The concurrent game semantics of Probabilistic PCF

Simon Castellan
Imperial College London
London, United Kingdom

Pierre Clairambault
Univ Lyon, CNRS, ENS de Lyon, UCB Lyon 1, LIP
Lyon, France

Hugo Paquet
University of Cambridge
Cambridge, United Kingdom

Glynn Winskel
University of Cambridge
Cambridge, United Kingdom

Abstract

We define a new games model of Probabilistic PCF (PPCF) by enriching *thin concurrent games with symmetry*, recently introduced by Castellan *et al*, with probability. This model supports two interpretations of PPCF, one *sequential* and one *parallel*. We make the case for this model by exploiting the causal structure of probabilistic concurrent strategies. First, we show that the strategies obtained from PPCF programs have a deadlock-free interaction, and therefore deduce that there is an interpretation-preserving functor 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. Finally, we propose a definition of probabilistic innocence and prove a finite definability result, leading to a second (independent) proof of full abstraction.

1 Introduction

What is the right setting for the denotational semantics of probabilistic programs? Numerous proposals exist. Early attemps [27, 18], in the setting of domain theory, involved the *probabilistic powerdomain*, with which it is notoriously difficult to obtain a satisfying cartesian closed category [19]. In 2002, Danos and Harmer [12] showed that making the model more *intensional* offers a much more mathematically tractable development: they construct a fully abstract games model for Probabilistic Algol, an extension of Plotkin's PCF [25] with ground mutable state and probabilistic choice. Later on, Danos and Ehrhard gave a model of Probabilistic PCF (PPCF) in *probabilistic coherence spaces* [11], stemming from work on Linear Logic and quantitative semantics [15], and later proved to be fully abstract [14]. In a different direction, recently Staton *et al* [28, 16] (followed even more recently by Ehrhard *et al* [13]) introduced denotational models for probabilistic programming, with a focus on *continuous distributions*, not previously supported.

This variety of models for a large part extends existing semantics for deterministic programs. However, without probability, *game semantics* [17, 3] has offered a more modular picture, accommodating in a single framework pure fonctional computation along with computational effects such as state [4, 2], control [20], and many others[1], following the well-known research programme pushed by Abramsky [1] under the name of *semantic cube*. Besides this modularity *w.r.t.* the available computational effects in the language, game semantics also offers tools to relate models. For instance, the standard cartesian closed category of Hyland-Ong games and *innocent strategies* embeds functorially in the relational model [6]. Under this *time-forgetting* operation, *points* of the relational model are understood as certain states reached by strategies, without any temporal information.

Of this nice picture however, little remains outside of the deterministic case. It is unclear how to equip Danos and Harmer's model [12] with a notion of probabilistic innocence extending the deterministic one, and how this model relates with alternative, less intensional semantics for probabilistic programs. In fact, even the preliminary question of non-deterministic innocence was unsolved until a few years ago [8, 30], when the important conceptual step was made to switch to a framework expressing *explicit branching* in strategies, representing more intensional behaviour. Adding quantitative information, this suggests the possibility of pushing the *semantic cube* towards probabilistic computation, yielding a valuable tool in our understanding of probabilistic programs.

In this paper, we make an important step in this direction. We draw on recent developments in so-called *concurrent game semantics* [7], a framework for game semantics built around the idea that the *causality* of computation (rather than plain temporal information) is primitive. In particular, we combine the *thin concurrent games with symmetry* [9, 10] of Castellan *et al*, used to build a parallel model of PCF [9], and the *probabilistic concurrent strategies* of [33]. We use this to build a games model of PPCF refining [12].

To support this model, we propose two further contributions. First, we give a quantitative extension of Boudes' theorem [6] and show that our model has a functorial collapse to the $\overline{\mathbb{R}}^+$-*weighted relational model* [21]. This builds on a key lemma independent of probabilities: that the condition of *visibility* from [9] ensures that composition of strategies is *deadlock-free*, and so inherently relational-like (an important precursor for that is Melliès' games model of Linear Logic [23]). As probabilistic coherence spaces embed faithfully in the weighted relational model, it follows by [14] that our model is intensionally fully abstract in the sense of Abramsky *et al* [3]. As a bonus we show that this holds both for a sequential interpretation of PPCF and for a *parallel* one, representing independence of sub-computations. However, definability fails.

Secondly, to get back definability we introduce a notion of *sequential probabilistic innocent strategy*, equivalent to standard innocent

LICS'18, Oxford, UK

978-1-4503-5583-4/18/07…$15.00
DOI: https://doi.org/10.1145/3209108.3209187

[1] This significant achievement led the authors of the seminal papers on game semantics to receive last year's *Alonzo Church Award for Outstanding Contributions to Logic and Computation*, awarded by SIGLOG, EACSL and the Kurt Gödel Society.

strategies in the deterministic case. Sequential probabilistic innocent strategies form a refined model of PPCF for which we prove finite definability (though only *w.r.t.* the sequential interpretation), yielding an independent proof of intensional full abstraction (in fact, unlike previously, *inequational full abstraction* holds).

Related work. Our probabilistic games are related to Tsukada and Ong's sheaf-based notion of innocence [29], though precise connections have not been investigated. That innocent strategies compose relationally is used in Melliès' work on game semantics for linear logic [5, 23], and exploited in Boudes' work on relating games with the relational model — our deadlock-free property generalises it to a non-sequential and non-innocent setting.

Outline. In Section 2 we introduce the semantics of probabilistic programs: we describe PPCF, its relational semantics, and the probabilistic event structures used to represent it in concurrent games. In Section 3 we develop the compositional aspects of the model, and prove the collapse to weighted relations. Finally in Section 4, we prove full abstraction: first as a consequence of the collapse, then (after adding innocence) via definability.

2 Semantics for Probabilistic Programs

2.1 Probabilistic PCF

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

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

The typing rules are standard and omitted – we assume that in **if** $M\ N_1\ N_2$, N_1 and N_2 have ground type (**Bool** or **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: **coin** $\rightarrow^{\frac{1}{2}} b$ for all $b \in \{\mathtt{tt}, \mathtt{ff}\}$. 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 $\prod_{1 \le i \le n} p_i$, and we define the coefficient $\Pr(M \rightarrow N)$ as $\sum\{w(\pi) \mid \pi \text{ is a path from } M \text{ to } N\}$.

Definition 2.1. Let M and N be PPCF terms such that $\Gamma \vdash M : A$ and $\Gamma \vdash N : A$. We write $M \lesssim_{\mathrm{ctx}} N$ if for every context $C[\cdot]$ such that $\vdash C[P]$: **Bool** for every $\Gamma \vdash P : A$,

$$\Pr(C[M] \rightarrow b) \le \Pr(C[N] \rightarrow b)$$

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

2.2 The weighted relational model

In [14], Ehrhard *et al* proved that *probabilistic coherence spaces* (**PCoh**) are **fully abstract** for PPCF: two PPCF terms are contextually equivalent iff they have the same denotation in **PCoh**. In fact, **PCoh** is cut down (via *biorthogonality*) from a more liberal model **PRel**, the $\overline{\mathbb{R}}^+$-*weighted relational model* [21], which we also refer to as the *probabilistic relational model.*

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

$$\begin{array}{ccccc} \mathcal{M}_f(\mathbb{B}) & \times & \mathcal{M}_f(\mathbb{B}) & \times & \mathbb{B} \\ ([\mathtt{tt}^2], & & [\,], & & \mathtt{tt}) \\ ([\mathtt{tt}, \mathtt{ff}], & & [\,], & & \mathtt{ff}) \\ ([\mathtt{ff}], & & [\mathtt{tt}], & & \mathtt{tt}) \\ ([\mathtt{ff}], & & [\mathtt{ff}], & & \mathtt{ff}) \end{array}$$

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.* [22] for an introduction. To construct the target cartesian closed category, we start with one of the simplest models of linear logic: the category **Rel** of sets and relations. In **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 $!X = \mathcal{M}_f(X)$. The cartesian closed category $\mathbf{Rel}_!$ is then the Kleisli category for the comonad !, see *e.g.* [24]. We omit the details of the interpretation of PCF in $\mathbf{Rel}_!$, which we will cover in the presence of probabilities.

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

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

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

The *pure* relational interpretation from before was based on the category **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 one may construct a probabilistic variant of **Rel** by simply replacing the boolean semiring $(\{0,1\}, \vee, \wedge)$ above by the semiring $(\overline{\mathbb{R}}_+, +, \times)$ where $\overline{\mathbb{R}}_+ = \mathbb{R}_+ \uplus \{\infty\}$ denotes the non-negative real numbers, with the infinity added to ensure convergence of the (potentially) infinite sum in the composition formula:

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

for $\varphi \in \mathbb{R}_+^{X \times Y}$, $\psi \in \mathbb{R}_+^{Y \times Z}$.

There is a category **PRel** with sets as objects, and as morphisms from X to Y the 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 **Rel**, **PRel** supports the structure of a model of linear logic with the constructions on objects the same as in **Rel** and analogous constructions on morphisms. We proceed to define the interpretation of PPCF in $\mathbf{PRel}_!$. As for **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 [22]. The interpretation of Y is also obtained in a standard way as a least upper bound of finite approximations, using that homsets of **PRel** are dcpos when ordered componentwise. We now focus on the interpretation of ground types and associated combinators.

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

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

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

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

Finally, the probabilistic primitive **coin** is interpreted as expected as having $[\![\mathbf{coin}]\!]_{\mathtt{tt}} = \frac{1}{2}$ and $[\![\mathbf{coin}]\!]_{\mathtt{ff}} = \frac{1}{2}$, completing the interpretation of PPCF.

In order to avoid infinite weights, the authors of [14] do not stop with **PRel**: they cut down the category using a biorthogonality construction and obtain another weighted model of linear logic, **PCoh**. In **PCoh** weights remain finite, and the interpretation of a term of ground type $M : \mathbb{X}$ yields a sub-probability distribution on $[\![\mathbb{X}]\!]$. In fact, the main result of [14] is that **PCoh** is *fully abstract*, *i.e.* for any M, N we have that $M \simeq_{\mathrm{ctx}} N$ iff $[\![M]\!]_{\mathbf{PCoh}} = [\![N]\!]_{\mathbf{PCoh}}$.

