What's in a game?

A theory of game models

Clovis Eberhart

Univ. Savoie Mont Blanc, CNRS, LAMA
Chambéry F-73000, France
clovis.eberhart@univ-smb.fr

Tom Hirschowitz

Univ. Savoie Mont Blanc, CNRS, LAMA
Chambéry F-73000, France
tom.hirschowitz@univ-smb.fr

Abstract

Game semantics is a rich and successful class of denotational models for programming languages. Most game models feature a rather intuitive setup, yet surprisingly difficult proofs of such basic results as associativity of composition of strategies. We set out to unify these models into a basic abstract framework for game semantics, *game settings*. Our main contribution is the generic construction, for any game setting, of a category of games and strategies. Furthermore, we extend the framework to deal with innocence, and prove that innocent strategies form a subcategory. We finally show that our constructions cover many concrete cases, mainly among the early models [6, 20] and the very recent sheaf-based ones [37].

1 Introduction

Game semantics has provided adequate models for a variety of (idealised) programming languages. We will here mainly be concerned with the numerous variations on *arena games*. This comprises, e.g., the original dialogue game model of PCF [20, 33], Abramsky et al.'s model for general references [5], Harmer and McCusker's model for finite nondeterminism [15, 17], Laird's model for control operators [24], and the recent model by Tsukada and Ong [37]. We will also briefly consider other models of PCF [6] and of linear logic [10].

This rich literature shares many features. E.g., all these models follow the same simple conceptual route: the types of the considered language are interpreted as games and programs as strategies. Games form the objects, and strategies the morphisms of a category, which is compared to the 'syntactic' category generated from the operational description of the language. However, as noted, e.g., in Harmer et al. [16], a less advantageous feature shared by all models is the surprising difficulty of certain proofs like associativity of composition or the fact that innocent strategies are closed under composition.

This raises the issue of unifying all these models into a satisfactory theory, with an emphasis on factoring out such difficult proofs. This is an ambitious goal, because although game models clearly share a lot of ideas, they are also rather diverse. E.g., depending on the considered language, various constraints are imposed upon strategies, like *innocence* or *well-bracketing*. Further sources of diversity have appeared with recent extensions, e.g., sheaf-based innocence [37], nominal models [32], tensorial logic [30], or concurrent extensions [18, 34].

This paper is an attempt at improving the situation. Focussing on the construction of game models, our main contributions are:

- We organise the basic data underlying typical game models into a simple categorical structure called a *game setting*,

emphasising its simplicial features. Each game setting gives rise to a candidate category of games and strategies.

- We delineate mild hypotheses under which games and strategies do form a category, heavily relying on high-level categorical techniques, including (presheaf) polynomial functors [23, Chapter 16] and Guitart's *exact squares* [14].
- We extend the framework to deal with *innocence*, an emblematic constraint put on strategies to capture purely functional computation. We enrich game settings with a notion of *view* and, under mild hypotheses, we derive a category of innocent strategies. Our approach exploits the recent recasting of innocence as a sheaf condition [18, 37], and again relies on advanced category theory to give high-level proofs.
- We show that a number of game models fall into our framework, namely variants of the original Hyland-Ong/Nickau (HON) model [15, 20, 26], AJM games [6] and Tsukada and Ong's model [37].
- We work out the limits of our techniques in two well-known dead ends of game semantics: non-associativity of composition in Blass games [4] and non-stability of innocent strategies under composition in the absence of determinism [15].
- In passing, we prove a categorical result of possibly independent interest, stating that certain commuting squares of categories and functors, called *local pushforward squares*, are *distributive*, in the sense that doing left then right (Kan) extension along one side is isomorphic to doing right then left along the other.

Our framework deals with various notions of composition and innocence. For clarity, let us readily fix some terminology. An important distinction is whether plays are considered as a poset (with the prefix ordering) or as a category. Another is whether the considered strategies are general or boolean presheaves. We annotate composition and innocence with the following codes.

Plays		Strategies	
poset	category	boolean	general
p	c	b	s

Example 1.1. Standard strategies, being prefix-closed sets of plays, are boolean presheaves on the prefix ordering, so their composition is pb-composition. Moreover, standard innocent strategies are innocent pb-strategies. Similarly, Tsukada and Ong [37] use proper categories of plays and their strategies are presheaves, hence their composition and innocent strategies are cs-composition and innocent cs-strategies. Unspecified items denote either possibility. E.g., s-composition means composition of presheaves, in either the poset-based or category-based setting.

Much effort is put into linking the different variants together, as summarised in the following table.

Section	relates	to
2.8	s-composition	b-composition
4.2	cs-innocence	ps-innocence
4.3	cs-innocence	pb-innocence

The established links are of various nature. E.g., in Section 2.8, we both infer associativity of b-composition from associativity of s-composition, and express b-composition in terms of s-composition.

1.1 Related work

Although significant work has been devoted to giving efficient proofs in particular models [16, 28, 30], the only attempts at abstraction we are aware of are our Hirschowitz [18] and Eberhart and Hirschowitz [12]. Both papers focus on the link between naive and innocent strategies as well as the interpretation of programs, but both fail to capture composition of strategies.

1.2 Plan

In Section 2, we gradually introduce game settings, following the successive steps for constructing a typical game model. We also state our main results for the basic setup along the way. We remain very informal about game semantics, and only start to consider the particulars of various game models in Section 3, where we establish that the announced game models fit into our framework. We then refine game settings to deal with innocence in Section 4, covering in passing Tsukada and Ong’s model [37]. Finally, we conclude in Section 5. Most proofs are deferred to appendices.

1.3 Notation and prerequisites

For all $n \in \mathbb{N}$, $[n]$ denotes the finite ordinal with n elements, i.e., the set $\{0, \dots, n-1\}$, and we sometimes use just n to denote the set $\{1, \dots, n\}$.

We assume some basic knowledge of category theory, namely categories, functors and natural transformations, as well as adjunctions. The category of *presheaves* over any category $\mathcal{C}$ is the functor category $[\mathcal{C}^{op}, \text{Set}]$ of contravariant functors to sets and natural transformations between them, which we denote by $\widehat{\mathcal{C}}$. For any presheaf $X: \mathcal{C}^{op} \rightarrow \text{Set}$, objects $c, c' \in \mathcal{C}$, morphism $f: c' \rightarrow c$, and element $x \in X(c)$, we use a right action notation $x \cdot f$ for $X(f)(x)$. By functoriality, we have $x \cdot f \cdot g = x \cdot (f \circ g)$, for any $g: c'' \rightarrow c'$. Replacing Set with $\mathbf{2}$, the ordinal 2 viewed as a category, we get the category $\widehat{\mathcal{C}}$ of *boolean presheaves*.

2 Game settings

2.1 Categories of plays in game semantics

In this section, we sketch several notions of play typically involved in the construction of a game model. We do this without referring to any particular model. In the next sections, we will organise this data into a coherent categorical structure, which we will then exploit to give an abstract construction of a category of games and strategies.

The construction of a typical game model relies on the definition of increasingly complex notions of play. There is first a notion of *game*. Each game A involves two players O (Opponent) and P (Proponent), and features in particular a set of *plays* $\mathbb{P}_A$, which may be endowed with the prefix ordering or with a more sophisticated notion of morphism, thus forming a category of plays. Such two-player games form the basis of the model.

The crucial step to view strategies as morphisms is to consider the *arrow game* $A \rightarrow B$, which intuitively describes the interaction of a middle player M playing Opponent against a left player L and Proponent against a right player R , as in

$$\begin{array}{ccccc}
 L & & M & & R \\
 & & \mathbb{B} & \xrightarrow{\quad} & \mathbb{B} \\
 & & & & q^R \\
 & q^M & & & \\
 & t^L & & & f^M.
 \end{array} \tag{1}$$

In this example, M plays like the negation function on booleans: R asks its return value by playing the move q^R ; M in turn asks L for the value of the argument by playing q^M , to which L answers ‘true’ by playing t^L ; M eventually answers the original question by playing f^M .

However, there is a subtlety: one often needs to restrict plays in $\mathbb{P}_{A \rightarrow B}$ with additional constraints, so that the relevant category is a subcategory $\mathbb{P}_{A,B} \hookrightarrow \mathbb{P}_{A \rightarrow B}$. Of course, there are projections to $\mathbb{P}_A$ and $\mathbb{P}_B$.

Example 2.1. We will provide more precise definitions later on, but for now, to fix intuition, in HON-style games (without bracketing) $\mathbb{P}_A$ would consist of all justified sequences, and $\mathbb{P}_{A,B}$ would restrict to alternating sequences of even length of $\mathbb{P}_{A \rightarrow B}$. The projections of a play in $\mathbb{P}_{A,B}$ to A and B may not be alternating, so it is crucial to be liberal in the choice of $\mathbb{P}_A$.

In order to define composition of strategies, the situation (1) is then scaled up to combinations of two such situations in which a first middle player M_1 plays on the right with a second one, say M_2 , as in

$$\begin{array}{ccccccc}
 L & & M_1 & & M_2 & & R \\
 & & \mathbb{B} & \xrightarrow{\quad} & \mathbb{B} & \xrightarrow{\quad} & \mathbb{B}.
 \end{array} \tag{2}$$

Plays in such combinations are standardly called *interaction sequences*, and typically form a subcategory $\mathbb{P}_{A,B,C} \hookrightarrow \mathbb{P}_{(A \rightarrow B) \rightarrow C}$. An important point is that interaction sequences admit projections to $\mathbb{P}_{A,B}$, $\mathbb{P}_{B,C}$ and $\mathbb{P}_{A,C}$, which satisfy the obvious equations w.r.t. further projections to $\mathbb{P}_A$, $\mathbb{P}_B$ and $\mathbb{P}_C$, e.g., the following square commutes:

$$\begin{array}{ccc}
 \mathbb{P}_{A,B,C} & \xrightarrow{\quad} & \mathbb{P}_{A,C} \\
 \downarrow & & \downarrow \\
 \mathbb{P}_{B,C} & \xrightarrow{\quad} & \mathbb{P}_C.
 \end{array}$$

Example 2.2. In HON games, $\mathbb{P}_{A,B,C}$ typically consists of alternating justified sequences on $(A \rightarrow B) \rightarrow C$ which end in A or C and whose projections to $A \rightarrow B$ and $B \rightarrow C$ are plays.

Finally, in order to prove associativity of composition, one defines *generalised interaction sequences* as a subcategory $\mathbb{P}_{A,B,C,D} \hookrightarrow \mathbb{P}_{((A \rightarrow B) \rightarrow C) \rightarrow D}$, again with projections satisfying the obvious equations.

2.2 Plays as a category-valued presheaf

Let us now organise all this data $(\mathbb{P}_A, \mathbb{P}_{A,B}, \mathbb{P}_{A,B,C}, \mathbb{P}_{A,B,C,D})$ into a simple categorical structure. First, as suggested by our notation,

What's in a game?

for all lists $L = (A_0, \dots, A_{n-1})$ of games, we may construct a category $\mathbb{P}_L$.

Example 2.3. In the HON case, we may take $\mathbb{P}_L$ to consist of alternating justified sequences on $(\dots (A_0 \rightarrow A_2) \rightarrow \dots) \rightarrow A_{n-1}$ whose projection to each $A_i \rightarrow A_{i+1}$ is a play, and which end in A_0 or A_{n-1} , the rightmost arena.

For the same reasons as before, we get projections $\delta_k: \mathbb{P}_L \rightarrow \mathbb{P}_{L \setminus A_k}$ (for ‘delete k ’), for all $k \in [n]$. A similar construction, relevant for defining identity strategies (so-called *copycat* strategies), is *insertions* $\iota_k: \mathbb{P}_L \rightarrow \mathbb{P}_{L+k}$ (for ‘insert k ’), where $k \in [n]$ and $L+k$ denotes L with the k th game duplicated. E.g., $\iota_1: \mathbb{P}_{A,B} \rightarrow \mathbb{P}_{A,B,B}$. Intuitively, this functor maps any play u in $\mathbb{P}_{A,B}$ to the interaction sequence in $\mathbb{P}_{A,B,B}$ which duplicates all moves on B . So in a situation like (2), M_2 would act as a ‘proxy’ between M_1 and R , repeating M_1 ’s moves to R and conversely. For a precise definition and an example in the case of HON games, see Section 3.1.

This may all be packed up into the comma category $\Delta/\mathbb{A}$, or more precisely $i/\ulcorner \mathbb{A} \urcorner$, where

- $\ulcorner \mathbb{A} \urcorner: 1 \rightarrow \text{Set}$ is the functor picking $\mathbb{A}$;
- $i: \Delta \hookrightarrow \text{Set}$ is the embedding of the simplicial category Δ into sets.

Let us recall that Δ has finite ordinals $[n]$ as objects, with monotone maps as morphisms. So concretely, $\Delta/\mathbb{A}$ has finite lists of games as objects, i.e., maps $L: [n] \rightarrow \mathbb{A}$ for some $n = \{0, \dots, n-1\}$, and as morphisms $(n, L) \rightarrow (n', L')$ all monotone maps f making the following triangle commute:

$$\begin{array}{ccc} [n] & \xrightarrow{f} & [n'] \\ & \searrow L \quad \swarrow L' & \\ & \mathbb{A} & \end{array}$$

Example 2.4. Let $d_k^n: [n] \rightarrow [n+1]$ miss $k \in [n]$, i.e., $d_k^n(i) = i$ for $i < k$ and $d_k^n(j) = j+1$ for $j \geq k$. Then, e.g., d_1^2 yields a map $(A, C) \rightarrow (A, B, C)$ for all games A, B, C . Similarly, consider $i_k^n: [n+1] \rightarrow [n]$ which collapses $k \in [n] \subseteq [n+1]$ and $k+1 \in [n+1]$. E.g., for $n = 2$ and $k = 0$, it yields a map $(A, A, B) \rightarrow (A, B)$ for all A and B .

