:

Definition 4. *A* ***probabilistic esp*** *consists of an esp* $(E, \leq_E, \mathit{Con}_E, \mathit{pol}_E)$ *and a* **valuation** $v : \mathcal{C}(E) \to [0,1]$ *satisfying* (1), (2) *above, plus* (3) *if* $y \subseteq^+ x_1, \ldots, x_n$*, then*

$$\mathrm{d}_v[y; x_1, \ldots, x_n] \geq 0$$

where the ***drop*** d_v *is defined as*

$$\mathrm{d}_v[y; x_1, \ldots, x_n] = v(y) - \sum_I (-1)^{|I|+1}\ v\left(\bigcup_{i \in I} x_i\right)$$

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

We pointed out in the beginning of Section II-C that the deterministic term M can be interpreted by either esp in Figure 1 – likewise, the probabilistic term M_+ can be interpreted by the probabilistic esp of Figure 2, or by some probabilistic version of the right hand side diagram of Figure 1. However, unlike for *sequential* probabilistic esps, for general ones the valuation cannot always be pushed to events and has to remain on configurations. Therefore we have to adjoin to the diagram on the right of Figure 1 the value of v for all configurations: in the case of M_+ a configuration has valuation $\frac{1}{2}$ if it contains the right occurrence of $\mathrm{q}_1^{(+,\mathcal{Q})}$, and 1 otherwise.

E. Games and strategies-as-esps

Until now, we have explained the formal nature of the strategies interpreting terms as (probabilistic) esps, but we have not said what *games* they play on. As usual in game semantics, the games will be abstract representation of *types*. and will describe the causal dependencies between computational events made available by a type.

1) Arenas and pre-strategies: The games (*arenas*) will themselves be certain esps – a type A will be interpreted by a arena $[\![A]\!]$, listing all the computational events existing in a call-by-name execution on this type and specifying the causality and compatibility constraints on these events. The arena will also remember the polarity of each event, and whether it is a question or an answer.

Consider the ground types $\mathbf{Bool}$ and $\mathbf{Nat}$. There are only two events available between an execution environment and a term of ground type: the environment starting the evaluation of the term (Opponent question) and the evaluation finishing (Player answer). Accordingly, the corresponding arenas are:

$[\![\mathbf{Bool}]\!] =$

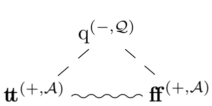

$[\![\mathbf{Nat}]\!] =$

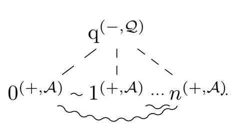

Again, the diagrams are read from top to bottom – immediate causality in arenas is represented by dashed lines rather than arrows, to keep it easily distinguishable from causality in strategies. Although the two notions have the same formal nature, they play a different role in the development.

In a typing judgment such as $\mathbf{Bool}_1 \vdash \mathbf{Bool}_2$ there are more computational events available: upon receiving the initial question on $\mathbf{Bool}_2$, Player might (as in Figures 1 and 2) interrogate $\mathbf{Bool}_1$, where polarity is reversed. In fact, in our running examples Player interrogates $\mathbf{Bool}_1$ *twice*, showing the need to create copies of $\mathbf{Bool}_1$. Accordingly, the sequent $\mathbf{Bool}_1 \vdash \mathbf{Bool}_2$ will get interpreted by the arena:

$[\![\mathbf{Bool}_1 \vdash \mathbf{Bool}_2]\!] =$

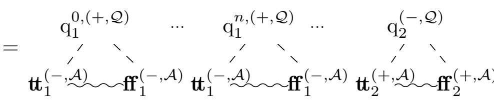

Note the new annotations $\mathrm{q}^{i,(+,\mathcal{Q})}$ in copies of the initial question of the argument. This **copy index** i is implicit in the moves $\mathrm{q}_1^{(+,\mathcal{Q})}$ in Figures 1 and 2. They will be introduced formally via an *exponential modality*.

We now give the general definition of arenas.

Definition 5. *An* ***arena*** *consists of a esp* A*, and a* ***labelling function*** $\lambda_A : A \to \{\mathcal{Q}, \mathcal{A}\}$ *such that:*
A is a forest: *if* $a_1 \leq a_3$ *and* $a_2 \leq a_3$*,* $a_1 \leq a_2$ *or* $a_2 \leq a_1$*.*
A is alternating: *if* $a_1 \rightarrow a_2$ *then* $\mathit{pol}(a_1) \neq \mathit{pol}(a_2)$*.*
Questions: *if* a_1 *is minimal or if* $a_1 \rightarrow a_2$ *then* $\lambda_A(a_1) = \mathcal{Q}$*.*
Answering is affine: *every* $x \in \mathcal{C}(A)$ *is* affine, i.e. *for every* $a_1 \in x$ *with* $\lambda_A(a_1) = \mathcal{Q}$*, there is at most one* $a_2 \in x$ *such that* $a_1 \rightarrow a_2$ *and* $\lambda_A(a_2) = \mathcal{A}$*.*

An arena A *(or esp in general) is* ***negative*** *if every minimal event in* A *is negative.*

2) Strategies: Now that we have our notion of games, we can finish making formal the diagrams of Figures 1 and 2 and define what is a (probabilistic) strategy on an arena.

As pointed out before, the diagrams of Figure 1 have to be understood as representing esps *labelled by the arena*, here $[\![\mathbf{Bool}_1 \vdash \mathbf{Bool}_2]\!]$. Modulo the (arbitrary) choice of copy indices for occurrences of $\mathrm{q}_1^{(+,\mathcal{Q})}$, this labelling function is implicit in the name of nodes of the diagram. However, not all such labelled esps make sense as strategies. In order to have a well-behaved notion of strategy, we will now give a number of further constraints, best introduced in multiple stages.

First, we introduce *pre-strategies*.

Definition 6. *A (probabilistic)* ***pre-strategy*** *on arena* A *is a (probabilistic) esp* S *along with a* labelling function

$$\sigma : S \to A$$

such that (1) *for all* $x \in \mathcal{C}(S)$*, the direct image* $\sigma\, x \in \mathcal{C}(A)$ *is a configuration of the game, and* (2) σ *is* locally injective*: for all* $s_1, s_2 \in x \in \mathcal{C}(S)$*, if* $\sigma\, s_1 = \sigma\, s_2$ *then* $s_1 = s_2$*.*

Conditions *(1)* and *(2)* amount to the fact that the function on events $\sigma : S \to A$ is also a **map of event structures** [28] from S to A (ignoring here the further structure on S and A).

Although pre-strategies give a reasonable mathematical description of concurrent processes performed under the rules of a game (or protocol) A, it is too general: in particular, the current definition ignores polarity. Even in a sequential world, we expect of a definition of strategy that *e.g.* Player cannot constrain the behaviour of Opponent further than what is specified by the rules of the game. For our strategies on event structures, Rideau and Winskel [26] proved that we need more in order to get a category. They define:

Definition 7. *A pre-strategy $\sigma : S \to A$ is a **strategy** iff it is **receptive**: for $x \in \mathcal{C}(S)$, if $\sigma x \subseteq^- y \in \mathcal{C}(A)$, there is a unique $x \subseteq x' \in \mathcal{C}(S)$* s.t. *$\sigma x' = y$; and **courteous**: for $s, s' \in S$, if $s \rightarrowtriangle_S s'$ and if $pol(s) = +$ or $pol(s') = -$, then $\sigma s \rightarrowtriangle_A \sigma s'$.*

Thus a strategy can only pick the *positive* events it wants to play, and for each of those, which Opponent moves need to occur before. It was proved in [26] and further detailed in [10] that strategies can be composed, and form a category (up to isomorphism) whose structure we will revisit in the next section, aiming for an interpretation of PPCF.

But for now we still have some definitions to give on strategies. Indeed although at this point the causal structure of strategies is well-behaved enough to fit in a compositional setting, as per usual in game semantics strategies have to be restricted further to ensure that they "behave like terms of PPCF". Typically, a set of further conditions on strategies is deemed adequate when it induces a *definability result*, leading to full abstraction. Here instead, our conditions will ensure that there is a functorial *collapse* operation from games to the already fully abstract probabilistic relational model.

Our further conditions are a subset of those of [12]. They crucially rely on the following definition.

Definition 8. *A **grounded causal chain (gcc)** on an esp S is $\rho = \{\rho_1, \dots, \rho_n\} \subseteq S$* s.t. *$\rho_1 \in \min(S)$ (ρ_1 is minimal), and*

$$\rho_1 \rightarrowtriangle_S \rho_2 \rightarrowtriangle_S \rho_3 \rightarrowtriangle_S \dots \rightarrowtriangle_S \rho_n$$

Note that some ρ_i may have dependencies that are not met in ρ. We write $\mathrm{gcc}(S)$ for the set of gccs in S.

Grounded causal chains give a notion of *thread* in this concurrent setting. The following definition ensures that each thread can be regarded as a standalone sequential program:

Definition 9. *A strategy $\sigma : S \to A$ is **visible** iff for all $\rho \in \mathrm{gcc}(S)$, we have $\sigma \rho \in \mathcal{C}(A)$.*

As arenas are forest-shaped, any non-minimal $a \in A$ has a unique predecessor $\mathrm{just}(a) \rightarrowtriangle_A a$. Likewise, by local injectivity of σ, for any $s \in S$ whose image is non-minimal there is a unique $s' \in S$ such that $\sigma\, s' \rightarrowtriangle_A \sigma\, s$, which we also refer to as $\mathrm{just}(s) = s'$ its **justifier**.

With that in mind, the visibility of $\sigma : S \to A$ can be equivalently stated by asking that for all $\rho \in \mathrm{gcc}(S)$, for each $\rho_i \in \rho$, we have $\mathrm{just}(\rho_i) \in \rho$ as well. This is reminiscent of the visibility condition in HO games, which states that the justifier of a Player move always happens within the P-view [18]. In our setting however, visibility means that a visible strategy can be regarded as a bag of sequential threads, sometimes forking with each other, sometimes merging, and sometimes conflicting. The strategy pictured in Figure 3 is non-visible, since the gcc $\mathrm{q}_{1'} \rightarrowtriangle \mathtt{tt}_1$ does not contain the justifier of $\mathtt{tt}_1$.

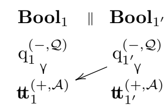

Fig. 3. A non-visible strategy on $\mathbf{Bool}_1 \parallel \mathbf{Bool}_{1'}$.