Interestingly this entails that, despite its drawbacks, **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 **PRel** and **PCoh**, the only difference being that the latter is more informative as it carries correctness information *w.r.t.* biorthogonality.

Although its proof is not reproduceable in **PRel**, the main theorem of [14] can be stated as:

Theorem 2.2. *For any terms $\Gamma \vdash M : A$ and $\Gamma \vdash N : A$ of PPCF, $M \simeq_{ctx} N$ iff $[\![M]\!]_{\mathbf{PRel}} = [\![N]\!]_{\mathbf{PRel}}$.*

Accordingly, in the rest of this paper, we will work only in **PRel** and ignore biorthogonality.

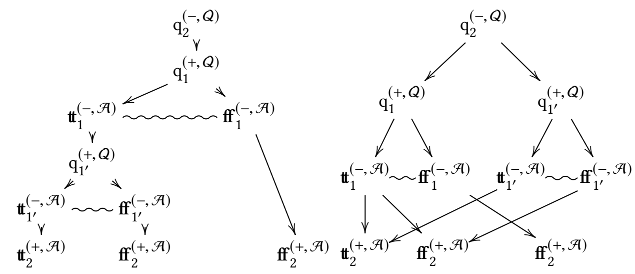

Figure 1. Two strategies for $b : \mathbf{Bool}_1 \vdash M = \mathbf{if}\, b\, b\, \mathtt{ff} : \mathbf{Bool}_2$.

2.3 Game semantics and event structures

The interpretation of a term in **PRel** "flattens out" its behaviour: it only displays the *multiplicity* of its use of resources, but forgets in what *order* these resources are 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 [9]), the term $b : \mathbf{Bool} \vdash M = \mathbf{if}\, b\, b\, \mathtt{ff} : \mathbf{Bool}$ 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. Moves are due to either Player (+) or Opponent (−), as indicated by their polarity, and are annotated by a **Question/Answer** labelling ($\mathcal{Q}/\mathcal{A}$): **questions** correspond to variable calls, whereas **answers** correspond to calls returning. 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 $\mathbf{Bool}_1$)?"
$\mathtt{tt}_1^{(-,\mathcal{A})}$	Opponent:	"The value of b is $\mathtt{tt}$."
$\mathrm{q}_1^{(+,\mathcal{Q})}$	Player:	"Then, what is, again, the value of b?"
$\mathtt{tt}_1^{(-,\mathcal{A})}$	Opponent:	"The value of b is $\mathtt{tt}$."
$\mathtt{tt}_2^{(+,\mathcal{A})}$	Player:	"Then, the output of M is $\mathtt{tt}$."

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 2.3. An **event structure** is $(E, \leq_E, \mathrm{Con}_E)$ with a set E of **events**, $\leq_E$ a partial order stipulating *causal dependency*, and Con_E a non-empty set of **consistent** subsets of E, such that

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

With an eye to game semantics, an **event structure with polarity (esp)** is an event structure E with a function $\mathrm{pol} : E \to \{-, +\}$.

Notations. Write $e \rightarrowtriangle e'$ for **immediate causality**, *i.e.* $e < e'$ with no events in between. Write $C(E)$ for the set of finite **configurations** of E, *i.e.* those finite $x \subseteq E$ such that $x \in \mathrm{Con}$ and x is down-closed, *i.e.* if $e \leq e' \in x$ then $e \in x$. Configurations of the form $[e]$, *i.e.* 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 $x, y \in C(E)$, 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, $X \in \mathrm{Con}$ 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 $e \rightsquigarrow e'$. 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 $\sim\!\sim$.

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 **Bool** $\vdash$ **Bool**, labelling from which the *annotations* $\mathrm{q}_2^{(-,\mathcal{Q})}, \mathtt{tt}_1^{(-,\mathcal{A})}, \ldots$ follow. But let us first discuss how *probability* is adjoined to event structures.

2.4 Event structures with probability

Sequential probabilistic 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 $x \in C(E)$, if x has several distinct extensions $x \cup \{e_1^+\}, x \cup \{e_2^+\} \in C(E)$ with positive events, then $x \cup \{e_1, e_2\} \notin C(E)$. This means that for every $x \in C(E)$, there is a set of positive extensions $\mathrm{ext}_E^+(x)$, all pairwise incompatible.

Sequential esps are easily enriched with probabilities, following the game semantics of Probabilistic Idealized Algol of Danos and Harmer [12]. The basic idea is that for each $x \in C(E)$, Player equips the set of extensions $\mathrm{ext}_E^+(x)$ with a *sub-probability distribution.* But

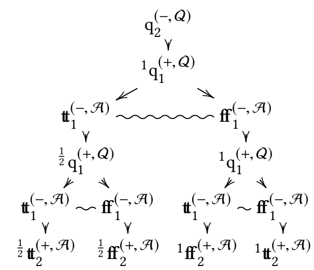

Figure 2. A probabilistic strategy for b : $\mathbf{Bool}_1 \vdash M_+ =$ **if** b (**if coin** b $\bot$) (**if** b **ff** **tt**) : $\mathbf{Bool}_2$

rather than having a sub-distribution for each probabilistic branching in an esp, it is more convenient to carry a single **valuation**

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

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 $x \in C(E)$,

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

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 2.2), with the probabilistic valuation written on the left of events.

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, as in [33]:

Definition 2.4. A **probabilistic esp** consists in an esp $(E, \leq_E, \mathrm{Con}_E, \mathrm{pol}_E)$ and a **valuation** $v : C(E) \to [0, 1]$ satisfying (1), (2) above, plus (3) if $y \subseteq^+ x_1, \ldots, x_n$, then

$$v(y) - \sum_I (-1)^{|I|+1}\, v\left(\bigcup_{i \in I} x_i\right) \geq 0$$

where the sum ranges over $\emptyset \neq I \subseteq \{1, \ldots, n\}$ s.t. $\bigcup_{i \in I} x_i \in C(E)$.

We pointed out in the beginning of Section 2.3 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 an event structure much like 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. Consider for instance how one may assign a valuation v to the esp

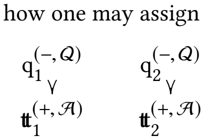

The configurations $\emptyset$, $\{q_1\}$, $\{q_2\}$ and $\{q_1, q_2\}$ necessarily have coefficient 1. Consider then letting $v(\{q_1, \mathtt{tt}_1\}) = v(\{q_1, q_2, \mathtt{tt}_2\}) = \frac{1}{2}$ and $v(\{q_2, \mathtt{tt}_2\}) = v(\{q_1, q_2, \mathtt{tt}_2\}) = \frac{1}{3}$: nothing forces $\mathtt{tt}_1$ and $\mathtt{tt}_2$ to be probabilistically independent events, *i.e.* we may have $v(\{q_1, q_2, \mathtt{tt}_1, \mathtt{tt}_2\}) \neq \frac{1}{6}$. In fact the axioms would allow any value $0 \leq p \leq \frac{1}{3}$. The assignment $v(\{q_1, q_2, \mathtt{tt}_1, \mathtt{tt}_2\}) = \frac{1}{3}$, for example, would indicate a probabilistic *dependence* between $\mathtt{tt}_1$ and $\mathtt{tt}_2$.

2.5 Games and strategies as esps

So far, we have explained the formal nature of the strategies interpreting terms as (probabilistic) esps, but we have not said what *games* they play on.

Arenas. 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 **Bool** and **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:

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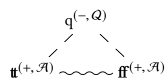 $[\![\mathbf{Nat}]\!] =$ 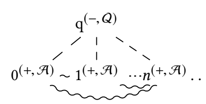

Again, the diagrams are read from top to bottom – immediate causality in arenas is represented by dashed lines rather than arrows, to keep it easily distinguishable from causality in strategies. Although the two notions have the same formal nature, they play a different role in the development.

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 interrogate $\mathbf{Bool}_1$, where polarity is reversed. In fact, in our running examples M and M_+ (from Figures 1 and 2), 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 be interpreted by the arena:

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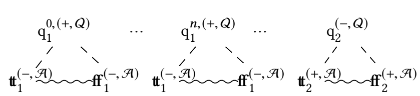

Note the new annotations $q^{i,(+,Q)}$ in copies of the initial question of the argument. This **copy index** i is implicit in the moves $q_1^{(+,Q)}$ in Figures 1 and 2. They will be introduced formally via an *exponential modality*. We now give the general definition of arenas.

Definition 2.5. An **arena** consists of a esp A, and a **labelling function** $\lambda_A : A \to \{Q, \mathcal{A}\}$ 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 $a_1 \rightarrowtriangle a_2$ then $\mathrm{pol}(a_1) \neq \mathrm{pol}(a_2)$.
- *A is race-free:* if $a_1 \sim\!\sim a_2$ then $\mathrm{pol}(a_1) = \mathrm{pol}(a_2)$.
- *Questions:* if a_1 is minimal or if $a_1 \rightarrowtriangle a_2$ then $\lambda_A(a_1) = Q$.
- *Answering is affine:* for every $a_1 \in x \in C(A)$ with $\lambda_A(a_1) = Q$, there is at most one $a_2 \in x$ *s.t.* $a_1 \rightarrowtriangle a_2$ and $\lambda_A(a_2) = \mathcal{A}$.

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

Strategies. Now that we have our notion of games, we can finish making formal the strategies displayed in Figures 1 and 2.