As promised, this yields a way to organise the various categories of plays involved in a typical game model into a coherent categorical structure. Indeed, we will show below that, for quite a few game models, the assignment $L \mapsto \mathbb{P}_L$ induces a category-valued presheaf on $\Delta/\mathbb{A}$, i.e., a functor $(\Delta/\mathbb{A})^{op} \rightarrow \text{Cat}$. Furthermore, the maps δ_k and ι_k introduced earlier will respectively be given by $\mathbb{P}(d_k)$ and $\mathbb{P}(i_k)$.

In the following, we will only need to use this structure up to lists of length 4:

Definition 2.5. For any $p \leq q$ and set $\mathbb{A}$, let $\mathbb{A}_{[p,q]}$ denote the full subcategory of $\Delta/\mathbb{A}$ spanning lists L of length between p and q .

In the next sections, we will define strategies, composition and copycat strategies abstractly, based on the category-valued presheaf $\mathbb{P}$ on $\mathbb{A}_{[1,4]}$. This is quite demanding, but we are rewarded with a higher-level view of composition, which yields abstract proofs of associativity and unitality, assuming a few additional properties of $\mathbb{P}$. We will define a game setting to consist of a set $\mathbb{A}$ and a category-valued presheaf satisfying these additional properties.

2.3 Notions of strategy

Let us now start our reconstruction of a game model from an arbitrary category-valued presheaf $\mathbb{P}$ on $\mathbb{A}_{[1,4]}$. Our first step is to define strategies. Standardly, a strategy $\sigma: A \rightarrow B$ is a prefix-closed set of plays in $\mathbb{P}_{A,B}$ (generally required to be non-empty). Equivalently, it is a functor $\mathbb{P}_{A,B}^{op} \rightarrow 2$, the ordinal 2 viewed as a category. In Tsukada and Ong’s model [37], $\mathbb{P}_{A,B}$ is a proper category, and strategies are generalised to *presheaves* on $\mathbb{P}_{A,B}$, i.e., functors $\mathbb{P}_{A,B}^{op} \rightarrow \text{Set}$. This is indeed a generalisation because 2 embeds into Set (more on this in Section 2.8).

The basis of our approach will be the general notion:

Definition 2.6. Let the category of *strategies* from A to B be $\widehat{\mathbb{P}_{A,B}}$. The category of *boolean strategies* is $\widehat{\mathbb{P}_{A,B}}$.

2.4 Polynomial functors

The next step in our reconstruction of a game model from $\mathbb{A}$ and $\mathbb{P}$ is to define identities and composition, which will rely on polynomial functors, which we now briefly recall.

Notation 1. Any functor $F: \mathcal{C} \rightarrow \mathcal{D}$ induces a restriction, or precomposition functor $\Delta_F: \widehat{\mathcal{D}} \rightarrow \widehat{\mathcal{C}}$ mapping any $X: \mathcal{D}^{op} \rightarrow \text{Set}$ to $X \circ F^{op}$, where $F^{op}: \mathcal{C}^{op} \rightarrow \mathcal{D}^{op}$ acts just as F but on opposite categories. When $\mathcal{C}$ and $\mathcal{D}$ are small, this restriction functor has both a left and a right adjoint, which we respectively denote by Σ_F and Π_F , as in

$$\begin{array}{ccc} & \Sigma_F & \\ \widehat{\mathcal{C}} & \begin{array}{c} \xrightarrow{\perp} \\ \Delta_F \\ \xleftarrow{\perp} \end{array} & \widehat{\mathcal{D}} \\ & \Pi_F & \end{array}$$

The left and right adjoints are respectively given by left and right extension, and enjoy explicit descriptions, both in terms of *coends* and *ends* and in terms of colimits and limits [25, 35]. A brief description of how they work is included for completeness in Appendix A.1, but most of the paper should be accessible without reading it.

Definition 2.7. A functor is *polynomial* iff it is isomorphic to some finite composite of functors of the form Δ_F , Π_F and Σ_F .

This definition is very close to Kock [23, Chapter 16] – it might even turn out to be the same if Kock’s presheaf polynomial functors are ever proved to be closed under composition.

2.5 Copycat as a polynomial functor

As a warm-up before considering composition, we would like to start with our abstract definition of identities, which are standardly given by *copycat* strategies. A natural way to define the copycat strategy $id_A: A \rightarrow A$ is to decree that it accepts all plays in $\mathbb{P}_{A,A}$ which are in the image of the insertion functor $\iota_0: \mathbb{P}_A \rightarrow \mathbb{P}_{A,A}$. Indeed, recalling (1) and according to the discussion of insertions, right after Example 2.3, such plays are precisely those in which M acts as a proxy between L and R , which agrees with the standard definition of copycat strategies.

This definition has the advantage of concreteness, but as announced we need to give an equivalent, polynomial definition. Because an object of a category $\mathcal{C}$ is the same as a functor $1 \rightarrow \mathcal{C}$, we may define id_A as a functor $1 \rightarrow \widehat{\mathbb{P}_{A,A}}$. Furthermore, 1 is a

presheaf category: indeed it is $\widehat{\emptyset}$, presheaves over the empty category. So we may view copycat over A as a functor $\widehat{\emptyset} \rightarrow \widehat{\mathbb{P}_{A,A}}$. In order to present it as a polynomial functor, we will need to assume that the insertion functor $\iota_0: \mathbb{P}_A \rightarrow \mathbb{P}_{A,A}$ is a discrete fibration. Let us recall the definition:

Definition 2.8. A functor $p: \mathcal{E} \rightarrow \mathcal{B}$ is a *discrete fibration* when for all objects $e \in \mathcal{E}$ and morphisms $f: b \rightarrow p(e)$ there exists a unique morphism $u: e' \rightarrow e$ such that $p(u) = f$. Such a morphism is called a *cartesian lifting* of e along f . Let $\text{DFib}_{\mathcal{B}}$ denote the full subcategory of $\text{Cat}/\mathcal{B}$ spanning discrete fibrations.

For us, the relevant property of discrete fibrations is the following characterisation of left extension along them:

Lemma 2.9. For any discrete fibration $p: \mathcal{E} \rightarrow \mathcal{B}$, presheaf $X \in \widehat{\mathcal{E}}$, and object $b \in \mathcal{B}$, we have $\sum_p(X)(b) \cong \sum_{e|p(e)=b} X(e)$, where $\sum$ means left extension in the left-hand side and disjoint union on the right-hand side.

Here is our polynomial presentation of copycat:

Proposition 2.10. If the insertion functor $\iota_0: \mathbb{P}_A \rightarrow \mathbb{P}_{A,A}$ is a discrete fibration, then the functor

$$\widehat{\emptyset} \xrightarrow{\Pi_!} \widehat{\mathbb{P}_A} \xrightarrow{\Sigma_{\iota_0}} \widehat{\mathbb{P}_{A,A}}$$

is isomorphic the copycat strategy id_A .

Proof. See Section B.1. $\square$

2.6 Composition as a polynomial functor

The next step is to express composition of strategies using the same language of polynomial functors. Let us first recall the standard definition in the boolean case: the composite $\sigma; \tau$ of two boolean strategies σ and τ over (A, B) and (B, C) respectively, is defined to accept all plays $p \in \mathbb{P}_{A,C}$ for which there exists $u \in \mathbb{P}_{A,B,C}$ such that $\delta_1(u) = p$ and

$$\delta_2(u) \in \sigma \quad \text{and} \quad \delta_0(u) \in \tau.$$

In Tsukada and Ong [37], this is extended to a polynomial functor $\widehat{\mathbb{P}_{A,B}} \times \widehat{\mathbb{P}_{B,C}} \rightarrow \widehat{\mathbb{P}_{A,C}}$, whose definition is essentially a proof-relevant version of the boolean one. If we understand ‘ σ accepts play p ’ as $\sigma(p) = \{\star\}$, then, e.g., $\delta_2(u) \in \sigma$ above becomes $\star \in \sigma(\delta_2(u))$. We get:

Definition 2.11. The composite $\sigma; \tau$ of two strategies σ and τ over (A, B) and (B, C) respectively, is defined to map any play $p \in \mathbb{P}_{A,C}$ to the set of triples (u, x, y) where $u \in \mathbb{P}_{A,B,C}$ is such that $\delta_1(u) = p$ and

$$x \in \sigma(\delta_2(u)) \quad \text{and} \quad y \in \tau(\delta_0(u)).$$

Let us present this polynomially. First, by universal property of coproduct we have $\widehat{\mathbb{P}_{A,B}} \times \widehat{\mathbb{P}_{B,C}} \cong \widehat{\mathbb{P}_{A,B} + \mathbb{P}_{A,B}}$, so we reduce to defining a functor $\widehat{\mathbb{P}_{A,B} + \mathbb{P}_{A,B}} \rightarrow \widehat{\mathbb{P}_{A,C}}$. Here is our candidate composition:

Definition 2.12. Let m denote the polynomial functor

$$\widehat{\mathbb{P}_{A,B} + \mathbb{P}_{B,C}} \xrightarrow{\Delta_{\delta_2 + \delta_0}} \widehat{\mathbb{P}_{A,B,C} + \mathbb{P}_{A,B,C}} \xrightarrow{\Pi_{[id, id]}} \widehat{\mathbb{P}_{A,B,C}} \xrightarrow{\Sigma_{\delta_1}} \widehat{\mathbb{P}_{A,C}}.$$

This definition is legitimated by:

Proposition 2.13. The functor m agrees with Definition 2.11, i.e., for all σ and τ , we have $(\sigma; \tau) \cong m[\sigma, \tau]$.

Proof sketch (see Section B.1 for a full proof). By discrete fibredness, the final Σ_{δ_1} says that for all $p \in \mathbb{P}_{A,C}$, the formula for $m[\sigma, \tau](p)$ will start by

$$\sum_{u \in \mathbb{P}_{A,B,C} | \delta_1(u)=p} \dots (u).$$

By a very general computation, the $\Pi_{[id, id]}$ says that, viewing the intermediate result in $\mathbb{P}_{A,B,C} + \mathbb{P}_{A,B,C}$ as a pair $[\sigma', \tau']$, the formula continues with $\sigma'(u) \times \tau'(u)$, which in our case directly instantiates to $\sigma(\delta_2(u)) \times \tau(\delta_0(u))$, as desired. $\square$

Remark 1. The discrete fibredness hypothesis is satisfied in most game models, with the notable exception of the saturated interpretation of AJM games (see Section 3.4), in which the projection is a non-discrete fibration. The construction still goes through, but we here stick to discrete fibrations for simplicity.

2.7 Game settings, associativity and unitality

We have now expressed copycat strategies and composition abstractly, relying only on the postulated category-valued presheaf. Let us now consider associativity. It has become standard in game semantics to prove associativity of composition using a *zipping* result [8] stating that both squares

$$\begin{array}{ccc} \mathbb{P}_{A,B,C,D} & \longrightarrow & \mathbb{P}_{A,B,D} \\ \downarrow \lrcorner & & \downarrow \\ \mathbb{P}_{B,C,D} & \longrightarrow & \mathbb{P}_{B,D} \end{array} \quad \begin{array}{ccc} \mathbb{P}_{A,B,C,D} & \longrightarrow & \mathbb{P}_{A,C,D} \\ \downarrow \lrcorner & & \downarrow \\ \mathbb{P}_{A,B,C} & \longrightarrow & \mathbb{P}_{A,C} \end{array} \quad (3)$$

are pullbacks on objects. This holds in all considered game models, and constitutes the last bit of our axiomatisation:

Definition 2.14. A *game setting* consists of a set $\mathbb{A}$ (whose elements we call *arenas*) and a category-valued presheaf $\mathbb{P}$ on $\mathbb{A}_{[1,4]}$ such that all projections $\mathbb{P}_{A,B,C} \rightarrow \mathbb{P}_{A,C}$ and insertions $\mathbb{P}_A \rightarrow \mathbb{P}_{A,A}$ are discrete fibrations, and all squares (3) are pullbacks on objects.

We call both squares (3) the *zipping* squares of $\mathbb{P}$.

One of our main results is:

Theorem 2.15. In any game setting, composition of strategies is associative up to isomorphism, and copycat strategies are units up to isomorphism.

Proof. See Section B.2. $\square$

2.8 The boolean case

Let us conclude this section by treating the boolean case: until now, our strategies were given by general presheaves (Definition 2.6). We would like to derive from Theorem 2.15 that boolean strategies also form a category.

The bridge to the boolean case is given by the embedding $r: 2 \hookrightarrow \text{Set}$ mapping $0 \leq 1$ to $\emptyset \rightarrow 1$. This functor has a left adjoint l mapping $\emptyset$ to 0 and collapsing all non-empty sets to 1 . Furthermore, r being fully faithful, we have in fact a full reflection (please note: this means that the right adjoint is fully faithful, not the left adjoint!), which induces a further one between presheaves and boolean presheaves:

Proposition 2.16. For any small category $\mathcal{C}$, post-composition by l and r yield a full reflection

$$\begin{array}{ccc} & l_! & \\ [\mathbb{C}^{op}, \text{Set}] & \xleftrightarrow{\perp} & [\mathbb{C}^{op}, 2] \\ & r_! & \end{array}$$

The left adjoint $l_!$ is called booleanisation.