Each of these sequential threads needs to respect the call-return discipline, in order to forbid strategies behaving like *e.g.* call/cc [19]. In a set $X \subseteq S$, we say that an answer $s_2^{\mathcal{A}} \in X$ (which is shortcut for $\lambda_A(\sigma\, s_2) = \mathcal{A}$) **answers** a question $s_1^{\mathcal{Q}} \in X$ iff $\sigma\, s_1 \rightarrowtriangle_A \sigma\, s_2$ (*i.e.*, $\mathrm{just}(s_2) = s_1$). If a gcc $\rho \in \mathrm{gcc}(S)$ has some unanswered questions, we say that its **pending question** is the latest unanswered question, *i.e.* the maximal unanswered question for $\leq_S$.

We import from HO games [18]:

Definition 10. *A visible strategy $\sigma : S \to A$ is **well-bracketed** iff for all $\rho = \{\rho_1 \rightarrowtriangle_S \dots \rightarrowtriangle_S \rho_{n+1}^{\mathcal{A}}\} \in \mathrm{gcc}(S)$, ρ_{n+1} answers the* pending question *of $\{\rho_1 \rightarrowtriangle_S \dots \rightarrowtriangle_S \rho_n\}$.*

The games model of [12] had these notions of visibility and well-bracketing, but also required a few others in order to achieve intensional full abstraction: determinism, innocence, and a further well-bracketing condition. Here we do not have determinism for obvious reasons. It is perhaps more surprising that we do not need either of innocence, the further well-bracketing conditions, or to restrict at all the shape of conflict. We would need them if we aimed for a definability result; but this is avoidable here as intensional full abstraction will not be proved via definability but instead via a collapse to $\mathbf{PRel}$.

III. Probabilistic Thin Concurrent Games

In the previous section we introduced strategies as static objects. However, in order to build a compositional interpretation of PPCF we need a *cartesian closed category* of probabilistic strategies. This means that we need to compose strategies, but also to overcome the intrinsically linear nature of basic concurrent strategies. So as to authorize multiple accesses to resources, as pointed out before we will duplicate them by adjoining *copy indices* – but as usual with copy indices [3] we will then need to consider strategies *up to symmetry*, *i.e.* quotient out the specific choice of copy indices.

Hence the main goal of this section is to marry the probabilistic strategies of [30] presented before with the *thin concurrent games with symmetry* developed in [12], [13] as a foundation for a cartesian closed category of concurrent games. We will then build on top of it a symmetric monoidal closed category, with a linear exponential comonad [24].

A. Concurrent games with symmetry

1) Symmetry in event structures: We first review the basics of *event structures with symmetry* [29], presented here as in [12] via *isomorphism families*.

Definition 11. *An **isomorphism family** on an event structure E is a set $\widetilde{E}$ of bijections $\theta : x \cong y$, where $x, y \in \mathcal{C}(E)$,* s.t.

(1) For all $x \in \mathcal{C}(E)$, $id_x : x \cong x \in \widetilde{E}$.
(2) If $\theta : x \cong y \in \widetilde{E}$ then $\theta^{-1} : y \cong x \in \widetilde{E}$.
(3) If $\theta : x \cong y$ and $\eta : y \cong z \in \widetilde{E}$ then $\eta \circ \theta : x \cong z \in \widetilde{E}$.
(4) If $\theta : x \cong y \in \widetilde{E}$ and $x \subseteq x' \in \mathcal{C}(E)$, then there exists $y \subseteq y' \in \mathcal{C}(E)$ and $\theta' : x' \cong y' \in \widetilde{E}$ such that $\theta \subseteq \theta'$.
(5) If $\theta : x \cong y \in \widetilde{E}$ and $x' \subseteq x \in \mathcal{C}(E)$, then there exists $y' \subseteq y \in \mathcal{C}(E)$ and $\theta' : x' \cong y' \in \widetilde{E}$ such that $\theta' \subseteq \theta$.

*An **event structure with symmetry (ess)** is a pair $\mathcal{E} = (E, \widetilde{E})$ where $\widetilde{E}$ is an isomorphism family on E. If E additionally has polarities, then the bijections in $\widetilde{E}$ are furthermore required to preserve them; $\mathcal{E}$ is then an **essp**.*

Conditions (1), (2) and (3) give $\widetilde{E}$ a groupoid structure, while (4) and (5) ensure that symmetric configurations have bisimilar future and isomorphic past. We regard bijections as sets of pairs, justifying the notation $\theta \subseteq \theta'$ (or $\subseteq^+$ and $\subseteq^-$ if E has polarities). If $\mathcal{E}$ and $\mathcal{F}$ are ess, a map of es $f : E \to F$ **preserves symmetry** if for every $\theta : x \cong_{\widetilde{E}} y$ (shorthand for $\theta : x \cong y \in \widetilde{E}$), the bijection $f\theta = \{(fe, fe') \mid (e, e') \in \theta\}$ is in $\widetilde{F}$; we write $f : \mathcal{E} \to \mathcal{F}$.

Symmetry and probability can be combined:

Definition 12. *A **probabilistic essp** is an essp $(E, \widetilde{E})$ and a valuation v on E such that $v(x) = v(y)$ whenever $\theta : x \cong_{\widetilde{E}} y$.*

In other words, symmetric configurations of a probabilistic essp must have the same probability valuation.

2) Thin concurrent games: We use $\mathcal{A}, \mathcal{B}, \mathcal{S}, \mathcal{T}, \dots$ to denote essps, keeping the underlying event structures $(A, B, \dots)$ and isomorphism families $(\widetilde{A}, \widetilde{B}, \dots)$ implicit.

The construction on games introducing symmetry, and which drives the notion of essps, is the exponential $!\mathcal{A}$. It is a *symmetric*, *infinitary* form of *parallel composition*:

Definition 13. *Given a family $(E_i)_{i \in I}$ of event structures, we define their **simple parallel composition** to have events:*

$$\|_{i \in I} E_i = \bigcup_{i \in I} \{i\} \times E_i$$

with componentwise causal ordering. *The* consistent sets *are the finite sets of the form $\|_{i \in I_0} X_i$ for $I_0 \subseteq I$ and $X_i \in Con_{E_i}$ for all $i \in I_0$.* Polarity*, if present, is inherited componentwise.* Symmetry*, if present, is given by $\|_{i \in I} \widetilde{E_i}$,* i.e. *with bijections $\theta : \|_{i \in I_0} x_i \cong \|_{i \in I_0} y_i$ induced by a family $(\theta_i : x_i \cong_{\widetilde{E_i}} y_i)_{i \in I_0}$ such that for all $(i, e) \in \|_{i \in I_0} x_i$, $\theta((i, e)) = (i, \theta_i\, e_i)$.*

We often write $\|_i\, x_i \in \mathcal{C}(\|_i\, E_i)$ for configurations of a parallel composition, or $x \parallel y \in \mathcal{C}(A \parallel B)$ in the binary case.

Definition 14. *Let $\mathcal{A}$ be a **negative** essp,* i.e. *A is negative. Then, $!\mathcal{A}$ is defined as $\|_{i \in \omega}\, \mathcal{A}$, with isomorphism family enriched to comprise the bijections $\theta : \|_{i \in I}\, x_i \cong \|_{j \in J}\, y_j$ such that there exists a permutation $\pi : I \cong J$ and a family $(\theta_i \in \widetilde{A})_{i \in I}$, with $\theta((i, a)) = (\pi\, i, \theta_i\, a)$ for all $(i, a) \in \|_{i \in I} x_i$.*

This is very similar to the equivalence relation on the game $!A$ in AJM games [3], and was also considered in [11]. Note that this ! operation is not the same as the one used in [12] and which duplicates all moves of the game "*in depth*" rather than just at the surface – in the spirit of HO games [18]. We prefer here this "*surface*" version, which allows an easier connection with the relational model as both cartesian closed categories are then obtained as Kleisli categories.

Very soon, strategies will be considered *up to the choice of copy indices*. But this is naively not preserved under composition – for it to be a congruence, strategies also have to be *uniform*: the behaviour of a strategy should not depend on the copy indices used by Opponent, although his choice of copy indices will. Constructing a framework of concurrent games where being "*the same up to copy indices*" is a congruence is quite challenging, see *e.g.* [13] for a discussion. One solution, used in [11], is to ask that all strategies are *saturated*, and play non-deterministically all possible copy indices. Another, introduced in [12] and detailed in [13], requires instead that strategies pick copy indices deterministically (are *thin*, see Definition 16). For *thin* strategies to behave well we also must constrain the games, and separate *Player* permutations and *Opponent* permutations, in a way that is very reminiscent of Melliès' notion of uniformity [21] by bi-invariance under the action of two groups of Opponent and Player permutations.

Definition 15. *A **thin concurrent game (tcg)** is $\mathcal{A} = (A, \widetilde{A}, \widetilde{A}_-, \widetilde{A}_+)$ where A is an esp, and $\widetilde{A}, \widetilde{A}_-$ and $\widetilde{A}_+$ are isomorphism families on A included in $\widetilde{A}$, such that:*

(1) If $\theta \in \widetilde{A}_+ \cap \widetilde{A}_-$ then $\theta = id_x$ for some $x \in \mathcal{C}(A)$,
(2) If $\theta \in \widetilde{A}_-$ and $\theta \subseteq^- \theta' \in \widetilde{A}$ then $\theta' \in \widetilde{A}_-$,
(3) If $\theta \in \widetilde{A}_+$ and $\theta \subseteq^+ \theta' \in \widetilde{A}$ then $\theta' \in \widetilde{A}_+$.

When $\mathcal{A}$ is a negative tcg, Opponent is responsible for the first layer of symmetry in $!\mathcal{A}$: the family $\widetilde{!A}_-$ comprises all $\theta : x \cong y$ such that for all $i \in \omega$, $\theta_i : x_i \cong y_{\pi(i)} \in \widetilde{A}_-$. On the other hand the family $\widetilde{!A}_+$ comprises all $\theta : x \cong y$ such that for all $i \in I$, $\pi\, i = i$ and $\theta_i \in \widetilde{A}_+$.

While the dual definition could also be given for *positive* A, candidates of $\widetilde{!A}_-$ and $\widetilde{!A}_+$ for A with minimal events of mixed polarities inevitably fail some axioms of tcgs (and their intended consequences) – building an exponential without any assumption on polarity requires *saturation* [6], [11].