As pointed out earlier, 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 $q_1^{(+,Q)}$, 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 2.6. 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 $x \in C(S)$, the direct image $\sigma\, x \in C(A)$ is a configuration of the game, and *(2)* σ is *locally injective*: for all $s_1, s_2 \in x \in C(S)$, 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** [31] 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 2.7. A pre-strategy $\sigma : S \to A$ is a **strategy** iff it is

- **receptive**: for $x \in C(S)$, if $\sigma x \subseteq^- y \in C(A)$, there is a unique $x \subseteq x' \in C(S)$ such that $\sigma x' = y$; and
- **courteous**: for $s, s' \in S$, if $s \rightarrowtriangle_S s'$ and if $\mathrm{pol}(s) = +$ or $\mathrm{pol}(s') = -$, then $\sigma s \rightarrowtriangle_A \sigma s'$.

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 [7] 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 sufficiently well-behaved 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 first ensure that there is a functorial *collapse* operation to the already fully abstract probabilistic relational model. We will add further conditions in Section 4 to prove definability.

Our conditions are a subset of those of [9]. They crucially rely on the following definition.

Definition 2.8. A **grounded causal chain (gcc)** in an esp S is a set $\rho = \{\rho_1, \ldots, \rho_n\} \subseteq S$ such that ρ_1 is minimal in S and $\rho_1 \rightarrowtriangle_S \rho_2 \rightarrowtriangle_S \rho_3 \rightarrowtriangle_S \ldots \rightarrowtriangle_S \rho_n$. Note that some ρ_i may have dependencies 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:

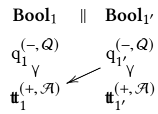

Figure 3. A non-visible strategy on $\mathbf{Bool}_1 \parallel \mathbf{Bool}_{1'}$.

Definition 2.9. A strategy $\sigma : S \to A$ is **visible** iff for all $\rho \in \mathrm{gcc}(S)$, we have $\sigma\,\rho \in C(A)$.

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 $s' \in S$, **its justifier**, such that $\sigma\, s' \rightarrowtriangle_A \sigma\, s$, which we also write to as $\mathrm{just}(s)$.

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 [17]. In our setting however, visibility says that a 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 justifier of $\mathtt{tt}_1$ is absent from the gcc $\mathrm{q}_{1'} \rightarrowtriangle \mathtt{tt}_1$.

Each of these sequential threads needs to respect the call-return discipline, in order to forbid strategies behaving like *e.g.* call/cc [20]. 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 $s_1^{Q} \in X$ 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 [17]:

Definition 2.10. 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 $\{\rho_1 \rightarrowtriangle_S \dots \rightarrowtriangle_S \rho_n\}$.

3 Compositional Structure and Collapse

3.1 A category of games and probabilistic strategies

We start by recalling some basic constructions on esps. Given an esp A, its **dual** is the esp $A^\perp$ whose events, causality and consistency are exactly those of A, but polarity is reversed: $\mathrm{pol}_{A^\perp}(a) = -\mathrm{pol}_A(a)$.

Given a family $(A_i)_{i\in I}$ of esps, we define their **simple parallel composition** to have events

$$\parallel_{i\in I} A_i = \bigcup_{i \in I} \{i\} \times A_i$$

with componentwise *causal ordering* and *polarity*. The *consistent sets* are the finite $\parallel_{i\in I_0} X_i$ for $I_0 \subseteq I$ and $X_i \in \mathrm{Con}_{A_i}$ for all $i \in I_0$.

These constructions extend to arenas with $\lambda_{A^\perp} = \lambda_A$ and $\lambda_{\parallel_{i\in I} A_i}$ defined componentwise. A (probabilistic) **strategy from** A **to** B is a (probabilistic) strategy on $A^\perp \parallel B$. Sometimes we write $\sigma : A \nrightarrow B$ for a strategy $\sigma : S \to A^\perp \parallel B$, keeping the S anonymous.

We now show how to *compose* strategies. As usual in game semantics composition involves two steps: interaction and hiding. We will first show them without probabilities, and then add it back.

Interaction of strategies. Let A, B and C be arenas, and $\sigma : S \to A^\perp \parallel B$ and $\tau : T \to B^\perp \parallel C$ be strategies. Intuitively, *states* of the interaction $\tau \circledast \sigma$ should correspond to so-called *synchronised* 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

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

However this seems inaccurate, because while σ wants to play $\mathtt{tt}_1$ *after* $\mathrm{q}_{1'}$, τ_l will only ask $\mathrm{q}_{1'}$ *after* σ plays $\mathtt{tt}_1$: 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 esps, which we know exists from [26, 7]:

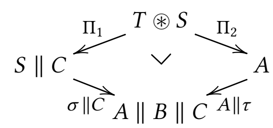

Either path around yields the **interaction** $\tau \circledast \sigma : T \circledast S \to A \parallel B \parallel C$, a *labelled event structure*, charaterised in *e.g.* [7]:

Lemma 3.1. *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 $(s,t) \lessdot_\varphi (s',t')$ if $s \leq s'$ or $t \leq t'$ is a partial order.*

In the interaction of σ and τ_l above, the state $(\{\mathrm{q}_1\}, \{\mathrm{q}_2, \mathrm{q}_1\})$ is maximal. It cannot be extended further, as we have a *deadlock*: strategies are waiting on each other. 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.* [34]. Accordingly our main result will rely on Lemma 3.7, which states that the composition of *visible* strategies is always deadlock-free.

Composition of strategies. Following [26, 7], from $\tau \circledast \sigma : T \circledast S \to A \parallel B \parallel C$, 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 $x \in C(T \odot S)$ has a **unique witness** $[x]_{T\circledast S} \in C(T \circledast S)$. *Polarities* in $T \odot S$ are set so that the restriction $\tau \odot \sigma : T \odot S \to A^\perp \parallel C$ preserves them. From this we get the **composition** of σ and τ, a strategy $\tau \odot \sigma : T \odot S \to A^\perp \parallel C$ [7].

Composition of probabilistic strategies. We turn to the probabilistic case. For the interaction $T \circledast S$, for $x \in C(T \circledast S)$ we set:

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

where $\Pi_1\, x = x_S \parallel x_C$ and $\Pi_2\, x = x_A \parallel x_T$. For $x \in C(T \odot S)$, we set $v_{T\odot S}(x) = v_{T\circledast S}([x]_{T\circledast S})$. From [33], we know that this makes $\tau \odot \sigma$ a probabilistic strategy. We have defined

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

a probabilistic strategy from A to C.

The probabilistic copycat strategy. The identity strategy on an arena A is the **copycat strategy**, $\text{cc}_A : \mathbb{C}_A \to A^\perp \parallel A$. The events, consistent subsets and polarity of $\mathbb{C}_A$ are those of $A^\perp \parallel A$, with causality relation $\leq_{\mathbb{C}_A}$ defined as the transitive closure of

$$\leq_{A^\perp \parallel A} \cup \left\{((1,a),(2,a)) \mid \text{pol}_{A^\perp}(1,a) = -\right\} \cup \left\{((2,a),(1,a)) \mid \text{pol}_A(2,a) = -\right\}.$$

Configurations of $\mathbb{C}_A$ are certain configurations $x_1 \parallel x_2 \in \mathcal{C}(A^\perp \parallel A)$. Being deterministic, copycat is easily made probabilistic by assigning probability 1 to every configuration [33]. Under these definitions the map $\text{cc}_A : \mathbb{C}_A \to A^\perp \parallel A$ is a probabilistic strategy.

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

Arenas, probabilistic strategies, and morphisms between them form a bicategory [26]. We will not use the 2-cells, so in what follows we work in the induced category (obtained by quotienting homsets). We are interested in a subcategory whose morphisms are the visible, well-bracketed strategies of Section 2.5, which are moreover **negative** (*i.e.* S is negative) and **well-threaded** (for all $s \in S$, $[s]$ has exactly one initial move). These additional conditions are needed for the categorical structure presented in the next section.

Definition 3.2. The category **PG** has

- *objects*: negative arenas;
- *morphisms from A to B*: negative, well-threaded, visible and well-bracketed probabilistic strategies, up to isomorphism.

3.2 A symmetric monoidal closed category

Monoidal structure. The **tensor** $A \otimes B$ is simply defined as $A \parallel B$, with unit $\mathbf{1}$ the empty arena. From $\sigma_1 : S_1 \to A_1^\perp \parallel B_1$ and $\sigma_2 : S_2 \to A_2^\perp \parallel B_2$, form $\sigma_1 \otimes \sigma_2 : S_1 \parallel S_2 \to (A_1 \otimes A_2)^\perp \parallel (B_1 \otimes B_2)$, as obvious from $\sigma_1 \parallel \sigma_2$; with $v_{S_1 \otimes S_2}(x_1 \parallel x_2) = v_{S_1}(x_1) \times v_{S_2}(x_2)$. Without probabilities, this yields a symmetric monoidal structure [7]; the extension with probabilities offers no difficulty.

Cartesian structure. The empty arena $\mathbf{1}$ is a terminal object. The **cartesian product** of arenas A and B, written $A \,\&\, B$, has events, causality, and polarity those of $A \parallel B$, and consistent subsets those finite $X = X_A \parallel \emptyset$ with $X_A \in \text{Con}_A$ or $X = \emptyset \parallel X_B$ with $X_B \in \text{Con}_B$. We have two **projections**:

$$\varpi_A : \mathbb{C}_A \to (A \,\&\, B)^\perp \parallel A \qquad \varpi_B : \mathbb{C}_B \to (A \,\&\, B)^\perp \parallel B$$

where one component of the & is not reached — this is compatible with receptivity since A and B are negative. From $\sigma : S \to A^\perp \parallel B$ and $\tau : T \to A^\perp \parallel C$, their **pairing**

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

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 B and C is key in ensuring local injectivity. Compatibility of pairing and projections, along with surjective pairing, are easy verifications.

Closed structure. Because our objects are *negative* arenas, $A^\perp \parallel B$ usually lies outside **PG**. So, inspired by the arrow construction in HO game semantics, we deviate from $A^\perp \parallel B$ by having A depend on $\min(B)$ the minimal events of B. If there are several of them, we copy A accordingly. As our setting is sensitive to linearity, we use consistency to ensure that this copying remains linear.