Proof. The functor $r_!$ is clearly a continuous full embedding, hence a right adjoint by the special adjoint functor theorem, and thus a full, reflective embedding. $\square$

Because 2 is complete and cocomplete, replacing Set with 2 in Notation 1 yields a notion of *boolean polynomial functor*.

Notation 2. Any functor $F: \mathcal{C} \rightarrow \mathcal{D}$ induces restriction, left extension and right extension functors between boolean presheaf categories $\widehat{\mathcal{C}}$ and $\widehat{\mathcal{D}}$, respectively denoted by $\overline{\Delta}_F$, $\overline{\Sigma}_F$ and $\overline{\Pi}_F$. Accordingly, the boolean version of any polynomial functor P will be denoted by $\overline{P}$.

We may thus transfer our polynomial definitions of copycat and composition to boolean strategies. Concrete examples of game settings will be considered in Section 3, for which we have:

Proposition 2.17. In all the game settings of Section 3, $\overline{m}$ coincides with standard composition.

As desired, we obtain:

Proposition 2.18. In any game setting, composition of boolean strategies is associative and unital up to isomorphism.

Proof. See Section B.4. $\square$

Remark 2. Please note that we have not claimed that boolean composition agrees with general, set-based composition, i.e., commutation of the left-hand diagram below.

$$\begin{array}{ccc} \widehat{\mathbb{P}_{A,B} + \mathbb{P}_{B,C}} & \xrightarrow{m} & \widehat{\mathbb{P}_{A,C}} \\ \uparrow r_! & & \uparrow r_! \\ \mathbb{P}_{A,B} + \mathbb{P}_{B,C} & \xrightarrow{\overline{m}} & \mathbb{P}_{A,C} \end{array}$$

In fact it does not in general, and this is the main cause for the failure of stability of boolean, innocent strategies under composition [15, Section 3.7.2]. What does hold, however, is

- commutation of booleanisation with composition as on the left below (this is the main idea for the proof of Proposition 2.18),
- the characterisation of boolean composition given below right, as set-based composition followed by booleanisation.

$$\begin{array}{ccc} \widehat{\mathbb{P}_{A,B} + \mathbb{P}_{B,C}} & \xrightarrow{m} & \widehat{\mathbb{P}_{A,C}} \\ \downarrow l_! & & \downarrow l_! \\ \mathbb{P}_{A,B} + \mathbb{P}_{B,C} & \xrightarrow{\overline{m}} & \mathbb{P}_{A,C} \end{array} \quad \begin{array}{ccc} \widehat{\mathbb{P}_{A,B} + \mathbb{P}_{B,C}} & \xrightarrow{m} & \widehat{\mathbb{P}_{A,C}} \\ \uparrow r_! & & \uparrow r_! \\ \mathbb{P}_{A,B} + \mathbb{P}_{B,C} & \xrightarrow{\overline{m}} & \mathbb{P}_{A,C} \end{array}$$

Let us move on to exhibit a few concrete game settings. We will return to the boolean case in Section 4.3 where we consider boolean innocence.

3 Applications

3.1 Hyland-Ong games and strategies

Let us now consider Hyland-Ong games in more detail, and show that they form a game setting. We mostly follow Harmer's [15] presentation. For simplicity, we adopt the following innocuous¹ modification of the standard notion of arena:

¹Tsukada and Ong use forest-shaped arenas.

Definition 3.1. An *arena* is a simple, countable, directed acyclic graph A equipped with a subset $\sqrt{A}$ of *initial* vertices, or *roots*, such that for all vertices m , all paths from m to some initial vertex have the same parity.

In particular, simple, upside-down forests form arenas. The intuition is that reachable vertices of an arena are moves in a two-player game, and that an edge $m \rightarrow m'$ in the forest means that m is *enabled*, or *justified* by m' . If the path from m to some root has even length, then O (for *Opponent*) is playing; otherwise P (for *Proponent*) is. E.g., all roots are played by O .

Example 3.2. A very simple arena, called o , is the single-vertex graph. For a less trivial example, the boolean type $\mathbb{B}$ may be interpreted as the arena

$$\begin{array}{c} q \\ \swarrow \quad \searrow \\ t \quad \quad f \end{array}$$

Now that we have defined arenas, let us move on to define plays. The idea, explained at length, e.g., in McCusker [27], is that plays are sequences of moves in which O and P take turns. But a subtlety is that moves may be played several times. So for any edge $m \rightarrow m'$ in the considered arena, there may be several occurrences of m and m' . We thus decorate sequences of moves with justification pointers matching those of the considered arena.

Definition 3.3. A *justified sequence* on any arena A consists of a natural number $n \in \mathbb{N}$, equipped with maps $f: n \rightarrow \text{ob}(A)$ and $\varphi: n \rightarrow \{0\} \uplus n$ (recalling Section 1.3) such that, for all $i \in n$,

- $\varphi(i) < i$,
- if $\varphi(i) = 0$ then $f(i) \in \sqrt{A}$, and
- if $\varphi(i) \neq 0$, then there is an edge $f(i) \rightarrow f(\varphi(i))$ in A .

Let $\mathbb{P}_A$ denote the poset of justified sequences on A , with prefix ordering.

We will draw justified sequences (n, f, φ) as the sequence of their $f(i)$'s, with arrows to denote φ , as is standard in game semantics.

Example 3.4. Here is a justified sequence in the boolean arena:

$$q \xleftarrow{\varphi} q \xleftarrow{\varphi} t \xleftarrow{\varphi} f, \text{ where times flows to the right.}$$

As mentioned in (1), game semantics proceeds by letting a middle player M play on two arenas A and B , with specific restrictions on when switching between A and B is allowed. For this, we form the compound arena $A \rightarrow B$:

Definition 3.5. For any two arenas A and B , let $A \rightarrow B$ denote the arena obtained by taking the disjoint union of A and B as directed graphs, adding an edge $m \rightarrow m'$ for all $m \in \sqrt{A}$ and $m' \in \sqrt{B}$, and taking $\sqrt{B}$ as $\sqrt{A \rightarrow B}$ (if B is not empty, otherwise we take $A \rightarrow B$ to be empty).

The switching constraints mentioned above are implemented by considering a subposet of $\mathbb{P}_{A \rightarrow B}$:

Definition 3.6. For any two arenas A and B , let $\mathbb{P}_{A,B}$ denote the poset of *plays* on (A, B) , i.e., *alternating* justified sequences of even length on $A \rightarrow B$.

Alternation here means that, for any $s = (n, f, \varphi)$, $f(i)$ is played by O iff i is odd.

Example 3.7. Recalling the arena o from Example 3.2, and calling its unique move r , here is an example play on $((\mathbb{B} \rightarrow o), o)$:

where time flows downwards (so the play really is $r'rqrqqr$) and arrows denote justification pointers.

We are now in a position to define the insertion functors $\iota_0 : \mathbb{P}_A \rightarrow \mathbb{P}_{A,A}$: they send any justified sequence $m_1 \dots m_n$ to some play in which P mimics the behaviour of O . The technical definition is not particularly illuminating, but the example following it should convince the reader that nothing really difficult is going on here.

Definition 3.8. For any justified sequence $p = (m_1 \dots m_n)$ on A , let $\iota_0(p)$ be the sequence $(m_1, k_1)(m_1, l_1) \dots (m_n, k_n)(m_n, l_n)$, where k_i and l_i denote either 0 or 1, according to the component of $A \rightarrow A$ in which the move is played. If m_i is an O -move, then $k_i = 1$ and $l_i = 0$; and otherwise $k_i = 0$ and $l_i = 1$. Pointers are as in p except for initial moves on the left, i.e., moves of the form $(m, 0)$ with $m \in \sqrt{A}$, which are justified by the corresponding $(m, 1)$.

Example 3.9. The justified sequence on $\mathbb{B}$ below left, which is not alternating, yields the copycat play on the right:

The next step is interaction sequences, for which the basic idea is: any play in $(A \rightarrow B) \rightarrow C$ may be projected to $\mathbb{P}_{A \rightarrow B}, \mathbb{P}_{B \rightarrow C}$, and even $\mathbb{P}_{A \rightarrow C}$, by prolongating pointers (i.e., $a \rightarrow b \rightarrow c$ becomes $a \rightarrow c$). Following Example 2.2, we put:

Definition 3.10. An *interaction sequence* is a justified sequence on $(A \rightarrow B) \rightarrow C$ ending in A or C , whose projections to $A \rightarrow B$ and $B \rightarrow C$ are plays. Let $\mathbb{P}_{A,B,C}$ denote the poset of interaction sequences with prefix ordering.

As desired, the projection to $A \rightarrow C$ is also a play, and we have monotone maps $\delta_k : \mathbb{P}_{A_0, A_1, A_2} \rightarrow \mathbb{P}_{A_i, A_j}$ with $i < j$ in $\{0, 1, 2\} \setminus \{k\}$.

We may define *generalised interaction sequences* similarly to obtain:

Proposition 3.11. The category-valued presheaf $\mathbb{P}$ defined by respectively taking $\mathbb{P}_A, \mathbb{P}_{A,B}, \mathbb{P}_{A,B,C}$ and $\mathbb{P}_{A,B,C,D}$ to be the posets of all justified sequences, plays, interaction sequences and generalised interaction sequences, for all arenas A, B, C, D , with projections and insertions as above, forms a game setting.

Proof. Projections $\delta_1 : \mathbb{P}_{A,B,C} \rightarrow \mathbb{P}_{A,C}$ are discrete fibrations: the restriction of any $u \in \mathbb{P}_{A,B,C}$ along any $p \leq \delta_1(u)$ may be taken to

be the shortest prefix of u whose projection is p (longer such prefixes do not end in A or C). The fact that squares (3) are pullbacks on objects is the standard zipping lemma. $\square$

3.2 Constraining strategies: local constraints

In the previous section, we consider a rather rough notion of play. Standardly, further constraints are considered on strategies, such as P -visibility, O -visibility, well-threadedness, and well-bracketing (when games are equipped with an appropriate question-answer discipline). E.g., a P -visible strategy is one which only accepts P -visible plays. One then needs to prove that such constraints are *robust*, i.e., are preserved by composition and satisfied by identities. This is done in a very clean and modular way in Harmer's thesis [15, Chapter 3].

In order for our framework to apply to such constrained strategies, we may start from the game setting for unconstrained plays and convert the proof of robustness of constraint c into the construction of a sub-game setting $\mathbb{P}^c$.

For each constraint $c \in \{P\text{-vis}, O\text{-vis}, wb, wt\}$, respectively denoting P -visibility, O -visibility, well-bracketing and well-threadedness:

- $\mathbb{P}_{A,B}^c$ is the full sub-poset of $\mathbb{P}_{A,B}$ consisting of plays satisfying the constraint c ;
- $\mathbb{P}_A^c$ is the full sub-poset of $\mathbb{P}_A$ whose insertions are in $\mathbb{P}_{A,A}^c$;
- $\mathbb{P}_{A,B,C}^c$ is the full sub-poset of $\mathbb{P}_{A,B,C}$ whose projections to $\mathbb{P}_{A,B}$ and $\mathbb{P}_{B,C}$ respectively factor through $\mathbb{P}_{A,B}^c$ and $\mathbb{P}_{B,C}^c$;
- $\mathbb{P}_{A,B,C,D}^c$ is the full sub-poset of $\mathbb{P}_{A,B,C,D}$ whose projections to $\mathbb{P}_{A,B}$, $\mathbb{P}_{B,C}$ and $\mathbb{P}_{C,D}$ respectively factor through $\mathbb{P}_{A,B}^c$, $\mathbb{P}_{B,C}^c$ and $\mathbb{P}_{C,D}^c$.

One delicate point is then to show that unmentioned projections also factor through the appropriate constrained posets.

Proposition 3.12. For all constraints c , $\mathbb{P}_{A,B,C}^c \rightarrow \mathbb{P}_{A,C}$ factors through $\mathbb{P}_{A,C}^c$.

Proof. For P -visibility, this is [15, Proposition 3.4.3]. The (implicit) proof of [26, Lemma 2.8] handles O -visibility and well-bracketing. Well-threadedness follows similarly to P -visibility. $\square$

Proposition 3.13. For all constraints c , $\mathbb{P}^c$ forms a sub-game setting of $\mathbb{P}$.

Proof. First, to see that projections δ_1 are discrete fibrations, it is enough to observe that all involved constraints are stable under prefix. It remains to show that constrained plays satisfy zipping. But constraints are merely imposed on the projections of interaction sequences, and thus are clearly stable under zipping. $\square$

Corollary 3.14. For all sets of constraints in $\{P\text{-vis}, O\text{-vis}, wb, wt\}$, arenas and strategies satisfying these constraints form a category.

As satisfactory as it seems, this result does not tell us that composition is the considered sub-game settings agrees with the original. Let us show that it does, by considering the general case. Consider any embedding $c : \mathbb{P}^c \hookrightarrow \mathbb{P}$ of game settings sharing a common set of arenas. We need to prove that the polynomial functor, say m^c , mimicking our polynomial composition m on $\mathbb{P}^c$ is compatible with the inclusion $c : \mathbb{P}^c \hookrightarrow \mathbb{P}$. In order to make this more precise, let us observe that each strategy $\sigma \in \widehat{\mathbb{P}_{A,B}^c}$ is canonically mapped to $\Sigma_c(\sigma) \in \widehat{\mathbb{P}_{A,B}}$. Intuitively, $\Sigma_c(\sigma)$ should act just as σ on $\mathbb{P}_{A,B}^c$ and be empty elsewhere. In order for this to hold, we merely

What's in a game?

need to assume that the considered constraint is stable under prefix, which in the general case amounts to requiring c to be a discrete fibration.