3) Probabilistic $\sim$-strategies: We now add probability to the uniform strategies of [12], [13], called *$\sim$-strategies*.

Definition 16. *A **probabilistic $\sim$-strategy** on a tcg $\mathcal{A}$ is a map of essps $\sigma : \mathcal{S} \to \mathcal{A}$ (where $\mathcal{A} = (A, \widetilde{A})$, ignoring $\widetilde{A}_+$ and $\widetilde{A}_-$)* s.t. *$\mathcal{S}$ is a probabilistic essp, $\sigma : S \to A$ a strategy, and:*

*(1) σ is **strong-receptive**: if $\theta \in \widetilde{S}$ and $\sigma\theta \subseteq^- \eta \in \widetilde{A}$, then there exists a unique $\theta \subseteq \theta' \in \widetilde{S}$ such that $\sigma\theta' = \eta$.*
*(2) $\mathcal{S}$ is **thin**: for $\theta : x \cong_{\widetilde{S}} y$* s.t. *$x' = x \cup \{s\} \in \mathcal{C}(S)$ with $pol(s) = +$, there is a* unique *$t \in S$* s.t. *$\theta \cup \{(s, t)\} \in \widetilde{S}$.*

The remaining concepts of Section II-C extend in the presence of symmetry: a **$\sim$-arena** is a tcg $\mathcal{A}$ with a $\mathcal{Q}/\mathcal{A}$ labelling λ on A, such that (A, λ) is an arena and every bijection in $\widetilde{A}$ preserves the action of λ. A $\sim$-strategy $\sigma : \mathcal{S} \to \mathcal{A}$ on a $\sim$-arena $\mathcal{A}$ is visible (resp. well-bracketed) when the underlying strategy $S \to A$ is visible (resp. well-bracketed).

We will first develop the categorical structure of $\sim$-strategies without $\mathcal{Q}/\mathcal{A}$ labelling, visibility and well-bracketing, which will only be reinstated in Section III-C when modelling PPCF.

B. A category

To define the morphisms of our category, we first give some constructions on tcgs. The **dual** of a tcg $\mathcal{A} = (A, \widetilde{A}, \widetilde{A}_-, \widetilde{A}_+)$ is the tuple $\mathcal{A}^\perp = (A^\perp, \widetilde{A}, \widetilde{A}_+, \widetilde{A}_-)$ where $A^\perp$ is the esp whose events, causality and consistency are exactly those of A, but polarity is reversed: $\mathrm{pol}_{A^\perp}(a) = -\mathrm{pol}_A(a)$. If $\mathcal{A}$ and $\mathcal{B}$ are tcgs, then $\mathcal{A} \parallel \mathcal{B}$ is a tcg with negative (resp. positive) isomorphism family set to $\widetilde{A}_- \parallel \widetilde{B}_-$ (*resp.* $\widetilde{A}_+ \parallel \widetilde{B}_+$).

A **$\sim$-strategy from $\mathcal{A}$ to $\mathcal{B}$** is defined to be a $\sim$-strategy on the tcg $\mathcal{A}^\perp \parallel \mathcal{B}$; we now investigate how to *compose* $\sim$-strategies. As usual in game semantics composing strategies involves two steps: interaction and hiding. We will first spell them out without probabilities, and then add it back.

1) Interaction of $\sim$-strategies: Let $\mathcal{A}, \mathcal{B}$ and $\mathcal{C}$ be tcgs, and $\sigma : \mathcal{S} \to \mathcal{A}^\perp \parallel \mathcal{B}$ and $\tau : \mathcal{T} \to \mathcal{B}^\perp \parallel \mathcal{C}$ be $\sim$-strategies. As a first approximation, *states* of the interaction $\tau \circledast \sigma$ should correspond to so-called *sychronised* pairs:

$$\{(x_S, x_T) \mid \sigma\, x_S = x_A \parallel x_B \;\&\; \tau\, x_T = x_B \parallel x_C\}$$

According to this, the interaction of σ of Figure 3 with either τ_l or τ_r from Figure 1 (regarded as strategies on $(\mathbf{Bool}_1 \parallel \mathbf{Bool}_{1'})^\perp \parallel \mathbf{Bool}_2$) would have the same maximal state

$$(\{\mathrm{q}_1, \mathrm{q}_{1'}, \mathbf{tt}_1, \mathbf{tt}_{1'}\}, \{\mathrm{q}_2, \mathrm{q}_1, \mathrm{q}_{1'}, \mathbf{tt}_1, \mathbf{tt}_{1'}, \mathbf{tt}_2\})$$

However this seems inaccurate, because while σ wants to play $\mathbf{tt}_1$ *after* $\mathrm{q}_{1'}$, τ_l will only ask $\mathrm{q}_{1'}$ *after* σ plays $\mathbf{tt}_1$: there is a *causal loop*. To get an ess whose configurations correspond to *causally reachable* pairs of synchronised configurations, we use the following *pullback* in the category of event structures with symmetry, which we know exists from [12], [13]:

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

From either path of the pullback we get the **interaction** $\tau \circledast \sigma : \mathcal{T} \circledast \mathcal{S} \to \mathcal{A} \parallel \mathcal{B} \parallel \mathcal{C}$, a *labelled ess*, whose underlying event structure is charaterised in *e.g.* [13]:

Lemma 17. *Configurations of $T \circledast S$ are in one-to-one correspondence with the synchronised pairs*

$$\{(x_S, x_T) \mid \sigma\, x_S = x_A \parallel x_B \;\&\; \tau\, x_T = x_B \parallel x_C\}$$

that are causally reachable. *Formally, the induced bijection*

$$\varphi : x_S \parallel x_C \simeq x_A \parallel x_T$$

*is **secured**,* i.e. *the relation on the graph of φ generated by $(s,t) \lhd_\varphi (s',t')$ if $s \leq s'$ or $t \leq t'$ is a partial order.*

The maximal state in the interaction of σ and τ_l above is $(\{\mathrm{q}_1\}, \{\mathrm{q}_2, \mathrm{q}_1\})$. It cannot be extended further, as we have a *deadlock*: both strategies are waiting for the other. Likewise, the *isomorphisms* between (x_S, x_T) and (y_S, y_T) are found to coincide with pairs (θ_S, θ_T) such that $\theta_S : x_S \cong_{\widetilde{S}} y_S$, $\theta_T : x_T \cong_{\widetilde{T}} y_T$, $\sigma\,\theta_S = \theta_A \parallel \theta_B$ and $\tau\,\theta_T = \theta_B \parallel \theta_C$.

This process of eliminating causal loops is the main difference between game semantics and relational semantics; and the reason why typically mapping game semantics to relational-like models is *not functorial*, as in *e.g.* [31]. Accordingly our main result will rely on Lemma 24, which states that the composition of *visible* strategies is always deadlock-free.

2) Composition of $\sim$-strategies: Following [12], [13], from $\tau \circledast \sigma : \mathcal{T} \circledast \mathcal{S} \to \mathcal{A} \parallel \mathcal{B} \parallel \mathcal{C}$, we set $T \odot S$ to comprise the events of $T \circledast S$ mapped to either A or C, with the data of an event structure inherited. Thus, each $x \in \mathcal{C}(T \odot S)$ has a **unique witness** $[x]_{T \circledast S} \in \mathcal{C}(T \circledast S)$. *Polarities* in $T \odot S$ are set so that the restriction $\tau \odot \sigma : T \odot S \to A^\perp \parallel C$ preserves them. The *isomorphism family* $\widetilde{T \odot S}$ is set to comprise the bijections $\theta : x \cong y$ $(x, y \in \mathcal{C}(T \odot S))$ such that $\theta \subseteq \theta'$ for $\theta' : [x]_{T \circledast S} \cong_{\widetilde{T \circledast S}} [y]_{T \circledast S}$. From this we get the **composition** of σ and τ, a $\sim$-strategy $\tau \odot \sigma : \mathcal{T} \odot \mathcal{S} \to \mathcal{A}^\perp \parallel \mathcal{C}$ [13].

3) Composition of probabilistic $\sim$-strategies: Assume now that σ and τ are probabilistic $\sim$-strategies. First, we make the interaction $\mathcal{T} \circledast \mathcal{S}$ probabilistic by setting, for $x \in \mathcal{C}(T \circledast S)$

$$v_{T \circledast S}(x) = v_S(x_S) \times v_T(x_T)$$

where $\Pi_1\, x = x_S \parallel x_C$ and $\Pi_2\, x = x_A \parallel x_T$. For $x \in \mathcal{C}(T \odot S)$, we set $v_{T \odot S}(x) = v_{T \circledast S}([x]_{T \circledast S})$. From [30], we know that this makes $\tau \odot \sigma$ a probabilistic strategy – it remains to check that the valuation $v_{T \odot S}$ is invariant under symmetry, which is immediate. We have defined

$$\tau \odot \sigma : \mathcal{T} \odot \mathcal{S} \to \mathcal{A}^\perp \parallel \mathcal{C},$$

a probabilistic $\sim$-strategy from $\mathcal{A}$ to $\mathcal{C}$.

4) The probabilistic copycat $\sim$-strategy: The identity strategy on a tcg $\mathcal{A}$ is a map $\text{æ}_{\mathcal{A}} : \mathbb{C}\mathbb{C}_{\mathcal{A}} \to \mathcal{A}^\perp \parallel \mathcal{A}$ called the **copycat $\sim$-strategy**. The event structure $\mathbb{C}\mathbb{C}_A$ has events, consistent subsets and polarity those of $A^\perp \parallel A$, and causality relation $\leq_{\mathbb{C}\mathbb{C}_A}$ defined as the transitive closure of

$$\begin{aligned} \leq_{A^\perp \parallel A} \cup \{((1,a),(2,a)) \mid \mathrm{pol}_{A^\perp}(1,a) = -\} \\ \cup \{((2,a),(1,a)) \mid \mathrm{pol}_A(2,a) = -\}. \end{aligned}$$

Configurations of $\mathbb{C}\mathbb{C}_A$ are certain configurations $x_1 \parallel x_2 \in \mathcal{C}(A^\perp \parallel A)$. The isomorphism family $\widetilde{\mathbb{C}\mathbb{C}_{\mathcal{A}}}$ comprises all

$$\theta = \theta_1 \parallel \theta_2 : x_1 \parallel x_2 \cong y_1 \parallel y_2$$

such that $\theta_1, \theta_2 \in \widetilde{A}$ and such that θ is an order-isomorphism.