Definition 3.3. Consider A, B two negative arenas. The arena $A \multimap B$ has as events $(\parallel_{b \in \min(B)} A^\perp) \parallel B$ 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

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

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

One may then check that there is a natural bijection $\mathbf{PG}(A \otimes B, C) \cong \mathbf{PG}(A, B \multimap C)$, *i.e.* **PG** is symetric monoidal closed.

Dcpo-enrichment. To interpret PPCF it will be necessary for **PG** to be *dcpo-enriched.* We equip the set of probabilistic strategies on a game A with a relation $\sqsubseteq$, as follows. For $\sigma : S \to A$ and $\tau : T \to A$ probabilistic strategies, set $\sigma \sqsubseteq \tau$ if $S \sqsubseteq T$, (*i.e.* $S \subseteq T$ and the structure of S is the restriction of that of T), and if moreover $v_S(x) \leq v_T(x)$ for any $x \in \mathcal{C}(S)$. It is clear that $\sqsubseteq$ is a partial order. The least upper bound (lub) of a directed set of probabilistic strategies is their union, with valuation given as $v(x) = \sup \{v_S(x) \mid (\sigma : S \to A) \in D \text{ and } x \in \mathcal{C}(S)\}$. The least element (up to isomorphism) is given by $\bot_A : \min(A) \to A$ (note that the map $\emptyset \to A$ is not receptive in general and so not a strategy).

3.3 Collapsing games and strategies

Though we have yet to introduce a linear exponential comonad on **PG** to break linearity, we find it better to delay its introduction, and give now the collapse of arenas and strategies to sets and relations. Its functoriality will be addressed in the next subsection.

Mapping arenas to sets. Unlike games, **PRel** only records the trace of the data returned by functions for successful executions. In games, the relevant information is captured by the **complete** configurations, *i.e.* those x where every question is answered in x.

Definition 3.4. Let A be an arena. Define $\downarrow A$ to be the set of nonempty, complete configurations of A.

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

Mapping strategies to matrices. Let $\sigma : S \to A$ be a (negative, well-threaded, visible, well-bracketed) probabilistic strategy. Our goal is to define a "vector" $\downarrow \sigma \in \overline{\mathbb{R}}_+^{\downarrow A}$ indexed by the nonempty and complete configurations of A.

Given $x \in \downarrow A$, the coefficient $(\downarrow \sigma)_x$ intuitively sums the probability coefficients of all the ways one can play x in S. This is formalised using the notion of *witness*:

Definition 3.5. Let $\sigma : S \to A$ be a strategy and $x \in \mathcal{C}(A)$. A **witness** for x in σ is $z \in \mathcal{C}(S)$ such that $\sigma z = x$, and such that all maximal moves of z have positive polarity (we say z is **+-covered**). Write $\text{wit}_S(x)$ for the set of all witnesses of x in S.

The requirement that witnesses should not have negative maximal moves is illustrated by the following strategy on the game $\mathbb{B} \multimap \mathbb{B}$, where Player calls its argument and returns independently:

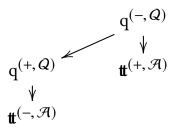

When flattening out this strategy, we must not include (tt, tt) as a possible execution, as this would cause functoriality to fail.

We can finally define the action of $\downarrow(\_)$ on strategies.

Definition 3.6. Let $\sigma : S \to A$ be a (negative, well-threaded, visible, well-bracketed) probabilistic strategy. For $x \in \downarrow A$, we let:

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

3.4 Functoriality of the collapse

Following the above a morphism $\sigma : S \to A^\perp \parallel B$ in **PG** collapses to a vector $\downarrow \sigma$ indexed by elements of $\downarrow(A^\perp \parallel B)$. This is not quite in $\mathbf{PRel}(\downarrow A, \downarrow B)$, which would instead be indexed by elements of $\downarrow A \times \downarrow B$, *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 σ might output a value without inspecting its argument: there may be witnesses to $\emptyset \parallel y$ in σ, so $(\downarrow \sigma)_{\emptyset \parallel y}$ may be non-zero. However because A, B and σ are *negative*, there can be no witnesses for $x \parallel \emptyset$ in σ, and the coefficient $(\downarrow \sigma)_{x \parallel \emptyset}$ is always zero.

These observations follow from **PG** being *affine*, whereas **PRel** is *linear*: a strategy can ignore its argument — and so can a morphism in the Kleisli category $\mathbf{PRel}_!$, but not in **PRel**. Thus the target of our collapse functor will not be **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 first describe the affine version of **PRel** and its relationship with $\mathbf{PRel}_!$. After that, we prove functoriality of the collapse.

The affine relational model. Following [24, §8.10] and decompose the ! of **PRel** into a weakening modality $!_w$ and a duplication modality $!_c$, each a comonad on **PRel**. For any set X, $!_c X$ contains its *nonempty* finite multisets: $!_c X = \mathcal{M}_f^{\mathrm{ne}}(X)$, while $!_w X$ has the set X along with the empty multiset: $!_w X = X + \{[\,]\}$. We omit details of their structure, induced from those of ! (found *e.g.* in [14]).

The Kleisli category $\mathbf{PRel}_{!_w}$ is now a model of affine logic, with structure defined in terms of the structure of **PRel**:

- *Products:* the same as in **PRel**, $X \,\&\, Y = X + Y$.
- *Monoidal structure:* $X \otimes_w Y = X \otimes Y + X + Y$, with unit $\emptyset$.
- *Closed structure:* $X \multimap_w Y = {!_w X} \multimap Y$.
- *Exponential modality:* the comonad $!_c$ lifted to $\mathbf{PRel}_{!_w}$.

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

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

preserving the structure required for the interpretation.

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

A functor. Consider $\tau : T \to B^\perp \parallel C$. To show the functoriality of $\downarrow$ we must relate $\downarrow(\tau \odot \sigma)$ to the Kleisli composition $\downarrow \tau \circ \downarrow \sigma$. For $x \in {!_w} \downarrow A$ and $z \in \downarrow C$, the latter is given as:

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

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

(1) witnesses w for $x \parallel z$ in $\tau \odot \sigma$, and
(2) pairs (w_S, w_T), where w_S is a witness for $x \parallel y$ in σ, and w_T for $y \parallel z$ in τ, for some $y \in {!_w} \downarrow B$,

which satisfies $v_{T \odot S}(w) = v_S(w_S) \times v_T(w_T)$. There are subtleties in both directions, which we proceed to define.

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 only lax functorial. Indeed, recall from Lemma 3.1 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 witnesses are synchronised with no securedness condition. The following crucial lemma states that, when composing *visible* strategies, securedness is redundant.

Lemma 3.7 (Deadlock-free lemma). *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 $\varphi : x_S \parallel x_C \simeq x_A \parallel x_T$ is* secured.

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

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 witnesses for $x \parallel y$ and $y \parallel z$. But it remains to check that $y \in {!_w} \downarrow B$, *i.e.* that it is complete. Well-bracketing ensures this.

Lemma 3.8. *If $w \in \mathrm{wit}_{T \odot S}(x \parallel z)$, for 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.*

Summing up. That this is bijective follows from +-coveredness of the witnesses; and the required equality is obtained by summing up on both sides following this bijection. The collapse preserves identities: for any arena A, $\downarrow \mathfrak{c}_A$ is the Kleisli identity ${!_w}(\downarrow A) \to (\downarrow A)$ (*i.e.* the counit for $!_w$). Therefore,

Theorem 3.9. $\downarrow : \mathbf{PG} \to \mathbf{PRel}_{!_w}$ *is a functor.*

Preservation of structure. This functor is well-behaved. One can easily check that it preserves the order structure on morphisms: if $\sigma \sqsubseteq \tau$ then $\downarrow \sigma \leq \downarrow \tau$, and furthermore $\downarrow(\bigvee_{\sigma \in D} \sigma) = \bigvee_{\sigma \in D}(\downarrow \sigma)$ for any directed set D — so in fact $\downarrow(\_)$ is itself dcpo-enriched. It behaves well also with respect to the categorical structure:

Lemma 3.10. *We have the natural isomorphisms in* $\mathbf{PRel}_{!_w}$*:*

$$\downarrow(A \,\&\, B) \cong \downarrow A \,\&\, \downarrow B \qquad \downarrow(A \parallel B) \cong \downarrow A \otimes_w \downarrow B$$

Moreover, when B has a unique initial move, we have $\downarrow(A \multimap B) \cong \downarrow A \multimap_w \downarrow B$. All associated structural morphisms are preserved.

3.5 Games and strategies with symmetry

In Section 2.5 we hinted at the need for moves to be duplicated, and adjoined *copy indices*. The necessity of expressing uniformity *w.r.t.* copy indices (see [10]) requires us to enrich our probabilistic games with a notion of *symmetry*.

Probabilistic thin concurrent games. *Event structures with symmetry*, introduced in [32], were applied to games in [8] and refined in [9]. For lack of space we only give an informal description.

Our category is a probabilistic enrichment of the *thin concurrent games* of [9]. The objects are **$\sim$-arenas**, consisting of an arena A and (among others) a set $\widetilde{A}$ of bijections $\theta : x \cong y$ between configurations $x, y \in \mathscr{C}(A)$, expressing that x and y are *interchangeable*, *i.e.* the same up to copy indices. This is subject to further axioms [10], and informs an equivalence relation on $\mathscr{C}(A)$. Likewise, **probabilistic $\sim$-strategies** are $\sigma : \mathcal{S} \to \mathcal{A}$ where $\mathcal{S}$ also has an isomorphism family preserved by σ, with the requirement that symmetric configurations should be assigned the same probability.