Lemma 3.15. *If c is a discrete fibration, then for any σ we have*

- $\sum_c(\sigma)(c(p)) \cong \sigma(p)$ for all $p \in \mathbb{P}_{A,B}^c$ and
- $\sum_c(\sigma)(q) = \emptyset$ for all $q \in \mathbb{P}_{A,B}$ outside the essential image of c .

Proof. By Lemma 2.9. $\square$

We may now express the desired compatibility of composition with c as $m(\sum_c(\sigma), \sum_c(\tau)) \cong \sum_c(m(\sigma, \tau))$. This however requires an additional hypothesis, saying that an interaction sequence $u \in \mathbb{P}_{A,B,C}$ is in $\mathbb{P}_{A,B}^c$ as soon as its projections to $\mathbb{P}_{A,B}$ and $\mathbb{P}_{B,C}$ are.

Example 3.16. To see why, imagine that there exists some interaction sequence $u \in \mathbb{P}_{A,B,C} \setminus \mathbb{P}_{A,B,C}^c$ such that $\delta_2(u)$ and $\delta_0(u)$ are in the respective essential images of $\mathbb{P}_{A,B}^c$ and $\mathbb{P}_{B,C}^c$, say as $c(p_1)$ and $c(p_2)$. Further assuming that σ and τ accept p_1 and p_2 , $m(\sum_c(\sigma), \sum_c(\tau))$ clearly accepts $\delta_1(u)$, while $\sum_c(m(\sigma, \tau))$ does not, because the interaction sequence which could witness it lies outside $\mathbb{P}^c$ (assuming that no interaction sequence from $\mathbb{P}^c$ projects to $\delta_1(u)$).

This hypothesis essentially says that the considered constraint may be checked locally, hence our terminology:

Definition 3.17. An embedding $c: \mathbb{P}^c \hookrightarrow \mathbb{P}$ between game settings sharing the same set of arenas is a *local constraint* iff its components are discrete fibrations and any $u \in \mathbb{P}_{A,B,C}$ is essentially in the image of c if $\delta_2(u)$ and $\delta_0(u)$ are.

Remark 3. *The locality constraint may be expressed concisely as a sheaf condition (see Lemma B.9).*

Proposition 3.18. *For all local constraints $c: \mathbb{P}^c \hookrightarrow \mathbb{P}$, the square*

$$\begin{array}{ccc} \widehat{\mathbb{P}_{A,B}^c} \times \widehat{\mathbb{P}_{B,C}^c} & \xrightarrow{m^c} & \widehat{\mathbb{P}_{A,C}^c} \\ \downarrow \Sigma_c \times \Sigma_c & & \downarrow \Sigma_c \\ \widehat{\mathbb{P}_{A,B}} \times \widehat{\mathbb{P}_{B,C}} & \xrightarrow{m} & \widehat{\mathbb{P}_{A,C}} \end{array}$$

commutes up to isomorphism.

Proof. See Section B.5. $\square$

3.3 Constraining strategies: predicates

Beyond the constraints mentioned above, a similar result may be proved for the refined notion of game in McCusker's thesis [26]. McCusker's games A are just like arenas, except that they come equipped with an abstract *validity* predicate P_A , which is a subset of the set L_A of *legal plays*, i.e., alternated, well-bracketed, P - and O -visible justified sequences. This predicate is only required to be non-empty, prefix-closed, and such that for all $p \in P_A$ and set I of occurrences of initial moves in p , the restriction $p|_I$ of p to moves hereditarily justified by some move in I is again in P_A .

McCusker then defines $P_{A \rightarrow B}$ to consist of legal plays in $L_{A \rightarrow B}$ whose projections to A and B are in P_A and P_B (instead of simply L_A and L_B), respectively. Interaction sequences are then defined as in the standard case, with the additional constraint that both projections to (A, B) and (B, C) are in $P_{A \rightarrow B}$ and $P_{B \rightarrow C}$, respectively.

Let us now show that this again yields a game setting. Denoting the new posets by $\mathbb{P}_{A,B}^P, \mathbb{P}_{A,B,C}^P$, etc., we have:

Proposition 3.19. *Projections $\mathbb{P}_{A,B,C}^P \rightarrow \mathbb{P}_{A,C}^P$ factors through $\mathbb{P}_{A,C}^P$.*

Proof. Being in $P_{A \rightarrow C}$ is only about projections being in P_A and P_C , which is taken care of by the condition on projections to (A, B) and (B, C) . $\square$

Similarly:

Proposition 3.20. *$\mathbb{P}^P$ is a game setting, and a local constraint of $\mathbb{P}$.*

Proof. As before, projections are discrete fibrations because validity predicates are stable under prefix. It only remains to show that zipping holds, which again follows from $P_{A \rightarrow B}$ being only about projections to A and B . $\square$

There are other kinds of constraints like innocence or single-threadedness, which may not be treated this way. We will deal with innocence in Section 4.

3.4 AJM games: a partial answer

Let us now briefly consider an alternative approach to game semantics by Abramsky et al. [6]. On the one hand, this approach is more elementary than Hyland and Ong's in that games do not feature justification pointers. So, e.g., composition of strategies is significantly simpler. On the other hand, games feature a partial equivalence relation between plays, which needs to be dealt with at the level of strategies.

In order to organise such games into a game setting, we have two sensible choices for the notion of morphism between plays: beyond the prefix ordering, we may also incorporate equivalence between plays. Presheaves then amount to so-called *saturated* strategies. We adopt Harmer's presentation [15].

Definition 3.21. A *game* A consists of two sets O_A and P_A , respectively of *Opponent* and *Proponent* moves, equipped with a partial, prefix-closed equivalence relation $\approx$ on alternated sequences of moves started by Opponent, for which any two equivalent plays have the same length, and such that

if $s \approx t$ and $sa \approx sa$, then there exists a' such that $sa \approx ta'$.

Let $\mathbb{P}_A$ consist of all alternated sequences of moves s started by Opponent, such that $s \approx s$. Then, for any games A and B , we form the game $A \rightarrow B$, which has $O_{A \rightarrow B} = P_A + O_B$ and $P_{A \rightarrow B} = O_A + P_B$ and $s \approx_{A \rightarrow B} t$ iff s and t play in the same component at each stage and their projections are equivalent in A , resp. B . The poset $\mathbb{P}_{A,B}$ may then be defined as the set of plays in $A \rightarrow B$ equipped with prefix ordering. Similarly defining $\mathbb{P}_{A,B,C}$ and $\mathbb{P}_{A,B,C,D}$ we obtain:

Proposition 3.22. *AJM games form a game setting.*

Proof. Squares (3) being pullbacks on objects is the standard zipping lemma. To show that projections $\delta_1: \mathbb{P}_{A,B,C} \rightarrow \mathbb{P}_{A,C}$ are discrete fibrations, we need to be able to canonically restrict any $u \in \mathbb{P}_{A,B,C}$ along any $s \leq \delta_1(u)$: just take the longest prefix of u mapped to s . $\square$

For saturated strategies, the idea is to incorporate for all A, B the partial equivalence relations $\approx_A$ and $\approx_B$ into the category of plays.

Proposition 3.23. *AJM games form a category-valued presheaf by mapping each list of games to the corresponding set of plays with as morphisms between any two plays u and v :*

- a singleton when there exists some play w such that $u \approx w \leq v$, or equivalently there exists w such that $u \leq w \approx v$;
- none otherwise.

However, the obtained category-valued presheaf is not a game setting, because projections $\mathbb{P}_{A,B,C} \rightarrow \mathbb{P}_{A,C}$ and insertions $\mathbb{P}_A \rightarrow \mathbb{P}_{A,A}$ are not discrete fibrations. Indeed, the fibres of $\mathbb{P}_{A,B,C} \rightarrow \mathbb{P}_{A,C}$ are proper groupoids in general, thus making it a non-discrete Grothendieck fibration. The case of $\mathbb{P}_A \rightarrow \mathbb{P}_{A,A}$ is worse: the restriction of a play in $s \in \mathbb{P}_A$ along a morphism $p \rightarrow q$ in $\mathbb{P}_{A,A}$ may at best be mapped to some p' isomorphic to p in general, thus making it a *Street fibration*. Our approach may generalise in this direction, but this will involve advanced categorical concepts such as stacks (which are to fibrations as sheaves are to discrete fibrations), so we leave it for future work.

3.5 A non-example: Blass games

In the previous sections, we have shown that several approaches to game semantics form game settings, with the exception of the saturated AJM setting. It may be instructive to consider Blass's games [9, 10], as they are well-known for their non-associative composition. Our account essentially follows Abramsky [4, Section 3], through the lens of game settings.

Definition 3.24. A Blass game consists of a family of rooted trees, together with a *polarity* in $\{P, O\}$.

Vertices are thought of as positions in the game, with alternating polarities. The given polarity indicates which player is to start the game, by choosing the initial position. The fact that Proponent may start is a notable difference with arena games. Another difference is that the given family of trees genuinely represents the 'game tree' – no move may be played twice. This determines the definition of $\mathbb{P}_A$, for any game $A = (T, \pi)$: it is the poset consisting of positions (i.e., vertices of T , plus a formal initial position), with $x \leq y$ when x is above y in T .

For $\mathbb{P}_{A,B}$ things become a bit more complicated. Strategies should be based on the *linear implication* game $A \rightarrow B$, which is constructed much as in, e.g., AJM games. First, one lets $A^\perp$ denote the game with the same family of trees as A but with opposite polarity. Then, one defines $A \rightarrow B$ by interleaving moves from $A^\perp$ and B with natural switching conditions: Opponent is to play as soon as possible. In other words, if the respective polarities in $A^\perp$ and B are OP , PO , or OO , then O is to play; otherwise P is. There is a catch, however: if the polarity is OO and Opponent plays, say in B , we reach a position with polarity OP , and Opponent is to play again, which breaks alternation. This is rectified by defining $A \rightarrow B$ to comprise *compound* moves from OO to PP , for each pair of moves in $A^\perp$ and B . This settles the definition of $\mathbb{P}_{A,B}$, up to some technicalities.

The next level is to define $\mathbb{P}_{A,B,C}$. Glossing over the details, this should consist of sequences of moves in A , B and C , whose projections to $\mathbb{P}_{A,B}$, $\mathbb{P}_{B,C}$ and $\mathbb{P}_{A,C}$ are well-defined. However, we may show that with these definitions, the squares (3) cannot both be pullbacks in general. Indeed, consider the case where the respective polarities of A , B , C and D are O , P , O and P , and A is non-empty. Then, let $\mathbb{P}_{A,B,C,D}^l$ denote the left-hand pullback and $\mathbb{P}_{A,B,C,D}^r$ denote the right-hand one. We will show that both pullbacks cannot be the same category by exhibiting a play in $\mathbb{P}_{A,B,C,D}^l$

which is not in $\mathbb{P}_{A,B,C,D}^r$. First, let us observe that the initial polarities from the respective points of view of $A \rightarrow B$, $B \rightarrow C$ and $C \rightarrow D$ are like so:

$$\begin{array}{ccccc} A & \longrightarrow & B & & B & \longrightarrow & C & & C & \longrightarrow & D \\ A^\perp & & B & & B^\perp & & C & & C^\perp & & D \\ P & & P & & O & & O & & P & & P. \end{array}$$

Letting a denote any root of A , the sequence a is then legal in $\mathbb{P}_{A,B,D}$ (the polarities are PP both in $A \rightarrow B$ and $A \rightarrow D$) and the empty sequence is legal in $\mathbb{P}_{B,C,D}$. Thus, a is legal in $\mathbb{P}_{A,B,C,D}^l$ by the left-hand pullback. However, if the two pullbacks were isomorphic, then by the properties of projections $a \in \mathbb{P}_{A,B,C,D}^l$ would be mapped to $a \in \mathbb{P}_{A,B,C,D}^r$ under the isomorphism. But $\mathbb{P}_{A,B,C,D}^r$ cannot contain a because this play is illegal in $\mathbb{P}_{A,B,C}$ (because the polarity is PO in $A \rightarrow C$).

4 Innocence

4.1 Concurrent innocence

In the previous sections, we have constructed a category of games and strategies parameterised over an arbitrary game setting which presents the advantage of unifying a number of such categories as instances of the same construction. However, in many game models, notably of purely functional languages, the relevant category is the identity-on-objects subcategory of innocent strategies. This raises the question of whether game settings are appropriate for an abstract account of innocence. More precisely, this should involve defining innocent strategies and proving that innocent strategies form a subcategory.

In order to achieve this, we will use the recent recasting of innocence as a sheaf condition [18, 19, 37]. Starting from Hyland-Ong games, the first step is to refine the posets $\mathbb{P}_A, \mathbb{P}_{A,B}, \dots$ into proper categories (with exactly the same objects), say $\mathbb{P}_A^+, \mathbb{P}_{A,B}^+, \dots$, with the crucial feature that for any play $p \in \mathbb{P}_{A,B}^+$ and move $m \in p$, there is a morphism $[p]_m \rightarrow p$ from the P -view of m to p^2 . This of course does not hold with the prefix ordering, as the view is rarely a prefix. This idea was introduced by Melliès [28] in a slightly different setting.

Passing from $\mathbb{P}$ to $\mathbb{P}^+$ raises the issue of how to extend the abstract framework. Should it now contain two category-valued presheaves? Or should we simply forget about prefix-based strategies and accept $\mathbb{P}^+$ as the new basic set up? We do not make any definitive choice here, but for simplicity and modularity reasons, we choose to first work with $\mathbb{P}^+$ only, and in a second round explore the connection with $\mathbb{P}$.