Finally copycat is made probabilistic. In fact copycat is *deterministic* [30], hence the constant function assigning probability 1 to every configuration is a valid valuation [30]. Under these definitions the map $\text{æ}_{\mathcal{A}} : \mathbb{C}\mathbb{C}_{\mathcal{A}} \to \mathcal{A}^\perp \parallel \mathcal{A}$ is a probabilistic $\sim$-strategy.

5) Equivalences of strategies: It is often not sensible to compare $\sim$-strategies up to strict equality; for instance the associativity and identity laws for composition only hold up to *isomorphism of $\sim$-strategies*. Let $\sigma : \mathcal{S} \to \mathcal{A}$ and $\tau : \mathcal{T} \to \mathcal{A}$ be probabilistic $\sim$-strategies on a tcg $\mathcal{A}$. A **strong morphism** from σ to τ is a map of essp $f : \mathcal{S} \to \mathcal{T}$ such that $\tau \circ f = \sigma$, and such that for all $x \in \mathcal{C}(S)$, $v_S(x) \leq v_T(fx)$. The probabilistic $\sim$-strategies σ and τ are *strongly isomorphic* if there are morphisms $f : \mathcal{S} \to \mathcal{T}$ and $g : \mathcal{T} \to \mathcal{S}$ of probabilistic $\sim$-strategies which are inverses as maps of essp.

Tcgs, $\sim$-strategies, and strong morphisms form a bicategory [13], which extends to *probabilistic* $\sim$-strategies. But strong isomorphisms do not exploit symmetry, and distinguish between strategies playing the same moves *up to copy indices*. We aim for a weaker notion of isomorphism of $\sim$-strategy which we will use to quotient our bicategory. We recall:

Definition 18. *Two maps $f, g : \mathcal{S} \to \mathcal{A}$ of ess are* ***symmetric****, written $f \sim g$, if for all $x \in \mathcal{C}(S)$, the bijection $\theta_x : \{(fs, gs) \mid s \in x\}$ is in $\widetilde{A}$. If moreover $\mathcal{A}$ is a tcg, say f and g are* ***positively symmetric****, written $f \sim^+ g$, if $\theta_x \in \widetilde{A}_+$ for all x.*

A **weak morphism** of probabilistic $\sim$-strategies from $\sigma : \mathcal{S} \to \mathcal{A}$ to $\tau : \mathcal{T} \to \mathcal{A}$ is a map of ess $f : \mathcal{S} \to \mathcal{T}$ such that $\tau \circ f \sim^+ \sigma$, and such that for all $x \in \mathcal{C}(S)$, $v_S(x) \leq v_T(fx)$. The induced notion of **weak isomorphism** yields a weaker notion of equivalence between $\sim$-strategies which we use to quotient our bicategory. A key result of [12], [13] is that weak isomorphism is preserved under composition, which crucially depends on the thinness axiom for $\sim$-strategies.

Hence there is a category PTCG with as objects the tcgs, and as morphisms from $\mathcal{A}$ to $\mathcal{B}$ the probabilistic $\sim$-strategies on $\mathcal{A}^\perp \parallel \mathcal{B}$, *up to weak isomorphism*. In fact, PTCG is *compact closed* – but we skip this construction, and restrict its tcgs and strategies to get a model of ILL.

C. A model of Intuitionistic Linear Logic and PPCF

We now build a subcategory $\mathbf{PG}$ of PTCG, adequate to interpret PPCF and perform the collapse operation. Its objects are negative $\sim$-arenas, and its morphisms from $\mathcal{A}$ to $\mathcal{B}$ are the probabilistic $\sim$-strategies $\sigma : \mathcal{S} \to \mathcal{A}^\perp \parallel \mathcal{B}$ (up to weak isomorphism), which are *visible and well-bracketed* (see Section II-E2), but also **negative** (*i.e.* $\mathcal{S}$ is negative) and **well-threaded** (for all $s \in S$, $[s]$ has exactly one initial move). Negativity is needed for the cartesian structure, and well-threadedness for monoidal closure. From [30], [12], these are stable under composition. We will sometimes write $\sigma : \mathcal{A} \overset{\mathbf{PG}}{\nrightarrow} \mathcal{B}$ for $\sigma : \mathcal{S} \to \mathcal{A}^\perp \parallel \mathcal{B}$, keeping the $\mathcal{S}$ anonymous.

1) Monoidal structure: The **tensor** $\mathcal{A} \otimes \mathcal{B}$ is simply defined as $\mathcal{A} \parallel \mathcal{B}$, with unit $\mathbf{1}$ the $\sim$-arena with no events. From $\sigma_1 : \mathcal{S}_1 \to \mathcal{A}_1^\perp \parallel \mathcal{B}_1$ and $\sigma_2 : \mathcal{S}_2 \to \mathcal{A}_2^\perp \parallel \mathcal{B}_2$, form

$$\sigma_1 \otimes \sigma_2 : \mathcal{S}_1 \parallel \mathcal{S}_2 \to (\mathcal{A}_1 \otimes \mathcal{A}_2)^\perp \parallel (\mathcal{B}_1 \parallel \mathcal{B}_2)$$

as obvious from $\sigma_1 \parallel \sigma_2$; with the valuation $v_{S_1 \otimes S_2}(x_1 \parallel x_2) = v_{S_1}(x_1) \times v_{S_2}(x_2)$. Without probabilities, we know from [12], [13] that this yields a symmetric monoidal structure; the extension with probabilities offers no difficulty.

2) Cartesian structure: The empty $\sim$-arena $\mathbf{1}$ is a terminal object. The **cartesian product** of $\sim$-arenas $\mathcal{A}$ and $\mathcal{B}$, written $\mathcal{A} \,\&\, \mathcal{B}$, has events, causality, and polarity those of $\mathcal{A} \parallel \mathcal{B}$, and consistent subsets those $X = X_A \parallel \emptyset$ with $X_A \in \mathrm{Con}_A$ or $X = \emptyset \parallel X_B$ with $X_B \in \mathrm{Con}_B$. The isomorphism family $\widetilde{A \,\&\, B}$, and its negative and positive subfamilies are obtained as restrictions of those of $\widetilde{A \parallel B}$. We have two **projections**:

$$\varpi_{\mathcal{A}} : \mathbb{C}_{\mathcal{A}} \to (\mathcal{A} \,\&\, \mathcal{B})^\perp \parallel \mathcal{A} \qquad \varpi_{\mathcal{B}} : \mathbb{C}_{\mathcal{B}} \to (\mathcal{A} \,\&\, \mathcal{B})^\perp \parallel \mathcal{B}$$

where one component of the $\&$ is not reached – this is compatible with receptivity since $\mathcal{A}, \mathcal{B}$ are negative.

From $\sigma : \mathcal{S} \to \mathcal{A}^\perp \parallel \mathcal{B}$ and $\tau : \mathcal{T} \to \mathcal{A}^\perp \parallel \mathcal{C}$, their **pairing**

$$\langle \sigma, \tau \rangle : \mathcal{S} \,\&\, \mathcal{T} \to \mathcal{A}^\perp \parallel (\mathcal{B} \,\&\, \mathcal{C})$$

is obtained from σ and τ in the obvious way. The valuation is $v_{S \& T}(x_S \parallel \emptyset) = v_S(x_S)$ and $v_{S\&T}(\emptyset \parallel x_T) = v_T(x_T)$. The incompatibility between $\mathcal{B}$ and $\mathcal{C}$ is key in ensuring local injectivity. Compatibility of pairing and projections, along with surjective pairing, are easy verifications.

3) Closed structure: A difficulty here, is that because our objects are *negative* $\sim$-arenas, $\mathcal{A}^\perp \parallel \mathcal{B}$ usually lies outside of $\mathbf{PG}$. So, inspired by the usual arena construction in HO game semantics, we shall deviate from $A^\perp \parallel B$ by having A depend on the set $\min(B)$ of minimal events of B. If there are several of them, we will copy A accordingly. But unlike HO games, our setting is sensitive to linearity – hence we will use consistency to ensure that this copying remains linear.

Definition 19. *Consider $\mathcal{A}$, $\mathcal{B}$ two negative $\sim$-arenas. The $\sim$-arena has as events $A \multimap B = (\parallel_{b \in \min(B)} A^\perp) \parallel B$ and polarity induced. The* causal order *is that above, enriched with pairs $((2, b), (1, (b, a)))$ for each $b \in \min(B)$ and $a \in A$.*

Notice that there is a function

$$\begin{array}{rccl} \chi_{A,B} & : & A \multimap B & \to & A^\perp \parallel B \\ & & (1, (b, a)) & \mapsto & (1, a) \\ & & (2, b) & \mapsto & (2, b) \end{array}$$

collapsing all copies. We set $\mathrm{Con}_{A \multimap B}$ so as to make $\chi_{A,B}$ a map of esps, i.e. *$(\parallel_{b \in \min(X_B)} X_b) \parallel X_B \in \mathrm{Con}_{A \multimap B}$ iff $X_B \in \mathrm{Con}_B$, $\bigcup_{b \in \min(X_B)} X_b \in \mathrm{Con}_A$, and this union is disjoint.*

Finally, its isomorphism family comprises those order-isomorphisms $\theta : x \cong y$ such that $\chi_{A,B}\, \theta \in \widetilde{A^\perp \parallel B}$. The families $\widetilde{A \multimap B}_+$ and $\widetilde{A \multimap B}_-$ are defined in the same way.

The fact that this indeed defines a closed structure with respect to the tensor relies on the following proposition, which informs directly the *currying* isomorphism.

Proposition 20. *For any $\sigma : \mathcal{S} \to \mathcal{A}^\perp \otimes \mathcal{B}^\perp \parallel \mathcal{C}$, there exists a* unique *$\sigma' : \mathcal{S} \to \mathcal{A}^\perp \parallel \mathcal{B} \multimap \mathcal{C}$* s.t. *$\sigma = (\mathcal{A}^\perp \parallel \chi_{\mathcal{B},\mathcal{C}}) \circ \sigma'$.*

The subtlety here is to determine, if σ maps $s \in S$ to B, which copy of B the event $\sigma' s$ should be in. But this is uniquely determined by well-threadedness: the unique minimal $s_0 \in [s]$ maps to $\sigma s_0 \in \min(C)$, specifying a unique copy.