Unlike **PG**, this category now supports a linear exponential comonad !, whose Kleisli category is, as usual, a ccc:

Lemma 3.11. *There is a cartesian closed category $\mathcal{PG}_!$ having*

- *objects: negative $\sim$-arenas;*
- *morphisms $\mathcal{A}$ to $\mathcal{B}$: (negative, well-threaded, visible, well-bracketed) probabilistic $\sim$-strategies $\sigma : \mathcal{S} \to !\mathcal{A}^\perp \parallel \mathcal{B}$, up to isomorphism and symmetry.*

Interpretation of PPCF. The interpretation of ground types as $\sim$-arenas was given in Section 2.5. It is extended to all types by setting $[\![A \Rightarrow B]\!] = ![\![A]\!] \multimap [\![B]\!]$. As a cartesian closed category, $\mathcal{PG}_!$ supports the interpretation of the simply-typed λ-calculus [22]: 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:

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

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 **if**: $[\![\mathbf{Bool}]\!] \,\&\, [\![\mathbf{Bool}]\!] \,\&\, [\![\mathbf{Bool}]\!] \twoheadrightarrow [\![\mathbf{Bool}]\!]$. 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 [9]. We omit the specific diagrams, hoping that they are easy to generalise from those of Figure 1. We denote the sequential and parallel interpretation by $[\![\_]\!]^s$ and $[\![\_]\!]^p$, respectively, and simply use $[\![\_]\!]$ when the choice does not matter: in particular, both $\sim$-strategies will collapse to the same weighted relation.

Finally constants are interpreted as in the following examples:

$$[\![\mathtt{tt}]\!] = \begin{array}{c} \mathbf{Bool} \\ \mathrm{q}^{(-,\mathcal{Q})} \\ \downarrow \\ \mathtt{tt}^{(+,\mathcal{A})} \end{array} \qquad [\![\mathbf{coin}]\!] = \begin{array}{c} \mathbf{Bool} \\ \mathrm{q}^{(-,\mathcal{Q})} \\ \swarrow \quad \searrow \\ \tfrac{1}{2}\,\mathtt{tt}^{(+,\mathcal{A})} \sim \tfrac{1}{2}\,\mathtt{ff}^{(+,\mathcal{A})} \end{array}$$

where configurations have probability 1 unless specified otherwise. For each $\sim$-arena $\mathcal{A}$, there is a (deterministic) *fixpoint combinator* $Y_{\mathcal{A}}$ on $(!(!\mathcal{A} \multimap \mathcal{A}))^\perp \parallel \mathcal{A}$ allowing us to interpret Y as the lub of a set of approximants, see [9] for details.

Relational collapse. The new subtlety in extending our functor $\downarrow$: $\mathbf{PG} \to \mathbf{PRel}_{\underset{w}{!}}$ from Section 3.4 is that moves in $!\mathcal{A}$ mention specific copy indices, while finite multisets $\mathcal{M}_f(A)$ only count multiplicity. To address that, we refine $\downarrow \mathcal{A}$ as the set of *$\cong$-equivalence classes* of non-empty and complete configurations of A (and similarly for $\downarrow \sigma$). The developments of Sections 3.3 and 3.4 adapt smoothly to the new framework, and we now have $\downarrow (!\mathcal{A}) \cong \mathcal{M}_f^{\mathrm{ne}}(\downarrow \mathcal{A})$.

Thus, $\downarrow$ takes $\sigma : !\mathcal{A} \twoheadrightarrow \mathcal{B}$ to $\downarrow \sigma$ in $\mathbf{PRel}_{\underset{w}{!}}(\downarrow !\mathcal{A}, \downarrow \mathcal{B})$, which is iso to $\mathbf{PRel}(!_{wc}^{\,!} \downarrow \mathcal{A}, \downarrow \mathcal{B}) \cong \mathbf{PRel}(! \downarrow \mathcal{A}, \downarrow \mathcal{B})$. Hence we can lift it:

Lemma 3.12. *There is a functor $\downarrow : \mathcal{PG}_! \to \mathbf{PRel}_!$.*

It is a straightforward verification that there is an isomorphism $\theta_A : \downarrow [\![A]\!]_{\mathbf{PG}} \cong [\![A]\!]_{\mathbf{PRel}}$ for any type A of PPCF. Moreover the functor preserves the interpretation of all PPCF primitives, so that:

Theorem 3.13. *For any PPCF term $\Gamma \vdash M : A$, up to the isomorphism $\theta_{\Gamma \vdash A}$, we have that $\downarrow [\![\Gamma \vdash M]\!]^s_{\mathcal{PG}} = \downarrow [\![\Gamma \vdash M]\!]^p_{\mathcal{PG}} = [\![\Gamma \vdash M]\!]_{\mathbf{PRel}}$.*

For instance, the probabilistic strategy for M_+ from Figure 2 collapses to its relational interpretation, given in Section 2.2.

Interestingly, the equational theory on PPCF induced by the parallel interpretation is strictly finer than that induced by $\mathbf{PRel}_!$.

4 Full Abstraction for PPCF

4.1 Full abstraction in $\mathcal{PG}_!$ by relational collapse

We import *adequacy* and *intensional full abstraction* from $\mathbf{PRel}_!$ to $\mathcal{PG}_!$ using the functor $\downarrow$. Let $\sigma : \mathcal{S} \to \mathbb{B}$ be a probabilistic $\sim$-strategy. Its **probability of convergence to** $b \in \{\mathtt{tt}, \mathtt{ff}\}$, written $\Pr(\sigma \to b)$, is $\sum_{\substack{x \in \mathscr{C}(S) \\ \text{s.t. } b \in \sigma x}} v_S(x)$. Applying Theorem 3.13 we get:

Theorem 4.1 (Adequacy)**.** *Let $\vdash M : \mathbf{Bool}$. Then, for $b \in \mathbb{B}$,*

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

In fact, $\mathcal{PG}_!$ is *intensionally fully abstract*, that is, contextual equivalence in the language coincides with contextual equivalence in the model. Let us now formally define the latter, by means of a contextual preorder. Note the similarity with Definition 2.1.

We start by defining a preorder $\leqq$ on ground type strategies: given $\sigma : S \to \mathbb{B}$ and $\tau : T \to \mathbb{B}$, write $\sigma \leqq \tau$ whenever $\Pr(\sigma \to b) \le \Pr(\tau \to b)$ for any $b \in \{\mathtt{tt}, \mathtt{ff}\}$. Observe that, writing $\equiv$ for the equivalence induced by $\leqq$, we have $\sigma \equiv \tau$ just in case $\downarrow \sigma = \downarrow \tau$.

Definition 4.2. If σ and τ are probabilistic $\sim$-strategies on an arbitrary $\sim$-arena $\mathcal{A}$, write $\sigma \lesssim_{\mathrm{ctx}} \tau$, if $\alpha \odot \Lambda(\sigma) \leqq \alpha \odot \Lambda(\tau)$ for every 'test' morphism $\alpha : \mathcal{A} \twoheadrightarrow \mathbb{B}$. The induced **contextual equivalence** is denoted $\simeq_{\mathrm{ctx}}$.

Theorems 3.13 and 4.1 imply full abstraction:

Theorem 4.3 (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 $[\![\Gamma \vdash M]\!]_{\mathcal{PG}} \simeq_{ctx} [\![\Gamma \vdash N]\!]_{\mathcal{PG}}$ (where $[\![\_]\!]$ is either $[\![\_]\!]^s$ or $[\![\_]\!]^p$).*

Visible and well-bracketed probabilistic $\sim$-strategies have thus precisely the same distinguishing power as PPCF contexts. But the model still contains "junk", *i.e.* $\sim$-strategies which do not behave like PPCF terms. In this section we impose a further condition on $\sim$-strategies (*sequential innocence*, defined in Section 4.2) in order to prove a finite definability result (Theorem 4.7). From there, a fully abstract model for PPCF follows using standard reasoning.

In what follows we simply use *strategies* to refer to the morphisms of $\mathcal{PG}$, *i.e.* the negative, well-threaded, visible, and well-bracketed probabilistic $\sim$-strategies, considered up to isomorphism.

4.2 Full abstraction by definability

In this paper we are only concerned with definability with respect to the sequential interpretation $[\![\_]\!]^s$ of PPCF.

Definition 4.4. A strategy $\sigma : \mathcal{S} \to \mathcal{A}$ is **sequential innocent** if

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

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

The first condition is necessary for definability as configurations with probability zero are not definable in PPCF.

Sequential strategies form a well-behaved class: they are stable under composition, and copycat is sequential innocent. Call $\mathcal{PG}^{\mathrm{si}}$ the subcategory of $\mathcal{PG}$ whose morphisms are (isomorphism classes of) sequential innocent strategies. We can use it to interpret PPCF:

Lemma 4.5. *For any PPCF term* $\Gamma \vdash M : A$, $[\![\Gamma \vdash M]\!]^s_{\mathrm{PG}}$ *is a sequential innocent strategy.*

So just like $\mathcal{PG}_!$, the category $\mathcal{PG}^{\mathrm{si}}_!$ provides an adequate model of PPCF. But it is a much smaller category, allowing us to prove intensional full abstraction via definability. Fix a $\sim$-arena $\mathcal{A} = [\![A]\!]$, for A some arbitrary PPCF type. As usual, *finite* definability will be sufficient for full abstraction:

Definition 4.6. A sequential innocent strategy $\sigma : \mathcal{S} \to \mathcal{A}$ is **finite** when:

- There is a bound to the length of gccs,
- For every $s^- \in S$, the set $\{t^+ \in S \mid s \rightarrowtriangle_S t\}$ is finite;
- $v(x) \in \mathbb{Q} \cap [0, 1]$ for every $x \in \mathcal{C}(S)$.

It is necessary for definability that configuration-valuations of finite strategies have rational coefficients (because of non-computable elements in $[0, 1]$), but $\mathbb{Q}$ being dense in $\mathbb{R}$, any configuration-valuation can be approximated by ones with rational values, and indeed finite strategies form a basis for the dcpo of innocent sequential strategies on $\mathcal{A}$. Finite innocent sequential strategies have an inductive tree structure, that we exploit for PPCF definability.

Theorem 4.7 (Finite definability)**.** *For any finite* $\sigma : \mathcal{S} \to \mathcal{A}$ *in* $\mathcal{PG}^{si}_!$*, there is a PPCF term* $\vdash M : A$ *such that* $[\![M]\!]^s_{\mathcal{PG}} = \sigma$.