Indeed, perhaps surprisingly, we have:

Proposition 4.1. *Tsukada and Ong's $\mathbb{P}^+$ forms a game setting.*

Proof. Follows from Lemmas 39, 46 and 47 of Tsukada and Ong [36]. $\square$

Returning to the abstract setting, the new data thus merely consists of a full subcategory $i_{A,B} : \mathbb{V}_{A,B} \hookrightarrow \mathbb{P}_{A,B}^+$, for all A, B , whose objects are called *views*.

Definition 4.2. The category of *innocent* strategies is the essential image of $\prod_{i_{A,B}} : \widehat{\mathbb{V}_{A,B}} \rightarrow \widehat{\mathbb{P}_{A,B}^+}$. The domain $\widehat{\mathbb{V}_{A,B}}$ is the category of *behaviours*.

²We omit the definition of views, as it is unnecessary for understanding the rest.

We now would like to establish that in any game setting equipped with such full embeddings, innocent strategies form a subcategory. However, our proof relies on two additional properties. The first, already observed by Tsukada and Ong [37, Lemma 32], states that one can reconstruct uniquely any interaction sequence from its projection to $\mathbb{P}_{A,C}$, say u , together with a compatible family, for each view v of u of an interaction sequence projecting to v . The second property essentially says that any morphism $v \rightarrow \delta_2(u)$ from a view $v \in \mathbb{V}_{A,B}$ to the projection of some $u \in \mathbb{P}_{A,B,C}^+$ factors canonically through the projection of some view (and similarly for δ_0).

Let us introduce both properties in more detail.

The first property essentially says that interaction is local. To state it, we need to recall the following standard construction and its prime property.

Definition 4.3. The category of elements $\mathbb{C}/X$ of any presheaf $X \in \widehat{\mathbb{C}}$ has as objects all pairs (c, x) with $x \in X(c)$ and as morphisms $(c, x) \rightarrow (c', x')$ all morphisms $f: c \rightarrow c'$ such that $x' \cdot f = x$. Let $p_X: \mathbb{C}/X \rightarrow \mathbb{C}$ map any (c, x) to c , and any f to itself.

Proposition 4.4. The assignment $X \mapsto (\mathbb{C}/X, p_X)$ extends to a functor $\text{el}: \widehat{\mathbb{C}} \rightarrow \text{DFib}_{\mathbb{C}}$ (recalling Definition 2.8).

Proof. The cartesian lifting of (c', x') along any $f: c \rightarrow c'$ is given by f itself, viewed as a morphism $(c, x' \cdot f) \rightarrow (c', x')$ in $\mathbb{C}/X$. $\square$

This functor el is in fact an adjoint equivalence.

Definition 4.5. Consider the functor $\text{sing}: \text{DFib}_{\mathbb{C}} \rightarrow \widehat{\mathbb{C}}$ mapping any discrete fibration $p: \mathbb{E} \rightarrow \mathbb{C}$ to the presheaf $\text{sing}(p)(c) = p^{-1}(c)$, with action on morphisms given by cartesian lifting.

Proposition 4.6. The functors

$$\begin{array}{ccc} \widehat{\mathbb{C}} & \xrightleftharpoons[\text{sing}]{\text{el}} & \text{DFib}_{\mathbb{C}} \\ & \perp & \\ & \text{sing} & \end{array}$$

form an adjoint equivalence.

Let us now state the first property we need to impose on game settings with embeddings $i_{A,B}: \mathbb{V}_{A,B} \hookrightarrow \mathbb{P}_{A,B}^+$. The projection $\mathbb{P}_{A,B,C} \rightarrow \mathbb{P}_{A,C}$, as a discrete fibration, induces a presheaf $\text{sing}(\mathbb{P}_{A,B,C})$ on $\mathbb{P}_{A,C}$ which we will require to be in the essential image of

$$\prod_{i_{A,B}}: \widehat{\mathbb{V}_{A,C}} \rightarrow \widehat{\mathbb{P}_{A,C}}.$$

This is equivalent to requiring that $\text{sing}(\mathbb{P}_{A,B,C})$ be a *sheaf* for the Grothendieck topology induced by the embedding $\mathbb{V}_{A,C} \rightarrow \mathbb{P}_{A,C}$. Similarly, we require the presheaf induced by $i_0: \mathbb{P}_A \rightarrow \mathbb{P}_{A,A}$ to be a sheaf for the Grothendieck topology induced by the embedding $\mathbb{V}_{A,A} \rightarrow \mathbb{P}_{A,A}$. Let us record this as:

Definition 4.7. A game setting $(A, \mathbb{P})$ equipped with full embeddings $i_{A,B}: \mathbb{V}_{A,B} \hookrightarrow \mathbb{P}_{A,B}$ is *local* iff $\text{sing}(\mathbb{P}_{A,B,C})$ and $\text{sing}(\mathbb{P}_A)$ are sheaves.

Proposition 4.8. Tsukada and Ong's $\mathbb{P}^+$ is local.

Proof. For $\mathbb{P}_{A,B,C}$, the result is precisely [37, Lemma 32]. For $\mathbb{P}_A$, just observe that a play is copycat iff all its views are. $\square$

So locality is the first property we need to require of our game settings with views. The second property has to do with projections, e.g., $\delta_2: \mathbb{P}_{A,B,C} \rightarrow \mathbb{P}_{A,B}$. It essentially says that any morphism $v \rightarrow \delta_2(u)$ with $v \in \mathbb{V}_{A,B}$ and $u \in \mathbb{P}_{A,B,C}$ factors ‘canonically’ through some $\delta_2(w)$ with $w \in \mathbb{V}_{A,B,C}$, where $\mathbb{V}_{A,B,C}$ denotes the full subcategory of $\mathbb{P}_{A,B,C}$ projecting to $\mathbb{V}_{A,C}$ (or otherwise said, $\mathbb{V}_{A,B,C} = \mathbb{P}_{A,B,C} \times_{\mathbb{P}_{A,C}} \mathbb{V}_{A,C}$). In order to define such canonicity, we appeal to the theory of analytic functors [21, 38, 39].

Definition 4.9 (Weber [38, 39]). A functor $T: \mathbb{C} \rightarrow \mathbb{D}$ admits *generic factorisations* relative to an object $d \in \mathbb{D}$ iff any $f: d \rightarrow Tc$ admits a factorisation as below left

$$\begin{array}{ccc} d & & Tb \\ g \downarrow & \searrow f & \uparrow Tg' \\ Ta & \xrightarrow{Th} & Tc \end{array} \quad \begin{array}{ccc} d & \xrightarrow{g'} & Tb \\ g \downarrow & \nearrow Tk & \downarrow Th' \\ Ta & \xrightarrow{Th} & Tc \end{array}$$

such that for all commuting squares as the exterior above right, there exists a lifting k as shown making the diagram commute, or more precisely such that $g' = Tk \circ g$ and $h = h'k$. The middle object a is called the *arity* of f – all such factorisations share the same a up to isomorphism.

For all subcategories $\mathbb{B} \hookrightarrow \mathbb{C}$ and $\mathbb{E} \hookrightarrow \mathbb{D}$, a functor $\mathbb{C} \rightarrow \mathbb{D}$ admitting generic factorisations relative to all objects of $\mathbb{E}$ with arities in $\mathbb{B}$ is called $(\mathbb{B}, \mathbb{E})$ -*analytic* [13].

Definition 4.10. A game setting $(A, \mathbb{P})$ equipped with full embeddings $i_{A,B}: \mathbb{V}_{A,B} \hookrightarrow \mathbb{P}_{A,B}$ is *view-analytic* when δ_2 is $(\mathbb{V}_{A,B,C}, \mathbb{V}_{A,B})$ -analytic and δ_0 is $(\mathbb{V}_{A,B,C}, \mathbb{V}_{B,C})$ -analytic.

Proposition 4.11. Tsukada and Ong's game setting [37] is view-analytic.

Proof. This follows from [36, Lemma 36]. $\square$

We may now state our main result about innocence:

Definition 4.12. An *innocent game setting* is a game setting $(A, \mathbb{P})$ equipped with full embeddings $i_{A,B}: \mathbb{V}_{A,B} \hookrightarrow \mathbb{P}_{A,B}$, which is both local and view-analytic.

Theorem 4.13. In any innocent game setting, innocent strategies form a subcategory.

Again, we defer the proof to the more technical Section B.6.

4.2 Prefix-based innocence

In the previous section, we have shown that innocent strategies behave well in any innocent game setting. However, our only concrete example of an innocent game setting for now is Tsukada and Ong's $\mathbb{P}^+$. There is in fact a further example, given by enriching arenas with bracketing information and restricting $\mathbb{P}_{A,B}^+$ to well-bracketed plays [37, Section VII]. This shows that innocence is stable under cs-composition. How about pb-composition? As mentioned before, innocence is not stable under pb-composition in general [15, Section 3.7.2], unless one restricts to deterministic strategies. In an attempt to better understand this phenomenon, we first move in this section from cs-composition to ps-composition, and prove that innocence remains stable. In the next section, we will explain why this does not carry over to pb-composition, although, as is well-known, it does on deterministic strategies.

We here proceed by first defining innocent, prefix-based strategies in an extended framework and then showing that our definition agrees with the standard one (which is only defined on

boolean behaviours). We then show that ps-innocent strategies include copycats and are closed under composition.

Definition 4.14. Consider game settings $(A, \mathbb{P}^+)$ and $(A, \mathbb{P})$ with the same set of arenas and $\mathbb{V}$ making $\mathbb{P}^+$ innocent, further equipped with a componentwise identity-on-objects natural embedding $k: \mathbb{P} \hookrightarrow \mathbb{P}^+$. Let a presheaf on $\mathbb{P}_{A,B}$ be *innocent via $\mathbb{P}^+$* , or *$\mathbb{P}^+$ -innocent*, iff it is in the essential image of $\widehat{\mathbb{V}_{A,B}} \xrightarrow{\Pi_{i_{A,B}}} \widehat{\mathbb{P}_{A,B}^+} \xrightarrow{\Delta_{k_{A,B}}} \widehat{\mathbb{P}_{A,B}}$. Similarly, let a presheaf on $\mathbb{P}_{A,A}$ be *$\mathbb{P}^+$ -copycat* iff it is in the essential image of $1 \cong \widehat{0} \xrightarrow{\Pi_1} \widehat{\mathbb{P}_A^+} \xrightarrow{\Sigma_{i_0}} \widehat{\mathbb{P}_{A,A}^+} \xrightarrow{\Delta_{k_{A,A}}} \widehat{\mathbb{P}_{A,A}}$.

Proposition 4.15. *A strategy in the standard HON sense is innocent iff it is $\mathbb{P}^+$ -innocent. It is copycat iff it is $\mathbb{P}^+$ -copycat.*

Proof. We prove the more specific statement that the innocent strategy associated to any boolean behaviour (i.e., in this case, a behaviour in the image of r_1) in standard HON games is given up to isomorphism by the functor $\Delta_{k_{A,B}} \circ \Pi_{i_{A,B}}$. Similarly, the copycat on A is given by the image of the unique element of 1 under $\Delta_{k_{A,A}} \circ \Sigma_{i_0} \circ \Pi_1$. For innocence, consider any boolean behaviour B . As is standard, right extension may in this case be computed as a conjunction, which entails that B is mapped to

$$B'(p) = \bigwedge_{\{v \in \mathbb{V}_{A,B} \mid \mathbb{P}_{A,B}^+(v,p) \neq \emptyset\}} B(v),$$

i.e., p is accepted iff B accepts all its views. The case of copycats is similar. $\square$

By the proposition, we may understand ps-innocence through cs-innocence. Let us now state the transfer result.

Proposition 4.16. *In the setting of Definition 4.14, if all naturality squares*

$$\begin{array}{ccc} \mathbb{P}_{A,B,C} & \longrightarrow & \mathbb{P}_{A,C} \\ \downarrow & & \downarrow \\ \mathbb{P}_{A,B,C}^+ & \longrightarrow & \mathbb{P}_{A,C}^+ \end{array} \quad \text{and} \quad \begin{array}{ccc} \mathbb{P}_A & \longrightarrow & \mathbb{P}_{A,A} \\ \downarrow & & \downarrow \\ \mathbb{P}_A^+ & \longrightarrow & \mathbb{P}_{A,A}^+ \end{array}$$

are pullbacks, then $\mathbb{P}^+$ -innocent strategies are closed under composition and comprise $\mathbb{P}^+$ -copycat strategies.

The proof is in Section B.7. Of course, both hypotheses are satisfied in the case of Hyland-Ong games.

4.3 Boolean innocence

We finally consider boolean innocence. As mentioned before, innocent boolean strategies are not closed under composition. One usually either imposes a further determinism constraint, or relaxes the innocence constraint. It might be instructive to see how trying to derive the boolean case from the set-based one using our methods directly points to the problem.