4) A linear exponential comonad: We already defined the action of ! on negative tcgs in Definition 14, we now define it on morphisms. From $\sigma : \mathcal{S} \to \mathcal{A}^\perp \parallel \mathcal{B}$, we define

$$!\sigma : !\mathcal{S} \to (!\mathcal{A})^\perp \parallel !\mathcal{B}$$

as the obvious map (easily checked to satisfy the conditions for a $\sim$-strategy), with probability valuation given by

$$v_{!S}(\|_{i\in I}\, x_i) = \prod_{i\in I} v_S(x_i)$$

yielding a probabilistic $\sim$-strategy $!\sigma$ from $!\mathcal{A}$ to $!\mathcal{B}$. This construction yields a functor $! : \mathbf{PG} \to \mathbf{PG}$.

By adjoining deterministic $\sim$-strategies corresponding to the standard copycat strategies of AJM games, ! has a comonad structure $(!, \delta, \varepsilon)$ satisfying the Seely axioms [24], turning $\mathbf{PG}$ into a model of ILL. In the next subsection, we show how to interpret PPCF into the Kleisli category $\mathbf{PG}_!$.

5) Interpretation of PPCF: The interpretation of ground types as $\sim$-arenas was given in Section II-E1. It is extended to all types by setting $[\![A \Rightarrow B]\!] = ![\![A]\!] \multimap [\![B]\!]$. As a cartesian closed category, $\mathbf{PG}_!$ supports the interpretation of the simply-typed λ-calculus [20]: as usual, a typed term $\Gamma \vdash M : B$, with $\Gamma = x_1 : A_1, \dots, x_n : A_n$, is interpreted as a morphism:

$$[\![M]\!] : !(\mathop{\&}_{1\le i\le n} [\![A_i]\!]) \overset{\mathbf{PG}}{\nrightarrow} [\![B]\!]$$

For each $\sim$-arena $\mathcal{A}$, there is a *fixpoint combinator* $Y_\mathcal{A}$ on $(!(!\mathcal{A} \multimap \mathcal{A}))^\perp \parallel \mathcal{A}$ allowing us to interpret Y. It is obtained as a supremum of a chain of finite approximations, see [12], [13] for details. It remains to interpret the primitives of PPCF. From $\Gamma \vdash M : \mathbf{Bool}$, $\Gamma \vdash N_1 : \mathbf{Bool}$, $\Gamma \vdash N_2 : \mathbf{Bool}$, we define $[\![\mathbf{if}\, M\, N_1\, N_2]\!]$ via composition with a deterministic $\sim$-strategy $\mathbf{if} : [\![\mathbf{Bool}]\!] \,\&\, [\![\mathbf{Bool}]\!] \,\&\, [\![\mathbf{Bool}]\!] \overset{\mathbf{PG}_!}{\nrightarrow} [\![\mathbf{Bool}]\!]$. There are in fact two possibilities for **if**. As in Figure 1, one is sequential and compatible with the usual interpretation of **if** in game semantics, while the other is the parallel strategy from [12]. We omit the specific diagrams by lack of space, hoping that they are easy to generalize from those of Figure 1. As both $\sim$-strategies will collapse to the same weighted relation, the actual choice does not matter for the results to come. Finally constants are interpreted as in the following examples:

$$[\![\mathtt{tt}]\!] = \begin{array}{c}\mathbf{Bool}\\ \mathrm{q}^{(-,\mathcal{Q})}\\ \curlyvee \\ \mathtt{tt}^{(+,\mathcal{A})}\end{array} \qquad [\![\mathbf{coin}]\!] = \begin{tikzcd}[column sep=small] & \mathbf{Bool} & \\ & \mathrm{q}^{(-,\mathcal{Q})} \arrow[dl] \arrow[dr] & \\ \frac{1}{2}\mathtt{tt}^{(+,\mathcal{A})} \arrow[rr, squiggly, no head] & & \frac{1}{2}\mathtt{ff}^{(+,\mathcal{A})}\end{tikzcd}$$

where we use the notation introduced in Section II-D2 for probabilities – no annotation means that the probability is one for all configurations. From the collapse operation presented in the next section, we will deduce that the interpretation just defined is sound, adequate, and intensionally fully abstract.

IV. Relational Collapse and Full Abstraction

In this section we show how to project games and probabilistic $\sim$-strategies to sets and weighted relations. We aim for a functor $\mathbf{PG}_! \to \mathbf{PRel}_!$ preserving the interpretation of PPCF, making the following diagram commute (up to iso):

$$\begin{tikzcd} & \mathrm{PPCF} \arrow[dl, "{[\![-]\!]_{\mathbf{PG}}}"'] \arrow[dr, "{[\![-]\!]_{\mathbf{PRel}}}"] & \\ \mathbf{PG}_! \arrow[rr, "\downarrow(-)"] & & \mathbf{PRel}_! \end{tikzcd}$$

Our notation $\downarrow(\_)$ emphasises that its purpose is to *flatten out* strategies, that is, forget the causal dependencies and the order in which events (*e.g.* argument calls) occurred.

A. Collapsing games and strategies

1) Mapping arenas to sets: The relational model does not record *function calls*, but only the trace of the data returned at the end of a successful execution. We first extract from games the corresponding configurations: a configuration x of a $\sim$-arena is **complete** if every question in x has an answer in x.

Likewise, the relational model records the *multiplicity* of calls, but not specific copy indices; to replicate that using the information in games we need to consider configurations up to symmetry. The isomorphism family $\widetilde{A}$ induces an equivalence relation on configurations: we say that configurations x and y are *symmetric* if there is a bijection $\theta : x \cong y$ in $\widetilde{A}$. In this case, because θ must preserve the $\mathcal{Q}/\mathcal{A}$ labelling, x is complete iff y is complete. Therefore the following is well-defined:

Definition 21. *Let $\mathcal{A}$ be a $\sim$-arena. Define $\downarrow\mathcal{A}$ to be the set of nonempty, complete configurations of A, up to symmetry.*

Consider for instance the $\sim$-arena $[\![\mathbf{Bool}]\!]_{\mathbf{PG}}$ for booleans. It has two nonempty and complete configurations, $\{q^-, \mathtt{tt}^+\}$ and $\{q^-, \mathtt{ff}^+\}$, and trivial symmetry, so the set $\downarrow[\![\mathbf{Bool}]\!]_{\mathbf{PG}}$ is isomorphic to the two-element set $\{\mathtt{tt}, \mathtt{ff}\} = [\![\mathbf{Bool}]\!]_{\mathbf{PRel}}$.

2) Mapping strategies to matrices: Let $\sigma : \mathcal{S} \to \mathcal{A}$ be a (negative, well-threaded, visible, well-bracketed) probabilistic $\sim$-strategy. Our goal is to define a "vector" $\downarrow\sigma \in \overline{\mathbb{R}}_+^{\downarrow\mathcal{A}}$ indexed by the nonempty and complete *symmetry classes* of configurations of A. We use bold letters $\mathbf{x}, \mathbf{y}, \dots$ to denote the symmetry classes of configurations $x, y, \dots$ respectively.

Given $\mathbf{x} \in \downarrow\mathcal{A}$, to compute the coefficient $(\downarrow\sigma)_\mathbf{x}$ we intuitively count the number of ways of *playing* $\mathbf{x}$ in $\mathcal{S}$, and sum up the probability coefficients for each of them. This is formalised using the notion of *witness*:

Definition 22. *Let $\sigma : \mathcal{S} \to \mathcal{A}$ be a $\sim$-strategy and $x \in \mathcal{C}(A)$. A **witness** for x in σ is $z \in \mathcal{C}(S)$ such that $\sigma z = x$.*

This extends to symmetry classes: if z is a witness for x we say that $\mathbf{z}$ is a witness for $\mathbf{x}$. Then, because σ preserves symmetry, each $z' \in \mathbf{z}$ is a witness for some $x' \in \mathbf{x}$.

Not every complete configuration is relevant: for instance the following $\sim$-strategy on the game $\mathbb{B} \Rightarrow \mathbb{B}$, in which Player calls its argument and returns independently:

$$\begin{tikzcd} \mathbb{B} & \Rightarrow & \mathbb{B} \\ & & \mathrm{q}^{(-,\mathcal{Q})} \arrow[dll] \arrow[d, "\curlyvee" description, no head] \\ \mathrm{q}^{(+,\mathcal{Q})} \arrow[d, "\curlyvee" description, no head] & & \mathtt{tt}^{(+,\mathcal{A})} \\ \mathtt{tt}^{(-,\mathcal{A})} & & \end{tikzcd}$$

Such a behaviour is not definable in PPCF, but this strategy nonetheless exists in our model and needs to be accounted for. When "flattening out" this strategy we must take care not to include $([\mathtt{tt}], \mathtt{tt})$ as a possible execution – this would cause functoriality to fail. In fact when defining $\downarrow\sigma$ we do not consider all configurations of $\mathcal{S}$, but only the **$+$-covered** ones, *i.e.* those whose maximal moves have positive polarity. In a $+$-covered configuration, any argument supplied by Opponent must be used (*i.e.* some Player action must depend on it). Since the property of being $+$-covered is preserved by symmetry, it extends to symmetry classes. We can finally define the action of $\downarrow(\_)$ on strategies.

Definition 23. *Let $\sigma : \mathcal{S} \to \mathcal{A}$ be a (negative, visible, well-bracketed) probabilistic $\sim$-strategy and let $\mathbf{x} \in \downarrow\mathcal{A}$. Define $(\downarrow\sigma)_{\mathbf{x}} = \sum\{v_S(\mathbf{z}) \mid \mathbf{z}$ is a $+$-covered witness for $\mathbf{x}\}$, where $v_S(\mathbf{z})$ is well-defined (v_S is invariant under symmetry).*

B. Functoriality

According to the above a morphism $\sigma : \mathcal{A} \overset{\mathbf{PG}}{\nrightarrow} \mathcal{B}$ collapses to a vector $\downarrow\sigma$ indexed by elements of $\downarrow(\mathcal{A}^\perp \parallel \mathcal{B})$. This is not quite a an element of $\mathbf{PRel}(\downarrow\mathcal{A}, \downarrow\mathcal{B})$, which would instead be indexed by elements of $\downarrow\mathcal{A} \times \downarrow\mathcal{B}$, *i.e.* pairs of nonempty configurations. For $x \parallel y \in \mathcal{C}(A^\perp \parallel B)$ to be nonempty it is enough for only one of x, y to be nonempty. And indeed it is possible for σ to output a value without inspecting its argument: there may be witnesses to $\emptyset \parallel y$ in σ, so the coefficient $(\downarrow\sigma)_{\emptyset\parallel\mathbf{y}}$ may be non-zero. However because $\mathcal{A}, \mathcal{B}$ and σ are *negative*, there can be no witnesses for $x \parallel \emptyset$ in σ, and the coefficient $(\downarrow\sigma)_{\mathbf{x}\parallel\emptyset}$ is always zero.