From here, deriving a fully abstract model is standard. We write $\lesssim^{\mathrm{si}}_{\mathrm{ctx}}$ to denote the contextual preorder in $\mathcal{PG}^{\mathrm{si}}_!$ defined by requiring the α of Definition 4.2 to be sequential innocent. We show:

Theorem 4.8. *Let* M, N *be PPCF terms such that* $\Gamma \vdash M : A$ *and* $\Gamma \vdash N : A$*. Then,* $M \lesssim_{ctx} N$ *iff* $[\![\Gamma \vdash M]\!]^s_{\mathcal{PG}} \lesssim^{si}_{ctx} [\![\Gamma \vdash N]\!]^s_{\mathcal{PG}}$.

Note that full abstraction holds in its stronger *inequational* form. Definability permits this while the relational collapse did not: inequational full abstraction does not hold in $\mathbf{PRel}_!$ [14].

5 Conclusion

In future work, we aim to rely on this to push further the *quantitative semantic cube*, studying interactions of probabilities with state and concurrency. The challenge, there, is to understand how probabilistic choice interacts with the nondeterminism of scheduling.

Acknowledgements. This work was supported by the LABEX MILYON (ANR-10-LABX-0070) of Université de Lyon, within the program "Investissements d'Avenir" (ANR-11-IDEX-0007) operated by the French National Research Agency (ANR). We also gratefully acknowledge the support of an EPSRC PhD studentship, ERC Advanced Grant ECSYM, EPSRC grants EP/K034413/1 and EP/K011715/1, and the Collegium de Lyon.

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. 1998. A Fully Abstract Game Semantics for General References. In *Proceedings, LICS'98, Indianapolis, Indiana, USA, June 21-24, 1998.* 334–344.
[3] Samson Abramsky, Radha Jagadeesan, and Pasquale Malacaria. 2000. Full Abstraction for PCF. *Inf. Comput.* 163, 2 (2000), 409–470.
[4] Samson Abramsky and Guy McCusker. 1996. Linearity, Sharing and State: a fully abstract game semantics for Idealized Algol with active expressions. *Electr. Notes Theor. Comput. Sci.* 3 (1996), 2–14.
[5] Samson Abramsky and Paul-André Melliès. Concurrent Games and Full Completeness. In *LICS'99, Trento, Italy, July 2-5, 1999.* 431–442.
[] 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.* 68–75.
[6] Pierre Boudes. 2009. Thick Subtrees, Games and Experiments. In *TLCA 2009, Brasilia, Brazil, July 1-3, 2009.* 65–79.
[7] Simon Castellan, Pierre Clairambault, Silvain Rideau, and Glynn Winskel. 2017. Games and Strategies as Event Structures. *LMCS* 13, 3 (2017).
[8] Simon Castellan, Pierre Clairambault, and Glynn Winskel. Symmetry in concurrent games. In *CSL-LICS '14, Vienna, 2014.* 28:1–28:10.
[9] Simon Castellan, Pierre Clairambault, and Glynn Winskel. 2015. The Parallel Intensionally Fully Abstract Games Model of PCF. In *Proceedings, LICS 2015, Kyoto, Japan, July 6-10, 2015.* 232–243.
[10] Simon Castellan, Pierre Clairambault, and Glynn Winskel. 2016. Concurrent Hyland-Ong Games. (2016). https://arxiv.org/abs/1409.7542.
[11] Vincent Danos and Thomas Ehrhard. 2011. Probabilistic coherence spaces as a model of higher-order probabilistic computation. *Inf. Comput.* 209, 6 (2011), 966–991. https://doi.org/10.1016/j.ic.2011.02.001
[12] Vincent Danos and Russell Harmer. 2002. Probabilistic game semantics. *ACM Trans. Comput. Log.* 3, 3 (2002), 359–382.
[13] Thomas Ehrhard, Michele Pagani, and Christine Tasson. 2018. Measurable cones and stable, measurable functions: a model for probabilistic higher-order programming. *PACMPL* 2, POPL, 59:1–59:28.
[14] 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.* 309–320.
[15] Jean-Yves Girard. 1988. Normal functors, power series and λ-calculus. *Annals of pure and applied logic* 37, 2 (1988), 129–177.
[16] Chris Heunen, Ohad Kammar, Sam Staton, and Hongseok Yang. 2017. A convenient category for higher-order probability theory. In *LICS'17, Reykjavik.* IEEE Computer Society, 1–12. https://doi.org/10.1109/LICS.2017.8005137
[17] J. M. E. Hyland and C.-H. Luke Ong. 2000. On Full Abstraction for PCF: I, II, and III. *Inf. Comput.* 163, 2 (2000), 285–408.
[18] Claire Jones and Gordon D. Plotkin. 1989. A Probabilistic Powerdomain of Evaluations. In *(LICS '89), Pacific Grove, California, USA.* IEEE Computer Society, 186–195. https://doi.org/10.1109/LICS.1989.39173
[19] Achim Jung and Regina Tix. 1998. The troublesome probabilistic powerdomain. *Electr. Notes Theor. Comput. Sci.* 13 (1998), 70–91.
[20] James Laird. 1997. Full Abstraction for Functional Languages with Control. In *LICS'97, Warsaw, Poland, June 29 - July 2, 1997.* 58–67.
[21] Jim Laird, Giulio Manzonetto, Guy McCusker, and Michele Pagani. Weighted relational models of typed lambda-calculi. In *LICS'13.*
[22] Joachim Lambek and Philip J Scott. 1988. *Introduction to higher-order categorical logic.* Vol. 7. Cambridge University Press.
[23] Paul-André Melliès. 2005. Asynchronous Games 4: A Fully Complete Model of Propositional Linear Logic. In *Proceedings, LICS 2005, Chicago, IL, USA.* 386–395.
[] Paul-André Mellies. 2003. Asynchronous games 1: A group-theoretic formulation of uniformity. *Manuscript, Available online* (2003).
[24] Paul-André Melliès. 2009. Categorical semantics of linear logic. *Panoramas et syntheses* 27 (2009), 15–215.
[25] Gordon D. Plotkin. 1977. LCF Considered as a Programming Language. *Theor. Comput. Sci.* 5, 3 (1977), 223–255.
[26] Silvain Rideau and Glynn Winskel. 2011. Concurrent Strategies. In *LICS '11, June 21-24, 2011, Toronto, Canada.* 409–418.
[27] N. Saheb-Djahromi. 1980. CPO'S of Measures for Nondeterminism. *Theor. Comput. Sci.* 12 (1980), 19–37. https://doi.org/10.1016/0304-3975(80)90003-1
[28] Sam Staton, Hongseok Yang, Frank D. Wood, Chris Heunen, and Ohad Kammar. 2016. Semantics for probabilistic programming: higher-order functions, continuous distributions, and soft constraints. In *Proceedings of LICS '16, New York, NY, USA, July 5-8, 2016.* 525–534.
[29] Takeshi Tsukada and C.-H. Luke Ong. 2014. Innocent Strategies are Sheaves over Plays - Deterministic, Non-deterministic and Probabilistic Innocence. *CoRR* abs/1409.2764 (2014). http://arxiv.org/abs/1409.2764
[30] Takeshi Tsukada and C.-H. Luke Ong. 2015. Nondeterminism in Game Semantics via Sheaves. In *30th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2015, Kyoto, Japan, July 6-10, 2015.* 220–231.
[31] Glynn Winskel. 1986. Event Structures. In *Advances in Petri Nets.* 325–392.

[32] Glynn Winskel. 2007. Event Structures with Symmetry. *Electr. Notes Theor. Comput. Sci.* 172 (2007), 611–652.
[33] Glynn Winskel. 2013. Distributed Probabilistic and Quantum Strategies. *Electr. Notes Theor. Comput. Sci.* 298 (2013), 403–425.
[34] Glynn Winskel. 2013. Strategies as Profunctors. In *Proceedings, FOSSACS '13, 16-24 March 2013, Rome, Italy.* 418–433.

A Games with Symmetry

A.1 Symmetry in event structures

We first review the basics of *event structures with symmetry* [32], presented here as in [9] via *isomorphism families.*

Definition A.1. An **isomorphism family** on an event structure E is a set $\widetilde{E}$ of bijections $\theta : x \cong y$, where $x, y \in \mathscr{C}(E)$, *s.t.*

(1) For all $x \in \mathscr{C}(E)$, $\mathrm{id}_x : x \cong x \in \widetilde{E}$.
(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 $x \subseteq x' \in \mathscr{C}(E)$, then there exists $y \subseteq y' \in \mathscr{C}(E)$ and $\theta' : x' \cong y' \in \widetilde{E}$ such that $\theta \subseteq \theta'$.
(5) If $\theta : x \cong y \in \widetilde{E}$ and $x' \subseteq x \in \mathscr{C}(E)$, then there exists $y' \subseteq y \in \mathscr{C}(E)$ 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 A.2. 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.

A.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 A.3. Given a family $\mathcal{A}_i, i \in I$ of essps, their **parallel composition** $\|_{i\in I}\ \mathcal{A}_i$ is $\|_{i\in I}\ A_i$ equipped with the isomorphism family $\|_{i\in I}\ \widetilde{A_i}$, with bijections $\theta : \|_{i\in I_0}\ x_i \cong \|_{i\in I_0}\ y_i$ induced by a family $(\theta_i : x_i \cong_{\widetilde{A_i}} y_i)_{i\in I_0}$ such that for all $(i, a) \in \|_{i\in I_0}\ x_i$, $\theta((i, a)) = (i, \theta_i\, a_i)$.

Definition A.4. 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 [8]. Note that this ! operation is not the same as the one used in [9] and which duplicates all moves of the game "*in depth*" rather than just at the surface – in the spirit of HO games [17]. 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.* [10] for a discussion. One solution, used in [8], is to ask that all strategies are *saturated*, and play non-deterministically all possible copy indices. Another, introduced in [9] and detailed in [10], requires instead that strategies pick copy indices deterministically (are *thin*, see Definition A.6). 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 [**?**] by bi-invariance under the action of two groups of Opponent and Player permutations.