Indeed, suppose given any innocent game setting $(A, \mathbb{P}, \mathbb{V}, i)$. We would like to show that two boolean polynomial functors, say

$$\overline{P}_1, \overline{P}_2: \widehat{\mathbb{V}_{A,B} + \mathbb{V}_{B,C}} \rightarrow \widehat{\mathbb{P}_{A,C}}$$

coincide. Here $\overline{P}_1$ is innocentisation followed by composition and $\overline{P}_2$ is the same, followed by another pass of innocentisation (and $\mathbb{P}$ could very well be replaced by $\mathbb{P}^+$). If we could show that each $\overline{P}_i$ factors as

$$\begin{array}{ccc} \widehat{\mathbb{V}_{A,B} + \mathbb{V}_{B,C}} & \xrightarrow{P_i} & \widehat{\mathbb{P}_{A,C}} \\ \uparrow r_i & & \downarrow l_i \\ \widehat{\mathbb{V}_{A,B} + \mathbb{V}_{B,C}} & \xrightarrow{\overline{P}_i} & \widehat{\mathbb{P}_{A,C}} \end{array}$$

where r_i and l_i are as in Proposition 2.16, then because we have already shown that $P_1 \cong P_2$, we would automatically get $\overline{P}_1 \cong \overline{P}_2$ as desired.

Now, in all cases, both functors have the form

$$\widehat{\mathbb{V}_{A,B} + \mathbb{V}_{B,C}} \xrightarrow{\overline{Q}} \widehat{\mathbb{P}_{A,B,C}} \xrightarrow{\overline{\Sigma_{\delta_1}}} \widehat{\mathbb{P}_{A,C}} \xrightarrow{\overline{R}} \widehat{\mathbb{P}_{A,C}},$$

where $\overline{Q}$ and $\overline{R}$ are only composed of $\overline{\Delta}$ s and $\overline{\Pi}$ s. Using the fact that $\overline{\Pi}$ s and $\overline{\Delta}$ s commute with r_i , we may thus hope to be able to prove the desired factorisation like so:

$$\begin{array}{ccccccc} \widehat{\mathbb{V}_{A,B} + \mathbb{V}_{B,C}} & \xrightarrow{\overline{Q}} & \widehat{\mathbb{P}_{A,B,C}} & \xrightarrow{\overline{\Sigma_{\delta_1}}} & \widehat{\mathbb{P}_{A,C}} & \xrightarrow{\overline{R}} & \widehat{\mathbb{P}_{A,C}} \\ \uparrow r_i & & \uparrow r_i & ? & \uparrow r_i & & \uparrow r_i \\ \widehat{\mathbb{V}_{A,B} + \mathbb{V}_{B,C}} & \xrightarrow{\overline{Q}} & \widehat{\mathbb{P}_{A,B,C}} & \xrightarrow{\overline{\Sigma_{\delta_1}}} & \widehat{\mathbb{P}_{A,C}} & \xrightarrow{\overline{R}} & \widehat{\mathbb{P}_{A,C}} \end{array}$$

where the right triangle commutes up to isomorphism because $l_i \dashv r_i$ is a reflection. The only problematic square is the marked one. And indeed, if the considered boolean strategies, say X_1 and X_2 , are non-deterministic, $\overline{\Sigma_{\delta_1}}(r_i(\overline{Q}[X_1, X_2]))$ may accept some $p \in \mathbb{P}_{A,C}$ in more than one way (see Harmer [15, 3.7.2]), which readily makes it non-isomorphic to any presheaf in the image of r_i .

Remark 4. *Exactly the same argument explains why boolean composition cannot agree with set-based composition in general (as was noted in Remark 2).*

Standardly, the problem is overcome by restricting to deterministic strategies, for which the marked square commutes. Again, this works indifferently in the ps or cs settings.

Finally, we proceed similarly in the case of copycat strategies. In this case, the problematic square

$$\begin{array}{ccc} \widehat{\mathbb{P}_A} & \xrightarrow{\Sigma_{i_0}} & \widehat{\mathbb{P}_{A,A}} \\ \uparrow r_1 & & \uparrow r_1 \\ \widehat{\mathbb{P}_A} & \xrightarrow{\overline{\Sigma_{i_0}}} & \widehat{\mathbb{P}_{A,A}} \end{array}$$

does commute, because the involved colimits are coproducts which are either empty or singleton:

Proposition 4.17. *In any innocent game setting $(A, \mathbb{P}, \mathbb{V}, i)$, boolean copycat strategies are innocent. This extends to the setting of Proposition 4.16, so that, as is standard, copycats are innocent pb-strategies.*

Proof. By discrete fibredness of i_0 , for all $p \in \mathbb{P}_{A,A}$, the comma category p/i_0 is either empty or connected. But $\overline{\Sigma_{i_0}}(r_1(X))(p)$ is the colimit of

$$(p/i_0)^{op} \xrightarrow{\text{cod}^{op}} \mathbb{P}_A^{op} \xrightarrow{X} 2 \xrightarrow{r} \text{Set},$$

which is thus $\emptyset$ or 1, hence isomorphic to $r_1(\overline{\Sigma_{i_0}}(X))(p)$. $\square$

5 Conclusion and perspectives

We have introduced game settings and their innocent variant, a categorical framework for game semantics, with the hope of facilitating the construction of new game models. A lot remains to be done, starting with the incorporation of further instances. The saturated

view of AJM games (Section 3.4) seems at hand, but will involve significantly more advanced category theory, as Street fibrations and stacks will replace discrete fibrations and sheaves. Less obvious is the treatment of more exotic game models [7, 11, 16, 28–31, 34], notably those based on event structures. Another direction is categorification: instead of reasoning up to isomorphism, we could refine our point of view and prove that games and strategies in fact form a bicategory, as, e.g., in Rideau and Winskel [34]. Finally, beyond game models, we should investigate game semantics, i.e., the correspondence with operational semantics, as initiated in Eberhart and Hirschowitz [12] in a different setting.

References

- [1] IEEE. *Proc. 12th Logic in Comp. Sci.* 1997.
- [2] IEEE. *Proc. 14th Logic in Comp. Sci.* 1999.
- [3] IEEE. *Proc. 30th Logic in Comp. Sci.* 2015.
- [4] Samson Abramsky. 2003. Sequentiality vs. Concurrency In Games And Logic. *MSCS* 13, 4 (2003), 531–565. <https://doi.org/10.1017/S0960129503003980>
- [5] Samson Abramsky, Kohei Honda, and Guy McCusker. 1998. A Fully Abstract Game Semantics for General References. In *LICS*. IEEE, 334–344.
- [6] Samson Abramsky, Radha Jagadeesan, and Pasquale Malacaria. 2000. Full Abstraction for PCF. *Inf. and Comp.* 163, 2 (2000), 409–470. <https://doi.org/10.1006/inco.2000.2930>
- [7] Samson Abramsky and Paul-André Melliès. 1999. Concurrent Games and Full Completeness, See [2], 431–442. <https://doi.org/10.1109/LICS.1999.782638>
- [8] Patrick Baillot, Vincent Danos, Thomas Ehrhard, and Laurent Regnier. 1997. Believe it or not, AJM's Games Model is a Model of Classical Linear Logic, See [1], 68–75. <https://doi.org/10.1109/LICS.1997.614933>
- [9] Andreas Blass. 1972. Degrees of indeterminacy in games. *Fundamenta Mathematica* LXXVII (1972), 151–162.
- [10] Andreas Blass. 1992. A game semantics for linear logic. *Annals of Pure and Applied Logic* 56, 1–3 (1992), 183–220. [https://doi.org/10.1016/0168-0072\(92\)90073-9](https://doi.org/10.1016/0168-0072(92)90073-9)
- [11] Simon Castellani, Pierre Clairambault, and Glynn Winskel. 2015. The parallel intensionally fully abstract games model of PCF, See [3].
- [12] Clovis Eberhart and Tom Hirschowitz. 2017. Game semantics as a singular functor, and definability as geometric realisation. (2017). <https://hal.archives-ouvertes.fr/hal-01527171> Preprint.
- [13] Richard Garner and Tom Hirschowitz. 2017. Shapely monads and analytic functors. *Journal of Logic and Computation* (2017), exx029. <https://doi.org/10.1093/logcom/exx029>
- [14] René Guitart. 1980. Relations et carrés exacts. *Annales des Sciences Mathématiques du Québec* 4, 2 (1980), 103–125.
- [15] Russell Harmer. 1999. *Games and Full Abstraction for Nondeterministic Languages*. Ph.D. Dissertation. Imperial College, University of London.
- [16] Russell Harmer, Martin Hyland, and Paul-André Melliès. 2007. Categorical Combinatorics for Innocent Strategies. In *Proc. 22nd Logic in Comp. Sci.* IEEE, 379–388. <https://doi.org/10.1109/LICS.2007.14>
- [17] Russell Harmer and Guy McCusker. 1999. A Fully Abstract Game Semantics for Finite Nondeterminism, See [2], 422–430. <https://doi.org/10.1109/LICS.1999.782637>
- [18] Tom Hirschowitz. 2014. Full abstraction for fair testing in CCS (expanded version). *LMCS* 10, 4 (2014). [https://doi.org/10.2168/LMCS-10\(4:2\)2014](https://doi.org/10.2168/LMCS-10(4:2)2014)
- [19] Tom Hirschowitz and Damien Pous. 2011. Innocent strategies as presheaves and interactive equivalences for CCS. In *Proc. of Interaction and Concurrency Experience*. EPTCS, 2–24. <https://doi.org/10.4204/EPTCS.59.2>
- [20] J. M. E. Hyland and C.-H. L. Ong. 2000. On Full Abstraction for PCF: I, II, and III. *Inf. and Comp.* 163, 2 (2000), 285–408. <https://doi.org/10.1006/inco.2000.2917>
- [21] André Joyal. 1986. Foncteurs analytiques et espèces de structure. In *Combinatoire énumérative (Montréal 1985) (LNM)*, Vol. 1234. Springer, 126–159.
- [22] G. M. Kelly. 1989. Elementary observations on 2-categorical limits. *Bulletin of the Australian Mathematical Society* 39 (1989), 301–317.
- [23] Joachim Kock. 2009. Notes on polynomial functors. (2009). <http://mat.uab.es/~kock/cat/polynomial.pdf> Working notes.
- [24] James Laird. 1997. Full Abstraction for Functional Languages with Control, See [1], 58–67. <https://doi.org/10.1109/LICS.1997.614931>
- [25] Saunders Mac Lane. 1998. *Categories for the Working Mathematician* (2nd ed.). Number 5 in Graduate Texts in Mathematics. Springer.
- [26] Guy McCusker. 1996. *Games and Full Abstraction for a Functional Metalanguage with Recursive Types*. Ph.D. Dissertation. Imperial College, University of London.
- [27] Guy McCusker. 2000. Games and Full Abstraction for FPC. *Inf. and Comp.* 160, 1–2 (2000), 1–61. <https://doi.org/10.1006/inco.1999.2845>
- [28] Paul-André Melliès. 2004. Asynchronous games 2: the true concurrency of innocence. In *Proc. 15th International Conference on Concurrency Theory (LNCS)*, Vol. 3170. Springer, 448–465. https://doi.org/10.1007/978-3-540-28644-8_29
- [29] Paul-André Melliès. 2005. Asynchronous Games 4: A Fully Complete Model of Propositional Linear Logic. In *Proc. 20th Logic in Comp. Sci.* IEEE, 386–395. <https://doi.org/10.1109/LICS.2005.6>
- [30] Paul-André Melliès. 2012. Game Semantics in String Diagrams. In *Proc. 27th Logic in Comp. Sci.* IEEE, 481–490. <https://doi.org/10.1109/LICS.2012.58>
- [31] Paul-André Melliès and Samuel Mimram. 2007. Asynchronous Games: Innocence Without Alternation. In *Proc. 19th International Conference on Concurrency Theory (LNCS)*, Vol. 4703. Springer, 395–411. https://doi.org/10.1007/978-3-540-74407-8_27
- [32] Andrzej S. Murawski and Nikos Tzevelekos. 2016. Nominal Game Semantics. *Foundations and Trends in Programming Languages* 2, 4 (2016), 191–269. <https://doi.org/10.1561/25000000017>
- [33] Hanno Nickau. 1994. Hereditarily Sequential Functionals. In *Proc. Logical Foundations of Computer Science (LNCS)*, Vol. 813. Springer, 253–264. https://doi.org/10.1007/3-540-58140-5_25
- [34] Sylvain Rideau and Glynn Winskel. 2011. Concurrent Strategies. In *Proc. 26th Logic in Comp. Sci.* IEEE, 409–418. <https://doi.org/10.1109/LICS.2011.13>
- [35] Emily Riehl. 2014. *Categorical Homotopy Theory*. Number 24 in New Mathematical Monographs. Cambridge University Press.
- [36] 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>
- [37] Takeshi Tsukada and C.-H. Luke Ong. 2015. Nondeterminism in Game Semantics via Sheaves. See [3].
- [38] Mark Weber. 2004. Generic morphisms, parametric representations and weakly cartesian monads. *Theory and Applications of Categories* 13 (2004), 191–234.
- [39] Mark Weber. 2007. Familial 2-functors and parametric right adjoints. *Theory and Applications of Categories* 18, 22 (2007), 665–732.

A Categorical background

A.1 Presheaves, restriction, extensions

This section is a brief introduction to the description of both adjoints to restriction, $\sum_F$ and $\prod_F$ (see Notation 1), in terms of coends and ends, respectively.