These observations are consequences of the fact that $\mathbf{PG}$ is *affine*, whereas $\mathbf{PRel}$ is *linear*: a strategy can ignore its argument – and so can a morphism in the Kleisli category $\mathbf{PRel}_!$, but not in $\mathbf{PRel}$. Therefore the target of our collapse functor will not be $\mathbf{PRel}$ but an affine version of it introduced below. Later, moving on to the cartesian closed category $\mathbf{PG}_!$, we will recover the usual relational model $\mathbf{PRel}_!$ of PPCF.

We will first describe the affine version of $\mathbf{PRel}$ and its relationship with $\mathbf{PRel}_!$. After that, we will prove the crucial property that visible $\sim$-strategies have a *deadlock-free interaction*, eventually leading to functoriality of the collapse.

1) The affine relational model: We follow [24, §8.10] and decompose the exponential modality $!$ of $\mathbf{PRel}$ into a weakening modality $\underset{w}{!}$ and a duplication modality $\underset{c}{!}$, each a comonad on $\mathbf{PRel}$. For any set X, $\underset{c}{!}X$ contains its *nonempty* finite multisets: $\underset{c}{!}X = \mathcal{M}_{\mathrm{f}}^{\mathrm{ne}}(X)$, while $\underset{w}{!}X$ has the set X along with the empty multiset: $\underset{w}{!}X = X + \{[\,]\}$. We omit the details of their structure, induced from those of $!$ (found *e.g.* in [17]).

The Kleisli category $\mathbf{PRel}_{\underset{w}{!}}$ is now a model of affine logic, with structure defined in terms of the structure of $\mathbf{PRel}$:
Products: the same as in $\mathbf{PRel}$, $X \mathbin{\&} Y = X + Y$.
Monoidal structure: $X \otimes_w Y = X \otimes Y + X + Y$, with unit $\emptyset$.
Closed structure: $X \multimap_w Y = \underset{w}{!}X \multimap Y$.
Exponential modality: the comonad $\underset{c}{!}$ lifted to $\mathbf{PRel}_{\underset{w}{!}}$.

Lifting the comonad $\underset{c}{!}$ to $\mathbf{PRel}_{\underset{w}{!}}$ exploits a distributive law $\underset{w}{!}\underset{c}{!} \to \underset{c}{!}\underset{w}{!}$, and the Kleisli category $(\mathbf{PRel}_{\underset{w}{!}})_{\underset{c}{!}}$ is isomorphic to $\mathbf{PRel}_!$. With this in place, the collapse will be a functor:

$$\downarrow : \mathbf{PG} \to \mathbf{PRel}_{\underset{w}{!}}$$

preserving the structure required for the interpretation.

We can now define the action of $\downarrow$ on a strategy $\sigma : \mathcal{A} \overset{\mathbf{PG}}{\nrightarrow} \mathcal{B}$: for $\mathbf{x} \in \underset{w}{!}(\downarrow\mathcal{A}), \mathbf{y} \in \downarrow\mathcal{B}$, we set $(\downarrow\sigma)_{[],\mathbf{y}}$ as $(\downarrow\sigma)_{\emptyset\parallel\mathbf{y}}$ and $(\downarrow\sigma)_{\mathbf{x},\mathbf{y}}$ as $(\downarrow\sigma)_{\mathbf{x}\parallel\mathbf{y}}$. We will now check that it is a functor, leaving the preservation of further structure for later.

2) A functor: Consider $\tau : \mathcal{B} \overset{\mathbf{PG}}{\nrightarrow} \mathcal{C}$. To show the functoriality of $\downarrow$ we must relate $\downarrow(\tau \odot \sigma)$ to the Kleisli composition $\downarrow\tau \circ \downarrow\sigma$. For $\mathbf{x} \in \underset{w}{!}\downarrow\mathcal{A}$ and $\mathbf{z} \in \downarrow\mathcal{C}$, the latter is given as:

$$(\downarrow\tau \circ \downarrow\sigma)_{\mathbf{x},\mathbf{z}} = \delta_{\mathbf{x},[\,]}(\downarrow\tau)_{[\,],\mathbf{z}} + \sum_{\mathbf{y}\in\downarrow\mathcal{B}} (\downarrow\sigma)_{\mathbf{x},\mathbf{y}}(\downarrow\tau)_{\mathbf{y},\mathbf{z}},$$

For $\downarrow(\tau \odot \sigma)_{\mathbf{x},\mathbf{z}} = (\downarrow\tau \circ \downarrow\sigma)_{\mathbf{x},\mathbf{z}}$, we use a bijection between:

(1) $+$-covered witnesses $\mathbf{w}$ for $\mathbf{x} \parallel \mathbf{z}$ in $\tau \odot \sigma$, and
(2) Pairs $(\mathbf{w}_S, \mathbf{w}_T)$, where $\mathbf{w}_S$ is a $+$-covered witness for $\mathbf{x} \parallel \mathbf{y}$ in σ, and $\mathbf{w}_T$ for $\mathbf{y} \parallel \mathbf{z}$ in τ, for some $\mathbf{y} \in \underset{w}{!}\downarrow\mathcal{B}$,

satisfying $v_{T\odot S}(\mathbf{w}) = v_S(\mathbf{w}_S) \times v_T(\mathbf{w}_T)$. There are subtleties in both directions, treated separately below. We ignore symmetry in the argument – all steps extend straightforwardly.

a) From (2) to (1): This direction is the most subtle, as it bumps against the reason why traditionally operations from dynamic to static semantics are just lax functorial. Indeed, recall from Lemma 17 that configurations of the interaction $T \circledast S$ correspond to synchronised pairs (w_S, w_T) for which the induced bijection is *secured*. This is in contrast with *(2)*, where witnessed are synchronised with no securedness condition.

The following crucial lemma states that, actually, when composing *visible* strategies, securedness is redundant.

Lemma 24. *Let $x_S \in \mathcal{C}(S)$ and $x_T \in \mathcal{C}(T)$ such that $\sigma\, x_S = x_A \parallel x_B$ and $\tau\, x_T = x_B \parallel x_C$. Then the induced bijection $x_S \parallel x_C \simeq x_A \parallel x_T$ is* secured.

Proof. In Appendix A. □

So, composing *visible* strategies is inherently relational, from which the direction from *(2)* to *(1)* is direct.

b) From (1) to (2): This direction is easier: given a witness w for $x \parallel z$ in $\tau \odot \sigma$, its down-closure $[w] \in \mathcal{C}(T \circledast S)$ satisfies $(\tau \circledast \sigma)[w] = x \parallel y \parallel z$ for some $y \in \mathcal{C}(B)$. It may look like we are done: writing $\Pi_1[w] = w_S \parallel z$ and $\Pi_2[w] = x \parallel w_T$ we obtain a pair (w_S, w_T) of $+$-covered witnesses for $x \parallel y$ and $y \parallel z$. But it remains to check that $y \in \underset{w}{!}\downarrow B$, *i.e.* that it is complete. Well-bracketing ensures this.

Lemma 25. *If $w \in \mathcal{C}(T \odot S)$ is a witness for $x \parallel z$ in the composition of well-bracketed visible strategies σ and τ, where x and z are complete, then the unique $y \in \mathcal{C}(B)$ such that $(\tau \circledast \sigma)[w] = x \parallel y \parallel z$ is also complete.*

c) Summing up: That this is bijective follows from +-coverdness; and the required equality is obtained by summing up on both sides following this bijection. The collapse preserves identities: for any $\sim$-arena $\mathcal{A}$, $\downarrow \mathrm{cc}_{\mathcal{A}}$ is the Kleisli identity $\underset{w}{!}(\downarrow \mathcal{A}) \to (\downarrow \mathcal{A})$ (*i.e.* the counit for $\underset{w}{!}$). Therefore,

Theorem 26. $\downarrow : \mathbf{PG} \to \mathbf{PRel}_{\underset{w}{!}}$ *is a functor.*

3) Preservation of structure: This functor is well-behaved:

Lemma 27. *We have the natural isomorphisms in* $\mathbf{PRel}_{\underset{w}{!}}$*:*

$$\downarrow(\mathcal{A} \,\&\, \mathcal{B}) \cong \downarrow \mathcal{A} \,\&\, \downarrow \mathcal{B}, \quad \downarrow(\mathcal{A} \parallel \mathcal{B}) \cong \downarrow \mathcal{A} \otimes_w \downarrow \mathcal{B}, \quad \downarrow(!\mathcal{A}) \cong \underset{c}{!}(\downarrow \mathcal{A})$$

Moreover, if B has a unique initial move, then additionally $\downarrow(\mathcal{A} \multimap \mathcal{B}) \cong \downarrow \mathcal{A} \multimap_w \downarrow \mathcal{B}$. All associated structural morphisms are also preserved by the collapse.

The collapse also preserves the interpretation of PPCF ground types. Since all $\sim$-arenas for PPCF types have a unique initial move, the interpretation of all types is preserved, so that there is an iso $\theta_A : \downarrow [\![A]\!]_{\mathbf{PG}} \cong [\![A]\!]_{\mathbf{PRel}}$ for any type A.

Because $\downarrow$ takes $\sigma : !\mathcal{A} \overset{\mathbf{PG}}{\nrightarrow} \mathcal{B}$ to $\downarrow \sigma \in \mathbf{PRel}_{\underset{w}{!}}(\downarrow !\mathcal{A}, \downarrow \mathcal{B}) \cong \mathbf{PRel}(\underset{w}{!}\underset{c}{!} \downarrow \mathcal{A}, \downarrow \mathcal{B}) \cong \mathbf{PRel}(! \downarrow \mathcal{A}, \downarrow \mathcal{B})$ we can lift it to $\mathbf{PG}_! \to \mathbf{PRel}_!$. The functor preserves the interpretation of all PPCF primitives, so that:

Theorem 28. *For any PPCF term $\Gamma \vdash M : A$, $\downarrow [\![M]\!]^{\Gamma}_{\mathbf{PG}} = [\![M]\!]^{\Gamma}_{\mathbf{PRel}}$, up to the isomorphism θ induced by Lemma 27.*

For instance, the probabilistic strategy for M_+ from Figure 2 collapses to its relational interpretation, given in Section II-B2.