Definition A.5. 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 $\theta = \mathrm{id}_x$ for some $x \in \mathscr{C}(A)$,
(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* [**?** 8].

We now add probability to the uniform strategies of [9, 10], called $\sim$*-strategies.*

Definition A.6. 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 $\widetilde{A}_-$ for now) such that $\mathcal{S}$ 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.* $x' = x \cup \{s\} \in \mathscr{C}(S)$ with $\mathrm{pol}(s) = +$, there is a *unique* $t \in S$ *s.t.* $\theta \cup \{(s, t)\} \in \widetilde{S}$.

The remaining concepts of Section 2.3 extend in the presence of symmetry: a $\sim$**-arena** is a tcg $\mathcal{A}$ with a $\boldsymbol{Q}/\boldsymbol{\mathcal{A}}$ 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).

A.3 A category of probabilistic $\sim$-strategies

The **parallel composition** of tcgs $\mathcal{A} = (A, \widetilde{A}, \widetilde{A}_-, \widetilde{A}_+)$ and $\mathcal{B} = (B, \widetilde{B}, \widetilde{B}_-, \widetilde{B}_+)$ is $\mathcal{A} \parallel \mathcal{B} = (A \parallel B, \widetilde{A} \parallel \widetilde{B}, \widetilde{A}_- \parallel \widetilde{B}_-, \widetilde{A}_+ \parallel \widetilde{B}_-)$. The dual of $\mathcal{A}$ is the tcg $\mathcal{A}^\perp = (A^\perp, \widetilde{A}, \widetilde{A}_+, \widetilde{A}_-)$. As usual, a **probabilistic $\sim$-strategy from $\mathcal{A}$ to $\mathcal{B}$** is one on $\mathcal{A}^\perp \parallel \mathcal{B}$.

A.3.1 Composition and copycat

Given $\sim$-strategies $\sigma : \mathcal{S} \to \mathcal{A}^\perp \parallel \mathcal{B}$ and $\tau : \mathcal{T} \to \mathcal{B}^\perp \parallel \mathcal{C}$, their **interaction** $\mathcal{T} \circledast \mathcal{S}$ is (just like its underlying event structure $T \circledast S$) defined as the pullback

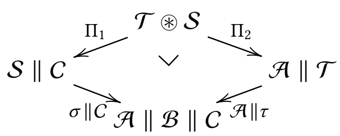

only this time in the category of event structures *with symmetry.*

To define the **composition** of σ and τ we equip the event structure $T \odot S$ (obtained from $T \circledast S$ after *hiding*) with the isomorphism family $\widetilde{T \odot S}$, 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 $\tau \odot \sigma : \mathcal{T} \odot \mathcal{S} \to \mathcal{A}^\perp \parallel \mathcal{C}$ [10].

When σ and τ are probabilistic $\sim$-strategies, the symmetry does not affect the definition of $v_{\widetilde{T \odot S}}$, which is easily shown to be invariant under the bijections in $\widetilde{T \odot S}$, making $\tau \odot \sigma$ a probabilistic $\sim$-strategy.

Finally, the copycat strategy on a $\sim$-arena $\mathcal{A}$ is also equipped with symmetry: the isomorphism family $\mathbb{C}_{\widetilde{A}}$ 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.

We can form a bicategory of tcgs, probabilistic $\sim$-strategies, and morphisms (where a morphism of $\sim$-strategies is one between the underlying strategies which additionally preserves symmetry). But 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.

A.3.2 Weak isomorphism

Definition A.7. 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 [9, 10] is that weak isomorphism is preserved under composition, which crucially depends on the thinness axiom for $\sim$-strategies.

The conditions on strategies introduced in Sections 2 and 3 do not rely on the affine nature of **PG**. They extend directly to the framework with symmetry, so that:

Definition A.8. There is a category $\mathcal{PG}$ having

- objects: negative $\sim$-arenas;
- morphisms: negative, well-threaded, visible and well-bracketed probabilistic $\sim$-strategies on $\mathcal{A}^\perp \parallel \mathcal{B}$.

A.3.3 A model of intuitionistic linear logic

The category $\mathcal{PG}$ is symmetric monoidal closed and cartesian, with all structure induced from that of **PG** in the obvious way (see [10] for details). In this setting however, we can define a linear exponential comonad !.

Given a $\sim$-arena $\mathcal{A}$, the essp $!\mathcal{A}$ was defined earlier, in Definition A.4. We now define the positive and negative isomorphism families. When $\mathcal{A}$ is a negative $\sim$-arena, 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 $\theta_i \in \widetilde{A}_+$.

The action of ! on morphisms is as follows: 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 $! : \mathcal{PG} \to \mathcal{PG}$.

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 $\mathcal{PG}$ into a model of ILL.

B Omitted Proofs

B.1 Proof of the deadlock-free lemma (Lemma 3.7)

The key property of visible strategies that we use to prove this result is the following lemma:

Lemma B.1. *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. □

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 B.2. *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** $\mathrm{depth}(a)$ 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:

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

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 $\leq_S$ and $\leq_T$ are transitive we can assume that $s_{2k} \leq s_{2k+1}$ and $t_{2k+1} \leq t_{2k+2}$ for all k. But then it follows by minimality that $\mathrm{pol}_S(s_{2k}) = -$ and $\mathrm{pol}_S(s_{2k+1}) = +$ 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 $t_{2k+1} \leq_T t_{2k+2}$ and $s_{2k+2} \leq_S s_{2k+3}$. The causal dependency $t_{2k+1} \leq_T t^-_{2k+2}$ decomposes into $t_{2k+1} \leq_T t \rightarrowtriangle_T t^-_{2k+2}$, with by courtesy $\tau\, t \rightarrowtriangle_A \tau\, t_{2k+2}$. Note that as A is alternating, this entails that $\mathrm{pol}_T(t) = +$. There must be some $(s, t) \in \varphi$, with $\mathrm{pol}_S(s) = -$. But since $\sigma\, s \leq_A \sigma\, s_{2k+2}$, we must have $s \leq_S s_{2k+2}$ 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 4.3, $\mathrm{just}(s_{2k+1})$ 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, $\mathrm{just}(t_{2k+2})$ is comparable and distinct from t_{2k+1}.

Assume that we have $s_{2k} < \mathrm{just}(s_{2k+1})$ for some k. Since $\mathrm{just}(s_{2k+1}) < s_{2k+1}$ and $\mathrm{just}(t_{2k+1}) < t_{2k+1} < t_{2k+2}$. 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

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

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

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

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

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

So we have proved that we always have $\mathrm{just}(s_{2k+1}) <_S \mathrm{just}(s_{2k})$ and $\mathrm{just}(t_{2k+2}) <_t \mathrm{just}(t_{2k+1})$. 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

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

which has the same length but smaller depth, absurd. □

The lemma above is the core of the proof. However, some more bureaucratic reasoning is necessary to reduce Lemma 3.7, 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 3.3, 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 B.3. *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:

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

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:

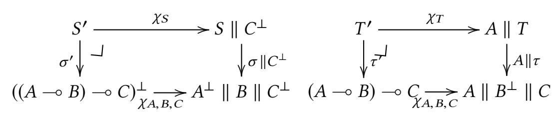

One can see $\sigma' : S' \to ((A \multimap B) \multimap C)^\perp$ and $\tau' : T' \to (A \multimap B) \multimap C$ simply as $\sigma \parallel C^\perp$ and $A \parallel \tau$, but with the added causality as in $(A \multimap B) \multimap C$, so that the games C, B, A are opened in that order. We have:

Lemma B.4. *So defined, σ and τ satisfy the conditions of Lemma 4.3,* 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.* [7]. □

We can finally wrap up:

Lemma 3.7. *Let $x_S \in \mathscr{C}(S)$ and $x_T \in \mathscr{C}(T)$ 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 4.3, we get $w \in \mathscr{C}((A \multimap B) \multimap C)$, and pairing w and $x_S \parallel x_C$ (resp. w and $x_A \parallel x_T$), along with the securedness property from Lemma 4.3, gives us $x_{S'} \in \mathscr{C}(S')$ (resp. $x_{T'} \in \mathscr{C}(T')$ such that $\sigma'\, x_{S'} = \tau'\, x_{T'}$. By Lemma 4.3, 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. □

B.2 Proof of Lemma 3.8

Lemma 3.8. *If $w \in \mathscr{C}(T \odot S)$ 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 $y \in \mathscr{C}(B)$ such that $(\tau \circledast \sigma)[w] = x \parallel y \parallel z$ is also complete.*

Proof. If y is empty, then $\mathrm{y} = [] \in\, !_w \downarrow \mathcal{B}$ is the witness. Otherwise, we must show that y is complete, so that its symmetry class $\mathrm{y} \in \downarrow \mathcal{B}$. So let $q \in [w]$ be a question mapped to y by $\tau \circledast \sigma$. Let e be a visible event (we call *visible* events those of w) such that $q < e$, chosen so that no event between q and e is visible. Then, by courtesy, $\mathrm{pol}(e) = +$ since there is a causal link to e in w which is not in the game. By assumption, maximal events of w are visible positive answers, so there is a gcc

$$\rho : \cdots \rightarrowtriangle q \rightarrowtriangle \cdots \rightarrowtriangle e^+ \rightarrowtriangle \cdots \rightarrowtriangle a^+.$$

We claim that ρ and a can be chosen in such a way that $\rho_{[e,a]}$ contains no visible negative questions. In order to show this we give a construction process, which necessarily terminates since all gccs are finite. Start the process at $\rho_i = e$. If e is a positive answer then we are done. If it is a positive question, then by receptivity there is a visible answer a'^- such that $e \rightarrowtriangle a'$, so let $\rho_{i+1} = a'$ and continue. If it is a negative answer then it is not maximal by assumption (w is +-covered), so continue down any gcc; the next move is either a positive visible event, and we can apply the steps above, or it is a hidden event of $[w]$, in which case we continue down any gcc until reaching a visible event (recall that there are no hidden maximal moves) and repeat the procedure.