Consider the left adjoint first, $\sum_F$. It is determined up to canonical isomorphism by the *coend*

$$\sum_F(X)(d) \cong \int^{c \in \mathcal{C}} X(c) \times \mathcal{D}(d, F(c)),$$

which we abbreviate as

$$\int^{c \in \mathcal{C}} X(c) \times [d, F(c)], \quad (4)$$

using the standard bracket notation for hom-sets (the ambient category should be clear from context). It will be enough to know that it is the quotient of the corresponding coproduct $\coprod_{c \in \mathcal{C}} X(c) \times \mathcal{D}(d, F(c))$ by the smallest equivalence relation $\sim$ satisfying

$$(x, F(f) \circ g) \sim (x \cdot f, g)$$

for all $x \in X(c')$, $g \in \mathcal{D}(d, F(c))$ and $f \in \mathcal{C}(c, c')$. Graphically, this relation may be visualised by the following commuting ‘diagram’:

$$\begin{array}{ccc} & 1 & \\ x \cdot f \swarrow & & \searrow x \\ X(c) & \xleftarrow{X(f)} & X(c') \\ & \xleftarrow{F(f)} & \\ F(c) & \xrightarrow{F(f)} & F(c') \\ & \nwarrow g & \nearrow F(f) \circ g \\ & d & \end{array} \quad (5)$$

We see that f acts in different directions, on $X(c')$ and $[d, F(c)]$, and the equivalence relation takes this into account.

The right adjoint works dually: it is determined up to canonical isomorphism by the *end*

$$\prod_F(X)(d) \cong \int_{c \in \mathcal{C}} \text{Set}(\mathcal{D}(F(c), d), X(c)),$$

which we abbreviate as

$$\int_{c \in \mathcal{C}} [[F(c), d], X(c)]. \quad (6)$$

In this case, the end admits a very simple description: the assignment $c \mapsto [F(c), d]$ defines a presheaf on $\mathcal{C}$, and the end is nothing but the set of natural transformations from this presheaf to X , so that we have

$$\prod_F(X)(d) \cong \widehat{\mathcal{C}}([F(-), d], X)$$

(which is in fact an instance of the universal property for ends [22]).

Still, it may be instructive to realise that such natural transformations are equally the elements α of the product

$$\prod_{c \in \mathcal{C}} [[F(c), d], X(c)]$$

satisfying a dual condition to (5), namely commutativity of

$$\begin{array}{ccccc} & & 1 & & \\ \alpha_c(g \circ F(f)) \swarrow & & & \searrow \alpha_{c'}(g) & \\ X(c) & \xleftarrow{X(f)} & & & X(c') \\ F(c) & \xrightarrow{F(f)} & & & F(c') \\ g \circ F(f) \searrow & & & \swarrow g & \\ & & d & & \end{array} \quad (7)$$

for all $f: c \rightarrow c'$ and $g: F(c') \rightarrow d$, which is equivalent to naturality, i.e., commutativity of the square

$$\begin{array}{ccc} [F(c'), d] & \xrightarrow{\alpha_{c'}} & X(c') \\ [F(f), d] \downarrow & & \downarrow X(f) \\ [F(c), d] & \xrightarrow{\alpha_c} & X(c). \end{array}$$

Beware: unlike in (5), where an element of the coend is a pair of arrows (e.g., $(x \cdot f, g)$), an element of the end is here a natural family of arrows $1 \rightarrow X(c)$ indexed by all $g: F(c) \rightarrow d$.

A.2 Fibrations

This section is a minimal introduction to (Grothendieck) fibrations and their relation to discrete fibrations.

Definition A.1. Consider any functor $p: \mathcal{E} \rightarrow \mathcal{C}$. A morphism $u: x \rightarrow y$ in $\mathcal{E}$ is *cartesian* when for all $v: z \rightarrow y$ in $\mathcal{E}$ and $k: p(z) \rightarrow p(x)$ making the bottom triangle commute

$$\begin{array}{ccccc} z & & & & \\ & \searrow v & & & \\ & & x & \xrightarrow{u} & y \\ & \swarrow w & & & \\ p(z) & & & & \\ & \searrow p(v) & & & \\ & & p(x) & \xrightarrow{p(u)} & p(y) \\ & \swarrow k & & & \end{array}$$

there exists a unique $w: z \rightarrow x$ such that $p(w) = k$ and $u \circ w = v$.

The functor p is a *fibration* when for all objects $e \in \mathcal{E}$ and morphisms $f: c \rightarrow p(e)$ there exists a cartesian morphism $u: e' \rightarrow e$ such that $p(u) = f$. Such a morphism is called a *cartesian lifting* of e along f .

Clearly, the discrete fibrations of Definition 2.8 are identical to fibrations whose cartesian liftings are unique, or equivalently whose *fibres* are discrete. The *fibre* of p at any object $c \in \mathcal{C}$ is the subcategory of $\mathcal{E}$ spanning objects in $p^{-1}(c)$ and morphisms in $p^{-1}(id_c)$.

A.3 Exact squares

An essential tool in our proofs will be Guitart's theory of *exact squares* [14], which we now recall.

Definition A.2. A *square* is a natural transformation

$$\begin{array}{ccc} A & \xrightarrow{T} & B \\ S \downarrow & \xRightarrow{\varphi} & \downarrow V \\ C & \xrightarrow{U} & D. \end{array} \quad (8)$$

of small categories and functors.

Any square yields by restriction a square as in the middle below, and so by adjunction (the so-called mate calculus) two further squares as on the left and right:

$$\begin{array}{ccc} \widehat{A} & \xleftarrow{\Delta_T} & \widehat{B} \\ \Sigma_S \downarrow & \Downarrow \Sigma_\varphi & \downarrow \Sigma_V \\ \widehat{C} & \xleftarrow{\Delta_U} & \widehat{D} \end{array} \quad \begin{array}{ccc} \widehat{A} & \xleftarrow{\Delta_T} & \widehat{B} \\ \Delta_S \uparrow & \Delta_\varphi \swarrow & \uparrow \Delta_V \\ \widehat{C} & \xleftarrow{\Delta_U} & \widehat{D} \end{array} \quad \begin{array}{ccc} \widehat{A} & \xrightarrow{\Pi_T} & \widehat{B} \\ \Delta_S \uparrow & \xRightarrow{\Pi_\varphi} & \uparrow \Delta_V \\ \widehat{C} & \xrightarrow{\Pi_U} & \widehat{D}. \end{array}$$

Definition A.3. A square φ is *exact* if and only if Σ_φ is an isomorphism.

The following is well-known:

Property 1. A square φ is exact if and only if Π_φ is an isomorphism.

The notion of exactness thus corresponds to commutation up to isomorphism of the two diagrams on the sides above.

Remark 5. In an attempt to spare the reader a few headaches trying to remember the directions of Σ 's, Δ 's and Π 's, as well as of the induced natural transformations, let us share our mnemonic: the original transformation φ points to the Σ 'ed functor if we are to reason about Σ_φ , and from the Π 'ed functor if we are to reason about Π_φ . Furthermore, induced natural transformations flow along left adjoints and against right adjoints (hence along Σ 's for Σ_φ , and against Π 's for Π_φ).

Notation 3. In the proofs, we will manipulate diagrams of restriction and its adjoints. In order to reduce notational clutter, hats will be omitted and arrows will point in the direction of underlying functors. This means that Δ arrows will point in the “wrong” direction, in the sense that if $F: X \rightarrow Y$, we will write $Y \xleftarrow{\Delta_F} X$ for $\widehat{Y} \xrightarrow{\Delta_F} \widehat{X}$.

Let us conclude this section with a few basic lemmas about exact squares. First, we recall that a square (8) is a *comma square* when it is the terminal lax cone to (U, V) . More formally, it is a terminal object in the category whose objects are spans $C \xleftarrow{S'} A' \xrightarrow{T'} B$ equipped with a natural transformation $\varphi': US' \rightarrow VT'$, and whose morphisms $(A', S', T', \varphi') \rightarrow (A'', S'', T'', \varphi'')$ are functors $m: A' \rightarrow A''$ such that $S'' \circ m = S'$, $T'' \circ m = T'$, and $\varphi'' \circ m = \varphi'$ (where $\circ$ here denotes whiskering). *Cocomma squares* are defined exactly dually, i.e., with D, U and V varying instead of A, S , and T . Guitart shows:

Lemma A.4. Any comma (resp. cocomma) square is exact.

Proof. See Guitart [14], Examples 1.14.2 and 1.14.3. $\square$

Lemma A.5. For any functor $f: A \rightarrow B$, the square below left is exact; furthermore, the square below right is exact iff f is fully faithful:

What's in a game?

$$\begin{array}{ccc} A & \xrightarrow{f} & B \\ \parallel & \xRightarrow{id_f} & \parallel \\ A & \xrightarrow{f} & B \end{array} \quad \begin{array}{ccc} A & \xRightarrow{id_f} & A \\ \parallel & \xRightarrow{id_f} & \downarrow f \\ A & \xrightarrow{f} & B \end{array}$$

Proof. See Guitart [14], Examples 1.14.1 and 1.14.4. $\square$

Lemma A.6. *Exact squares are stable under horizontal composition.*

Proof. See Guitart [14], Theorem 1.8. $\square$

Lemma A.7. *Any square (8) in which φ and S are identities and V is fully faithful is exact.*

Proof. We obtain the given square as the horizontal composite

$$\begin{array}{ccccc} & & T & & \\ & \nearrow & & \searrow & \\ A & \xrightarrow{T} & B & \xRightarrow{id} & B \\ \parallel & & \parallel & & \downarrow V \\ A & \xrightarrow{T} & B & \xrightarrow{V} & C \\ & \searrow & & \nearrow & \\ & & U & & \end{array}$$

which is exact by Lemma A.6, because both squares are exact by Lemma A.5. $\square$

Guitart provides the following necessary and sufficient criterion for exactness:

Lemma A.8 (Guitart [14], Theorem 1.2 (zig-zag criterion)). *A square as in (8) is exact if and only if, for all objects c of C , d of D , and morphisms $f: Uc \rightarrow Vd$:*

- *there is an object a of A such that f factors as $Uc \xrightarrow{Ug} USa \xrightarrow{\varphi_a} VTa \xrightarrow{Vh} Vd$,*
- *and for all two such factorisations $Uc \xrightarrow{Ug} USa \xrightarrow{\varphi_a} VTa \xrightarrow{Vh} Vd$ and $Uc \xrightarrow{Ug'} USa' \xrightarrow{\varphi_{a'}} VTa' \xrightarrow{Vh'} Vd$, there exists a commuting diagram (which Guitart calls a lantern)*

$$\begin{array}{ccccccc} & & c & & & & \\ & \nearrow g & & \searrow g' & & \nearrow & \\ Sa & \xleftarrow{St_0} \cdot \xleftarrow{St_1} \dots \xleftarrow{St_{n-1}} \cdot \xleftarrow{St_n} & Sa' & & & & \\ Ta & \xleftarrow{Tt_0} \cdot \xleftarrow{Tt_1} \dots \xleftarrow{Tt_{n-1}} \cdot \xleftarrow{Tt_n} & Ta' & & & & \\ & \searrow h & & \nearrow h' & & \searrow & \\ & & d & & & & \end{array}$$

Let us conclude this section with an apparently new (though not surprising) result showing that pullbacks of (op)fibrations are exact squares (not in the same direction!).

Lemma A.9. *Any pullback square (8) with V a fibration is exact. Similarly, if U is an opfibration, then (8) is exact.*

Proof. We proceed by the zig-zag criterion. Consider any $f: Uc \rightarrow Vb$. Let us first establish existence of the desired factorisation, by considering any cartesian lifting $l: b_0 \rightarrow b$ of b along f . We obtain $Uc = Vb_0$, hence by universal property of pullback (By the standard categorification/nerve adjunction, limits of categories are computed as limits of underlying simplicial sets, hence pointwise.)

a unique a such that $c = Sa$ and $b_0 = Ta$. The original f thus factors as

$$Uc \xrightarrow{U(id)} USa \xrightarrow{id} VTa \xrightarrow{Vl} Vb.$$

We now need to show that any factorisation is connected to this one by some lantern. So consider any factorisation of f as

$$Uc \xrightarrow{Ug} USa' \xrightarrow{id} VTa' \xrightarrow{Vh} Vb.$$

Because fibrations are stable under pullback, S is a fibration, so we may pick a cartesian lifting, say $k: a'' \rightarrow a'$, of a' along g , so that in particular $Sa'' = c$ and $Sk = g$. Now, we have

$$VTa'' = USa'' = Uc = USa = VTa$$

hence a commuting triangle as below left

$$\begin{array}{ccc} VTa'' & \xrightarrow{VTk} & VTa' \xrightarrow{Vh} Vb \\ \parallel & \nearrow VI & \\ VTa & & \end{array} \quad \begin{array}{ccc} Ta'' & \xrightarrow{Tk} & Ta' \xrightarrow{h} b \\ u \downarrow & \nearrow l & \\ Ta & & \end{array}$$

(because $VTk = Usk = Ug$ so $Vh \circ VTk = f = VI$). So by cartesian-ness of l , we obtain a (unique) u making the above right triangle commute. We thus obtain the following commuting diagrams:

$$\begin{array}{ccc} & c & \\ & \parallel & \searrow g \\ Sa & \xRightarrow{id} & Sa'' \xrightarrow{Sk} Sa' \\ & & \\ Ta & \xleftarrow{u} & Ta'' \xrightarrow{Tk} Ta' \\ & \searrow l & \nearrow h \\ & & b. \end{array}$$

But since $Vu = id$, we have that $U(id) = Vu$ so by universal property of pullback again there exists a unique $w: a'' \rightarrow a$ such that $Sw = id$ and $Tw = u$. The above diagram thus yields a lantern, as desired. $\square$

A.4 Distributivity

As we saw, exact squares ensure commutation of restriction with left (resp. right) extension. In some of our proofs, we will also encounter squares for which we will need left extension to commute with right extension. This is an instance of what we will call distributive squares, which we now introduce.