C. Full abstraction for PPCF

Finally, we import *adequacy* and *intensional full abstraction* from $\mathbf{PRel}_!$ to $\mathbf{PG}_!$. Let $\sigma : S \to \mathbb{B}$ be a probabilistic $\sim$-strategy. Its **probability of convergence** $\mathrm{Pr}(\sigma \to b)$ is $\sum\{v_S(x) \mid x \in \mathcal{C}(S) \text{ and } \exists e \in x \text{ s.t. } \sigma e = b\}$, for any b. Applying Theorem 28 we immediately get:

Theorem 29 (Adequacy)**.** *Let $\vdash M : \mathbf{Bool}$. Then, for $b \in \mathbb{B}$,*

$$\mathrm{Pr}(M \to b) = \mathrm{Pr}([\![M]\!]_{\mathbf{PG}} \to b)$$

In fact, $\mathbf{PG}_!$ is *intensionally fully abstract*. We say that $\sigma : S \to \mathbb{B}$ and $\tau : T \to \mathbb{B}$ are **observationally equivalent at ground type**, written $\sigma \equiv \tau$, when $\mathrm{Pr}(\sigma \to b) = \mathrm{Pr}(\tau \to b)$ for any $b \in \mathbb{B}$. Observe that $\sigma \equiv \tau$ just in case $\downarrow \sigma = \downarrow \tau$. If σ and τ are probabilistic $\sim$-strategies on an arbitrary $\sim$-arena $\mathcal{A}$, σ and τ are **observationally equivalent**, written $\sigma \approxeq \tau$, if $\alpha \odot \Lambda(\sigma) \equiv \alpha \odot \Lambda(\tau)$ for every 'test' morphism $\alpha : \mathcal{A} \nrightarrow \mathbb{B}$.

From Theorems 28, 29, using standard reasoning, we get:

Theorem 30 (Intensional full abstraction)**.** *Let M and N be PPCF terms such that $\Gamma \vdash M : A$ and $\Gamma \vdash N : A$. Then $M \simeq_{ctx} N$ if and only if $[\![M]\!]^{\Gamma}_{\mathbf{PG}_!} \approxeq [\![N]\!]^{\Gamma}_{\mathbf{PG}_!}$.*

In particular, after quotienting the homsets in $\mathbf{PG}_!$ by the relation $\approxeq$, we get a fully abstract model.

References

[1] Samson Abramsky. Game semantics for programming languages (abstract). In *MFCS'97, Bratislava, Slovakia, August 25-29, 1997*.

[2] Samson Abramsky, Kohei Honda, and Guy McCusker. A fully abstract game semantics for general references. In *Proceedings, LICS'98, Indianapolis, Indiana, USA, June 21-24, 1998*, pages 334–344, 1998.

[3] Samson Abramsky, Radha Jagadeesan, and Pasquale Malacaria. Full abstraction for PCF. *Inf. Comput.*, 163(2):409–470, 2000.

[4] Samson Abramsky and Guy McCusker. Linearity, sharing and state: a fully abstract game semantics for idealized algol with active expressions. *Electr. Notes Theor. Comput. Sci.*, 3:2–14, 1996.

[5] Samson Abramsky and Paul-André Melliès. Concurrent games and full completeness. In *LICS'99, Trento, Italy, July 2-5, 1999*, pages 431–442.

[6] Patrick Baillot, Vincent Danos, Thomas Ehrhard, and Laurent Regnier. Believe it or not, AJM games model is a model of classical linear logic. In *LICS '97, Warsaw, Poland, June 29 - July 2, 1997*, pages 68–75.

[7] Patrick Baillot, Vincent Danos, Thomas Ehrhard, and Laurent Regnier. Timeless games. In *Proceedings, CSL '97, Aarhus, Denmark, August 23-29, 1997, Selected Papers*, pages 56–77, 1997.

[8] G. Berry and Pierre-Louis Curien. Sequential algorithms on concrete data structures. *Theor. Comput. Sci.*, 20:265–321, 1982.

[9] Pierre Boudes. Thick subtrees, games and experiments. In *TLCA 2009, Brasilia, Brazil, July 1-3, 2009.*, pages 65–79, 2009.

[10] Simon Castellan, Pierre Clairambault, Silvain Rideau, and Glynn Winskel. Games and strategies as event structures. 2016. Submitted, https://hal.inria.fr/hal-01302713.

[11] Simon Castellan, Pierre Clairambault, and Glynn Winskel. Symmetry in concurrent games. In *CSL-LICS '14, Vienna, 2014*, pages 28:1–28:10.

[12] Simon Castellan, Pierre Clairambault, and Glynn Winskel. The parallel intensionally fully abstract games model of PCF. In *Proceedings, LICS 2015, Kyoto, Japan, July 6-10, 2015*, pages 232–243, 2015.

[13] Simon Castellan, Pierre Clairambault, and Glynn Winskel. Concurrent hyland-ong games. 2016. https://arxiv.org/abs/1409.7542.

[14] Vincent Danos and Thomas Ehrhard. Probabilistic coherence spaces as a model of higher-order probabilistic computation. *Inf. Comput.*, 209(6):966–991, 2011.

[15] Vincent Danos and Russell Harmer. Probabilistic game semantics. *ACM Trans. Comput. Log.*, 3(3):359–382, 2002.

[16] Thomas Ehrhard. The Scott model of linear logic is the extensional collapse of its relational model. *Theor. Comput. Sci.*, 424:20–45, 2012.

[17] Thomas Ehrhard, Christine Tasson, and Michele Pagani. Probabilistic coherence spaces are fully abstract for probabilistic PCF. In *Proceedings, POPL '14, San Diego, USA, January 20-21, 2014*, pages 309–320.

[18] J. M. E. Hyland and C.-H. Luke Ong. On full abstraction for PCF: i, ii, and III. *Inf. Comput.*, 163(2):285–408, 2000.

[19] James Laird. Full abstraction for functional languages with control. In *LICS'97, Warsaw, Poland, June 29 - July 2, 1997*, pages 58–67, 1997.

[20] Joachim Lambek and Philip J Scott. *Introduction to higher-order categorical logic*, volume 7. Cambridge University Press, 1988.

[21] Paul-André Mellies. Asynchronous games 1: A group-theoretic formulation of uniformity. *Manuscript, Available online*, 2003.

[22] Paul-André Melliès. Asynchronous games 4: A fully complete model of propositional linear logic. In *Proceedings, LICS 2005, 26-29 June 2005, Chicago, IL, USA*, pages 386–395, 2005.

[23] Paul-André Melliès. Asynchronous games 2: The true concurrency of innocence. *Theor. Comput. Sci.*, 358(2-3):200–228, 2006.

[24] Paul-André Melliès. Categorical semantics of linear logic. *Panoramas et syntheses*, 27:15–215, 2009.

[25] Gordon D. Plotkin. LCF considered as a programming language. *Theor. Comput. Sci.*, 5(3):223–255, 1977.

[26] Silvain Rideau and Glynn Winskel. Concurrent strategies. In *LICS '11, June 21-24, 2011, Toronto, Canada*, pages 409–418, 2011.

[27] Takeshi Tsukada and C.-H. Luke Ong. Innocent strategies are sheaves over plays - deterministic, non-deterministic and probabilistic innocence. *CoRR*, abs/1409.2764, 2014.

[28] Glynn Winskel. Event structures. In *Advances in Petri Nets*, pages 325–392, 1986.

[29] Glynn Winskel. Event structures with symmetry. *Electr. Notes Theor. Comput. Sci.*, 172:611–652, 2007.

[30] Glynn Winskel. Distributed probabilistic and quantum strategies. *Electr. Notes Theor. Comput. Sci.*, 298:403–425, 2013.

[31] Glynn Winskel. Strategies as profunctors. In *Proceedings, FOSSACS '13, 16-24 March 2013, Rome, Italy*, pages 418–433, 2013.

APPENDIX

In this section, we provide a detailed proof of the deadlock-free lemma (Lemma 24). The key property of visible strategies that we use to prove this result is the following lemma:

Lemma 31. *Let $\sigma : S \to A$ be a visible strategy and let $s < s'$ be events of S. Then the justifier of s' is comparable to s.*

Proof. Since $s < s'$, there exists a gcc ρ of S such that s and s' occur in ρ. By visibility of σ, just(s') occurs in ρ. Since ρ is a total-order, just(s') must be comparable to s. $\square$

We first prove the lemma for *dual* visible strategies, on a game A with only *negative* minimal events. So consider visible $\sigma : S \to A$ (necessarily negative), and $\tau : T \to A^\perp$ (necessarily non-negative). We assume moreover that events in S (resp. T) that map to minimal events of A are minimal.

In such a situation, we have:

Lemma 32. *In a situation as above, for any $x \in \mathcal{C}(S), y \in \mathcal{C}(T)$ such that $\sigma\, x = \tau\, y$, the bijection $\varphi : x \simeq \sigma x = \tau y \simeq y$, induced by local injectivity, is secured.*

Proof. Observe first that because $\sigma s = \tau(\varphi(s))$, it follows that φ preserves justifier: $\varphi(\text{just}(s)) = \text{just}(\varphi s)$. We recall that φ is secured when the relation $(s, t) \lhd_\varphi (s', t')$ defined on graph of φ as $s <_S s'$ or $t <_T t'$ is acyclic. Suppose it is not, and consider a cycle $((s_1, t_1), \ldots, (s_n, t_n))$ with

$$(s_1, t_1) \lhd_\varphi (s_2, t_2) \lhd_\varphi \ldots \lhd_\varphi (s_n, t_n) \lhd_\varphi (s_1, t_1)$$

Let us first give a measure on such cycles. The **length** of a cycle as above is n. For $a \in A$, the **depth** $\text{depth}(a)$ of a is the length of the path to a minimal event of the arena – so the depth of a minimal event is 0. Then, the **depth** of the cycle above is the sum:

$$d = \sum_{1 \le i \le n} \text{depth}(\sigma\, s_i)$$

Cycles are well-ordered by the lexicographic ordering on (n, d); let us now consider a cycle which is minimal for this well-order. Note: in this proof, all arithmetic computations on indices are done modulo n (the length of the cycle).