So we have defined $\rho : \cdots \rightarrowtriangle q \rightarrowtriangle \cdots \rightarrowtriangle e^+ \rightarrowtriangle \cdots \rightarrowtriangle a^+$. By visibility, a^+ points to some negative q' in ρ, which is necessarily visible since $(\tau \circ \sigma)a$ and $(\tau \circledast \sigma)q'$ are in the same component. Therefore, by construction of ρ, q' must occur before q in ρ. By well-bracketing of τ and σ (which implies the well-bracketing of gccs in $T \circledast S$), all questions of $\rho_{[q',a^+]}$ must be answered in $\rho_{[q',a^+]}$, including q. So in particular q has an answer in w, and y is complete. □

B.3 Proof of intensional full abstraction via collapse (Theorem 4.3)

Theorem 4.3 (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 $[\![\Gamma \vdash M]\!]_{\mathcal{PG}} \simeq_{ctx} [\![\Gamma \vdash N]\!]_{\mathcal{PG}}$.*

Proof. (Only if). By the full abstraction result in $\mathbf{PRel}_!$, $M \simeq N$ implies $[\![\Gamma \vdash M]\!]_{\mathbf{PRel}} = [\![\Gamma \vdash N]\!]_{\mathbf{PRel}}$, which by Theorem 3.13 is the same as saying that $\downarrow [\![\Gamma \vdash M]\!]_{\mathcal{PG}} = \downarrow [\![\Gamma \vdash N]\!]_{\mathcal{PG}}$. Suppose there exists $\alpha : ([\![\Gamma \Rightarrow A]\!]) \twoheadrightarrow \mathbb{B}$ such that $\alpha \odot \Lambda([\![\Gamma \vdash M]\!]_{\mathcal{PG}}) \not\equiv \alpha \odot \Lambda([\![\Gamma \vdash N]\!]_{\mathcal{PG}})$. This implies in particular that $\downarrow(\alpha \odot \Lambda([\![\Gamma \vdash M]\!]_{\mathcal{PG}})) \neq \downarrow(\alpha \odot \Lambda([\![\Gamma \vdash N]\!]_{\mathcal{PG}}))$. Because $\downarrow$ is a structure-preserving functor, this is equivalent to $\downarrow \alpha \odot \Lambda(\downarrow [\![\Gamma \vdash M]\!]_{\mathcal{PG}}) \neq \downarrow \alpha \odot \Lambda(\downarrow [\![\Gamma \vdash N]\!]_{\mathcal{PG}})$, a contradiction since $\downarrow [\![\Gamma \vdash M]\!]_{\mathcal{PG}} = \downarrow [\![\Gamma \vdash N]\!]_{\mathcal{PG}}$. So no such α can exist, and $[\![\Gamma \vdash M]\!]_{\mathcal{PG}} \simeq_{\mathrm{ctx}} [\![\Gamma \vdash N]\!]_{\mathcal{PG}}$.

(If). Suppose now that $[\![\Gamma \vdash M]\!]_{\mathcal{PG}} \simeq_{\mathrm{ctx}} [\![\Gamma \vdash N]\!]_{\mathcal{PG}}$. Let $C[\cdot]$ be a context such that $C[M]$ and $C[N]$ are closed terms of type **Bool**. Then $[\![C[\cdot]]\!]_{\mathcal{PG}}$ is a probabilistic $\sim$-strategy $([\![\Gamma]\!] \Rightarrow [\![A]\!]) \twoheadrightarrow \mathbb{B}$, and therefore $[\![C[M]]\!]_{\mathcal{PG}} \equiv [\![C[N]]\!]_{\mathcal{PG}}$ since $[\![\Gamma \vdash M]\!]_{\mathcal{PG}}$ and $[\![\Gamma \vdash N]\!]_{\mathcal{PG}}$ are observationally equivalent. By adequacy (Theorem 4.1), we have $\Pr(C[M] \to \underline{b}) = \Pr(C[N] \to \underline{b})$ for all b. So $M \simeq N$. □

C Comparing equational theories

In this final section, we compare the different equational theories induced on terms of PPCF by the sequential interpretation, the parallel interpretation, and the interpretation in **PRel**.

If $\Gamma \vdash M, N : A$ are terms of PPCF, we introduce three notions of equivalences between them. We write $M \equiv^s_{\mathcal{PG}} N$ iff $[\![M]\!]^s_{\mathcal{PG}} = [\![N]\!]^s_{\mathcal{PG}}$, likewise we write $M \equiv^p_{\mathcal{PG}} N$ iff $[\![M]\!]^p_{\mathcal{PG}} = [\![N]\!]^p_{\mathcal{PG}}$, and finally $M \equiv_{\mathbf{PRel}} N$ iff $[\![M]\!]_{\mathbf{PRel}} = [\![N]\!]_{\mathbf{PRel}}$. The three induced equational theories on terms of PPCF are ordered as follows:

$$\equiv_{\mathbf{PRel}}$$

$$\equiv^s_{\mathcal{PG}} \hookrightarrow \equiv_{\mathbf{PRel}} \hookleftarrow \equiv^p_{\mathcal{PG}}$$

where we emphasize that the inclusions are strict, and that $\equiv^s_{\mathcal{PG}}$ and $\equiv^p_{\mathcal{PG}}$ are incomparable. The non-strict inclusions are immediate consequences of the collapse functor.

$\equiv^s_{\mathcal{PG}} \subset \equiv_{\mathbf{PRel}}$. Observe the two following terms.

$$M_1 = \textbf{if } x \textbf{ then } \textbf{if } y \textbf{ then } \mathtt{tt} \textbf{ else } \bot \textbf{ else } \bot$$

$$M_2 = \textbf{if } y \textbf{ then } \textbf{if } x \textbf{ then } \mathtt{tt} \textbf{ else } \bot \textbf{ else } \bot$$

These are equal in $\equiv_{\mathbf{PRel}}$ (and $\equiv^p_{\mathcal{PG}}$), but not in $\equiv^s_{\mathcal{PG}}$, where we observe the evaluation order.

$\equiv^p_{\mathcal{PG}} \subset \equiv_{\mathbf{PRel}}$. Technically it suffices to observe that the two terms **if coin then tt else tt** and **tt** have a distinct interpretation as $\mathcal{PG}$ will remember the nondeterministic branching and represent the former with two conflicting events. However, although it is not introduced in the paper, there is a simple equivalence (called "rigid image equivalence") that eliminates idempotent probabilistic choice and would assimilate these two in $\mathcal{PG}$ as well. So more convincingly, we propose the more robust example below.

$$M_3 = \textbf{if coin then } \textbf{if } x \textbf{ then } \mathtt{tt} \textbf{ else } \mathtt{ff} \textbf{ else } \textbf{if } x \textbf{ then } \mathtt{ff} \textbf{ else } \mathtt{tt}$$

$$M_4 = \textbf{if } x \textbf{ then coin else coin}$$

Those two terms are equal in **PRel**, but are distinguished by $\equiv^p_{\mathbf{PRel}}$ where for M_3 we observe two immediate parallel calls to x, in contrast with M_4 where there is only one.

$\equiv^p_{\mathcal{PG}} \not\subseteq \equiv^s_{\mathcal{PG}}$. Indeed, $M_1 \equiv^p_{\mathcal{PG}} M_2$ but $M_1 \not\equiv^s_{\mathcal{PG}} M_2$.

$\equiv^s_{\mathcal{PG}} \not\subseteq \equiv^p_{\mathcal{PG}}$. This comes from $\equiv^p_{\mathcal{PG}}$ not being "sensible": it observes part of the computation that will never be used. More precisely, the term

$$M_5 = \textbf{if } \bot \textbf{ then } x \textbf{ else } \bot$$

satisfies $M_5 \equiv^s_{\mathcal{PG}} \bot$, but $M_5 \not\equiv^p_{\mathcal{PG}} \bot$: following the parallel interpretation we see the call to x performed "speculatively", though of course it will yield no result.

D Deadlock-free Lemma

In this section, we provide a detailed proof of the deadlock-free lemma (Lemma 3.7). The key property of visible strategies that we use to prove this result is the following lemma:

Lemma 4.3. *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. □

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 4.3. *In a situation as above, for any $x \in \mathscr{C}(S), y \in \mathscr{C}(T)$ 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 $((s_1,t_1),\dots,(s_n,t_n))$ with

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

Let us first give a measure on such cycles. The **length** of a cycle as above is n. For $a \in A$, the **depth** 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 4.3, 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

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

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

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

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

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

So we have proved that we always have $\mathrm{just}(s_{2k+1}) <_S \mathrm{just}(s_{2k})$ and $\mathrm{just}(t_{2k+2}) <_t \mathrm{just}(t_{2k+1})$. 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

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

which has the same length but smaller depth, absurd. □

The lemma above is the core of the proof. However, some more bureaucratic reasoning is necessary to reduce Lemma 3.7, 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 3.3, 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 4.3. *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:

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

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:

$$\begin{array}{ccc} S' & \xrightarrow{\chi_S} & S \parallel C^\perp \\ {\scriptstyle \sigma'}\downarrow & \lrcorner & \downarrow{\scriptstyle \sigma \parallel C^\perp} \\ ((A \multimap B) \multimap C)^\perp & \xrightarrow[\chi_{A,B,C}]{} & A^\perp \parallel B \parallel C^\perp \end{array} \qquad \begin{array}{ccc} T' & \xrightarrow{\chi_T} & A \parallel T \\ {\scriptstyle \tau'}\downarrow & \lrcorner & \downarrow{\scriptstyle A \parallel \tau} \\ (A \multimap B) \multimap C & \xrightarrow[\chi_{A,B,C}]{} & A \parallel B^\perp \parallel C \end{array}$$

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 4.3. *So defined, σ and τ satisfy the conditions of Lemma 4.3,* 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.* [7]. □

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 4.3, 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 4.3, 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 4.3, 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. □