Since $\le_S$ and $\le_T$ are transitive we can assume that $s_{2k} \le s_{2k+1}$ and $t_{2k+1} \le t_{2k+2}$ for all k. But then it follows by minimality that $\text{pol}_S(s_{2k}) = -$ and $\text{pol}_S(s_{2k+1}) = +$ so that the cycle is alternating. Indeed, assume

$$(s_{2k+1}, t_{2k+1}) \lhd_\varphi (s^+_{2k+2}, t^-_{2k+2}) \lhd_\varphi (s_{2k+3}, t_{2k+3})$$

with $t_{2k+1} \le_T t_{2k+2}$ and $s_{2k+2} \le_S s_{2k+3}$. The causal dependency $t_{2k+1} \le_T t^-_{2k+2}$ decomposes into $t_{2k+1} \le_T t \rightarrowtriangle_T t^-_{2k+2}$, with by courtesy $\tau\, t \rightarrowtriangle_A \tau\, t_{2k+2}$. Note that as A is alternating, this entails that $\text{pol}_T(t) = +$. There must be some $(s, t) \in \varphi$, with $\text{pol}_S(s) = -$. But since $\sigma\, s \le_A \sigma\, s_{2k+2}$, we must have $s \le_S s_{2k+2}$ as well, therefore we can replace the cycle fragment above with

$$(s_{2k+1}, t_{2k+1}) \lhd_\varphi (s^-, t^+) \lhd_\varphi (s_{2k+3}, t_{2k+3})$$

which has the same length but smaller depth, absurd. By the dual reasoning, events with odd index must have polarity as in (s^+_{2k+1}, t^-_{2k+1}) as well.

Now, we remark that the cycle cannot contain events that are minimal in the game. Indeed, by hypothesis a synchronised event (s, t) such that $\sigma\, s = \tau\, t \in A$ is minimal in A is such that $s \in S$ and $t \in T$ are minimal as well, so (s, t) is a root for $\lhd_\varphi$ and cannot be in a cycle. Therefore, all events in the cycle have a predecessor in the game, *i.e.* a justifier.

Since $s_{2k} <_S s_{2k+1}$, by Lemma 31, just(s_{2k+1}) is comparable with s_{2k} in S. They have to be distinct, as otherwise we would have $\sigma s_{2k} \rightarrowtriangle_A \sigma s_{2k+1}$ which in turn implies $t_{2k} <_T t_{2k+1}$. This gives $t_{2k-1} <_T t_{2k+2}$ hence (s_k, t_k) and (s_{k+1}, t_{k+1}) can be removed without breaking the cycle, contradicting its minimality. By a similar reasoning, just(t_{2k+2}) is comparable and distinct from t_{2k+1}.

Assume that we have $s_{2k} < \text{just}(s_{2k+1})$ for some k. Since $\text{just}(s_{2k+1}) < s_{2k+1}$ and $\text{just}(t_{2k+1}) < t_{2k+1} < t_{2k+2}$. Therefore, we can replace the cycle fragment

$$(s_{2k}, t_{2k}) \lhd_\varphi (s_{2k+1}, t_{2k+1}) \lhd_\varphi (s_{2k+2}, t_{2k+2})$$

with the cycle fragment

$$(s_{2k}, t_{2k}) \lhd_\varphi (\text{just}(s_{2k+1}), \text{just}(t_{2k+1})) \lhd_\varphi (s_{2k+2}, t_{2k+2})$$

which has the same length but smaller depth, absurd. So we must have $\text{just}(s_{2k+1}) < s_{2k}$. Similarly, we must have $\text{just}(t_{2k+2}) < t_{2k+1}$ for all k.

So we have that for all k, $\text{just}(s_{2k+1}) < s_{2k}$ with $\text{pol}_S(s_{2k}) = -$. By courtesy and the fact that A is alternating, this has to factor as

$$\text{just}(s_{2k+1}) <_S \text{just}(s_{2k})^+ \rightarrowtriangle_S s^-_{2k}$$

By the dual reasoning, we have that $\text{just}(t_{2k+2}) <_T \text{just}(t_{2k+1})$ (note that $\text{just}(s_{2k+1}) \neq \text{just}(s_{2k})$ and $\text{just}(t_{2k+1}) \neq \text{just}(t_{2k+2})$ as they have different polarities).

So we have proved that we always have $\text{just}(s_{2k+1}) <_S \text{just}(s_{2k})$ and $\text{just}(t_{2k+2}) <_t \text{just}(t_{2k+1})$. That means that we can replace the full cycle

$$(s_1, t_1) \lhd_\varphi (s_2, t_2) \lhd_\varphi \ldots \lhd_\varphi (s_n, t_n) \lhd_\varphi (s_1, t_1)$$

with the cycle

$$\begin{array}{l}(\text{just}(s_1), \text{just}(t_1)) \lhd_\varphi (\text{just}(s_n), \text{just}(t_n)) \lhd_\varphi \\ (\text{just}(s_{n-1}), \text{just}(t_{n-1})) \lhd_\varphi \cdots \lhd_\varphi (\text{just}(s_1), \text{just}(t_1))\end{array}$$

which has the same length but smaller depth, absurd. $\square$

The lemma above is the core of the proof. However, some more bureaucratic reasoning is necessary to reduce Lemma 24, which does not talk of two dual visible strategies on one arena of fixed polarity, to the one above.

Consider $\sigma : S \to A^\perp \parallel B$ and $\tau : T \to B^\perp \parallel C$ which are both visible, well-threaded negative strategies with A, B and C negative arenas. We cannot use transparently the lemma above, because the interaction of σ and τ involves the closed interaction of $\sigma \parallel C^\perp : S \parallel C^\perp \to A^\perp \parallel B \parallel C^\perp$ and

$A \parallel \tau : A \parallel T \to A \parallel B^\perp \parallel C$, and the arena $A \parallel B^\perp \parallel C$ is not negative.

Instead, we will use that the *same* interaction can be replayed in the arena with enriched causality $(A \multimap B) \multimap C$. Remark that as in Definition 19, we have a map:

$$\chi_{A,B,C} : ((A \multimap B) \multimap C) \to A \parallel B^\perp \parallel C$$

Using the fact that σ and τ are well-threaded, these additional causal links in the games are compatible with the interaction:

Lemma 33. *Let $x_S \in \mathcal{C}(S)$ and $x_T \in \mathcal{C}(T)$ such that $\sigma\, x_S = x_A \parallel x_B$ and $\tau\, x_T = x_B \parallel x_C$, and consider the induced bijection (not yet known to be secured):*

$$\varphi : x_S \parallel x_C \simeq x_A \parallel x_T$$

Then, there is $w \in \mathcal{C}((A \multimap B) \multimap C)$ such that $\chi_{A,B,C}\, w = x_A \parallel x_B \parallel x_C$ and the induced bijections:

$$x_S \parallel x_C \simeq w \qquad\qquad x_A \parallel x_T \simeq w$$

are secured.

Proof. By well-threadedness, each $t \in x_T$ mapping to B has a unique minimal causal dependency mapping to C, informing the copy of $A \multimap B$, hence the event of $(A \multimap B) \multimap C$ it should be sent to. Likewise, each $s \in x_S$ has a unique minimal causal dependency $s' \in S$ mapping to B, and there is some synchronisation $((1, s'), (2, t'))$ where t' in turn has a unique minimal causal dependency mapping to C – this informs the event of $(A \multimap B) \multimap C$ that s should be sent to.

Securedness is immediate from the observation that the only immediate causal links added have the form $c \rightarrowtriangle b$ or $b \rightarrowtriangle a$ for a, b, c minimal respectively in A, B, C; in both cases spanning a parallel composition in $S \parallel C$ or $A \parallel T$. ☐

We now need to modify $\sigma \parallel C^\perp$ and $A \parallel \tau$ so that they are dual playing on $((A \multimap B) \multimap C)^\perp$ and $(A \multimap B) \multimap C$ respectively. We do that via the following two pullbacks:

$$\begin{tikzcd}
S' \arrow[r, "\chi_S"] \arrow[d, "\sigma'"'] \arrow[dr, phantom, "\lrcorner", very near start] & S \parallel C^\perp \arrow[d, "\sigma \parallel C^\perp"] \\
((A \multimap B) \multimap C)^\perp \arrow[r, "\chi_{A,B,C}"'] & A^\perp \parallel B \parallel C^\perp
\end{tikzcd}
\qquad
\begin{tikzcd}
T' \arrow[r, "\chi_T"] \arrow[d, "\tau'"] \arrow[dr, phantom, "\lrcorner", very near start] & A \parallel T \arrow[d, "A \parallel \tau"] \\
(A \multimap B) \multimap C \arrow[r, "\chi_{A,B,C}"'] & A \parallel B^\perp \parallel C
\end{tikzcd}$$

One can see $\sigma' : S' \to ((A \multimap B) \multimap C)^\perp$ and $\tau' : T' \to (A \multimap B) \multimap C$ simply as $\sigma \parallel C^\perp$ and $A \parallel \tau$, but with the added causality as in $(A \multimap B) \multimap C$, so that the games C, B, A are opened in that order. We have:

Lemma 34. *So defined, σ and τ satisfy the conditions of Lemma 32,* i.e. *they are visible and events mapping to minimal events of $(A \multimap B) \multimap C$ are minimal.*

Proof. Immediate from standard arguments on the analysis of immediate causality in a pullback, see *e.g.* [10]. ☐

We can finally wrap up:

Lemma 24. *Let $x_S \in \mathcal{C}(S)$ and $x_T \in \mathcal{C}(T)$ such that $\sigma\, x_S = x_A \parallel x_B$ and $\tau\, x_T = x_B \parallel x_C$. Then, the induced bijection*

$$x_S \parallel x_C \simeq x_A \parallel x_T$$

is secured.

Proof. By Lemma 33, we get $w \in \mathcal{C}((A \multimap B) \multimap C)$, and pairing w and $x_S \parallel x_C$ (resp. w and $x_A \parallel x_T$), along with the securedness property from Lemma 33, gives us $x_{S'} \in \mathcal{C}(S')$ (resp. $x_{T'} \in \mathcal{C}(T')$ such that $\sigma'\, x_{S'} = \tau'\, x_{T'}$. By Lemma 32, the induced bijection

$$x_{S'} \simeq x_{T'}$$

is secured. But this entails that the composite bijection

$$x_S \parallel x_C \overset{\chi_S}{\simeq} x_{S'} \simeq x_{T'} \overset{\chi_T}{\simeq} x_A \parallel x_T$$

is secured as well, as the constraints are weaker. ☐