Consider any exact square as (8). The inverse of Σ_φ yields a natural transformation as below left

$$\begin{array}{ccc} \hat{A} & \xleftarrow{\Delta_T} & \hat{B} \\ \Sigma_S \downarrow & \xleftarrow{\Sigma_\varphi^{-1}} & \downarrow \Sigma_V \\ \hat{C} & \xleftarrow{\Delta_U} & \hat{D} \end{array} \quad \begin{array}{ccc} \hat{A} & \xleftarrow{\Pi_T} & \hat{C} \\ \Sigma_S \downarrow & \xleftarrow{\tilde{\varphi}} & \downarrow \Sigma_V \\ \hat{B} & \xleftarrow{\Pi_U} & \hat{D}, \end{array}$$

which by adjunction induces a natural transformation as on the right.

Definition A.10. An exact square is *distributive* when $\tilde{\varphi}$ is an isomorphism.

Here is a useful construction of distributive squares. Recalling Proposition 4.6, consider any functors $A \xrightarrow{S} B \xrightarrow{U} D$ with S a discrete fibration U fully faithful. Let $C = \text{el}(\prod_U(\text{sing}(S)))$ denote the category of elements of the right extension of $\text{sing}(S)$ along

U , and let V denote its projection to D . Because extension along a fully faithful functor is fully faithful, the counit of the adjunction $\Delta_U \dashv \prod_U$ is an isomorphism. This yields a pullback square

$$\begin{array}{ccc} A & \xrightarrow{T_{S,U}} & D/\prod_U(\text{sing}(S)) \\ S \downarrow & \lrcorner & \downarrow P\prod_U(\text{sing}(S)) \\ B & \xrightarrow{U} & D \end{array} \quad (9)$$

which is exact by Lemma A.9.

Definition A.11. Any square obtained in this way is called a *local pushforward square*.

Lemma A.12. *Local pushforward squares are distributive.*

Proof. Consider any local pushforward square (8). We prove the corresponding result for the equivalent categories of discrete fibrations. There, because V and hence S are discrete fibrations, Σ_S becomes post-composition with S , and similarly for Σ_V . So the result essentially states that for any discrete fibration $P: E \rightarrow A$,

$$\prod_U(SP) \cong \prod_U(S) \circ \prod_T(P) \cong V \circ \prod_T(P),$$

where the first isomorphism follows from a simple diagram chase showing that $\prod_U(S) \circ \prod_T(P)$ enjoys the universal property of $\prod_U(SP)$. $\square$

B Main proofs

B.1 Adequacy

Proof of Proposition 2.10. First, because $\prod_!$ is right adjoint to $\Delta_!$, it preserves the terminal object (which is the unique object of $\widehat{\mathcal{O}}$), hence maps 1 to the terminal presheaf on $\mathbb{P}_A$, defined to map any play in $\mathbb{P}_A$ to 1. So we reduce to showing that $\text{id}_A \cong \Sigma_{i_0}(1)$.

Discrete fibredness of i_0 entails that left extension may be expressed as a mere coproduct: we have, for any presheaf X on $\mathbb{P}_A$ and $p \in \mathbb{P}_{A,A}$,

$$\left(\sum_{i_0} X\right)(p) \cong \sum_{\{q \in \mathbb{P}_A \mid i_0(q)=p\}} X(q).$$

So in particular when $X = 1$ we get

$$\sum_{i_0}(1)(p) \cong \sum_{\{q \in \mathbb{P}_A \mid i_0(q)=p\}} 1 \cong \begin{cases} 1 & \text{if } p \in \text{Im}(i_0) \\ \emptyset & \text{otherwise,} \end{cases}$$

as desired. $\square$

Proof of Proposition 2.13. Starting from the end, consider Σ_{δ_1} . Thanks to discrete fibredness of δ_1 , we have for any presheaf X'' on $\mathbb{P}_{A,B,C}$ and $p \in \mathbb{P}_{A,C}$:

$$\sum_{\delta_1} X''(p) \cong \sum_{\{u \in \mathbb{P}_{A,B,C} \mid \delta_1(u)=p\}} X''(u).$$

The next step is to unfold the right extension $\prod_{[id, id]}$. In this case, $[id, id]$ is easily seen to be a discrete opfibration, which entails that right extension may be expressed as the product

$$\prod_{[id, id]} (X')(u) \cong \prod_{\{v \in (\mathbb{P}_{A,B,C} + \mathbb{P}_{A,B,C}) \mid [id, id](v)=u\}} X'(v),$$

i.e., unfolding X' as the copairing $[X'_0, X'_1]$, v as the pair (i, u') and using the constraint $[id, id](v) = [id, id](i, u') = u' = u$:

$$\prod_{[id, id]} [X'_0, X'_1](u) \cong \prod_{i \in \{0,1\}} (X'_i(u)) \cong X'_0(u) \times X'_1(u).$$

Finally, because $\Delta_{\delta_2 + \delta_0}[X_0, X_1] = [\Delta_{\delta_2}(X_0), \Delta_{\delta_0}(X_1)]$, we obtain that the candidate composition of Definition 2.12 maps any $[X_0, X_1]$ to

$$\begin{aligned} & \Sigma_{\delta_1}(\prod_{[id, id]}(\Delta_{\delta_2 + \delta_0}[X_0, X_1]))(p) \\ &= \sum_{\{u \in \mathbb{P}_{A,B,C} \mid \delta_1(u)=p\}} \prod_{[id, id]}[\Delta_{\delta_2}(X_0), \Delta_{\delta_0}(X_1)](u) \\ &= \sum_{\{u \in \mathbb{P}_{A,B,C} \mid \delta_1(u)=p\}} \Delta_{\delta_2}(X_0)(u) \times \Delta_{\delta_0}(X_1)(u) \\ &= \sum_{\{u \in \mathbb{P}_{A,B,C} \mid \delta_1(u)=p\}} X_0(\delta_2(u)) \times X_1(\delta_0(u)), \end{aligned}$$

which clearly coincides with Tsukada and Ong's definition. $\square$

B.2 Associativity

In this section, we prove Theorem 2.15. Let us start with an alternative description of composition, which relies on the following intermediate category:

Definition B.1. For any triple of arenas A, B, C , let $\mathbb{P}_{(A,B),(B,C)}$ denote the *lax colimit* [22] of

$$\mathbb{P}_{A,B} \longleftarrow \mathbb{P}_{A,B,C} \longrightarrow \mathbb{P}_{B,C},$$

i.e., the initial category equipped with natural transformations

$$\begin{array}{ccccc} \mathbb{P}_{A,B} & & \mathbb{P}_{A,B,C} & & \mathbb{P}_{B,C} \\ & \searrow \lambda & \downarrow & \swarrow \rho & \\ & & \mathbb{P}_{(A,B),(B,C)} & & \end{array}$$

The obtained category has as objects the disjoint union of objects from $\mathbb{P}_{A,B}$, $\mathbb{P}_{B,C}$, and $\mathbb{P}_{A,B,C}$. It inherits the corresponding morphisms, and has additional morphisms $\lambda_u: \delta_2 u \rightarrow u$ and $\rho_u: \delta_0 u \rightarrow u$ for all u satisfying the obvious naturality requirements:

$$\lambda_{u'} \circ \delta_2 f = f \circ \lambda_u \quad \text{and} \quad \rho_{u'} \circ \delta_0 f = f \circ \rho_u,$$

for all $f: u \rightarrow u'$ in $\mathbb{P}_{A,B,C}$. We have:

Proposition B.2. *Composition is isomorphic to the polynomial functor*

$$\mathbb{P}_{A,B} + \mathbb{P}_{B,C} \xrightarrow{\Pi} \mathbb{P}_{(A,B),(B,C)} \xleftarrow{\Delta} \mathbb{P}_{A,B,C} \xrightarrow{\Sigma} \mathbb{P}_{A,C}.$$

Proof. This will follow from Lemma A.4 by observing that the square

$$\begin{array}{ccc} \mathbb{P}_{A,B,C} + \mathbb{P}_{A,B,C} & \xrightarrow{\nabla} & \mathbb{P}_{A,B,C} \\ \delta_2 + \delta_0 \downarrow & \Longrightarrow & \downarrow \\ \mathbb{P}_{A,B} + \mathbb{P}_{B,C} & \longrightarrow & \mathbb{P}_{(A,B),(B,C)} \end{array}$$

is a cocomma square, where ∇ denotes the copairing $[id, id]$. Indeed, both categories have the same universal property, expressed differently. $\square$

Proving associativity thus reduces to showing that the perimeter of the top diagram of Figure 1 commutes up to isomorphism. In order to do this, we introduce the category $\mathbb{P}_{(A,B),(B,C),(C,D)}$, similar to $\mathbb{P}_{(A,B),(B,C)}$ but with four arenas, which is constructed as the lax colimit of

$$\begin{array}{ccccc} & & \mathbb{P}_{B,C} & & \\ & \nearrow & & \nwarrow & \\ \mathbb{P}_{A,B} & \leftarrow \mathbb{P}_{A,B,C} & & \mathbb{P}_{B,C,D} & \rightarrow \mathbb{P}_{C,D} \\ & \nwarrow & \mathbb{P}_{A,B,C,D} & \nearrow & \end{array}$$

Figure 1. Diagram for associativity and zoom into the bottom left heptagon

In the top diagram of Figure 1, both little squares commute up to isomorphism because the underlying squares do. It thus suffices to show that both heptagons commute up to isomorphism. Both cases are symmetric, so we only treat the bottom left one. We then need to introduce yet another category, $\mathbb{P}_{(A,B),(B,C),(C,D)}$, which is like $\mathbb{P}_{(A,B),(B,C),(C,D)}$, except that $\mathbb{P}_{B,C,D}$ is not decomposed into $\mathbb{P}_{B,C}$ and $\mathbb{P}_{C,D}$: it is the lax colimit of

$$\mathbb{P}_{A,B} \leftarrow \mathbb{P}_{A,B,C,D} \rightarrow \mathbb{P}_{B,C,D}.$$

These categories are related by full embeddings

$$\mathbb{P}_{(A,B),(B,C)} \hookrightarrow \mathbb{P}_{(A,B),(B,C),(C,D)} \hookrightarrow \mathbb{P}_{(A,B),(B,C,D)}.$$

The crucial reason why the heptagon commutes is:

Lemma B.3. *The square*

$$\begin{array}{ccc} \mathbb{P}_{A,B} + \mathbb{P}_{B,C,D} & \longrightarrow & \mathbb{P}_{(A,B),(B,C,D)} \\ \mathbb{P}_{A,B} + \delta_1 \downarrow & & \downarrow \mathbb{P}_{(A,B),\delta_1} \\ \mathbb{P}_{A,B} + \mathbb{P}_{B,D} & \xrightarrow{[\lambda,\rho]} & \mathbb{P}_{(A,B),(B,D)} \end{array}$$

is a local pushforward square, i.e., it is a pullback, $\mathbb{P}_{A,B} + \delta_1$ is a discrete fibration, $[\lambda,\rho]$ is fully faithful, and

$$\mathbb{P}_{(A,B),\delta_1} \cong \prod_{[\lambda,\rho]} (\mathbb{P}_{A,B} + \delta_1).$$

Proof sketch. Everything is direct, except the last isomorphism, which says that $\mathbb{P}_{(A,B),(B,C,D)}$ is a sheaf for the topology induced by the full embedding $\mathbb{P}_{A,B} + \mathbb{P}_{B,D} \hookrightarrow \mathbb{P}_{(A,B),(B,D)}$. This reduces to the fact that any $u \in \mathbb{P}_{A,B,C,D}$ is entirely determined by giving its projections $u_1 = \delta_2(\delta_2(u)) \in \mathbb{P}_{A,B}$ and $u_2 = \delta_0(\delta_2(u)) \in \mathbb{P}_{B,D}$, plus some $p \in \mathbb{P}_{B,C,D}$ such that $\delta_1(p) = u_2$, which holds by zipping.

In more detail, the considered topology has canonical sieves: for objects in the image of the embedding, the canonical sieve is just the total sieve; for $u \in \mathbb{P}_{A,B,D}$, the canonical sieve is generated by the cospan

$$\delta_2(u) \rightarrow u \leftarrow \delta_0(u).$$

Thus, the sheaf condition for $\mathbb{P}_{(A,B),(B,C,D)}$ reduces to saying that, implicitly coercing discrete fibrations $p: \mathbb{E} \rightarrow \mathbb{P}_{(A,B),(B,C)}$ into the corresponding presheaf $\text{sing}((\mathbb{E}, p))$, the restriction map

$$\mathbb{P}_{(A,B),(B,C,D)}(u) \rightarrow \mathbb{P}_{A,B}(\delta_2(u)) \times \mathbb{P}_{B,C,D}(\delta_0(u))$$

is an isomorphism. But $\mathbb{P}_{A,B}(\delta_2(u)) \cong 1$, so this further reduces to the restriction map

$$\mathbb{P}_{(A,B),(B,C,D)}(u) \rightarrow \mathbb{P}_{B,C,D}(\delta_0(u))$$

being an isomorphism. This is only non-trivial when $u \in \mathbb{P}_{A,B,D}$ and in that case it is equivalent to

$$\begin{aligned} \forall u \in \mathbb{P}_{A,B,D}, v \in \mathbb{P}_{B,C,D}, (\delta_0(u) = \delta_1(v)) \\ \Rightarrow \exists! w \in \mathbb{P}_{A,B,C,D}, \delta_2(w) = u \wedge \delta_0(w) = v, \end{aligned}$$

i.e., the left square of (3) being a pullback on objects. $\square$

This leads us to fill the heptagon as at the bottom of Figure 1. The top right square commutes by Lemmas A.12 and B.3. The top triangle commutes up to isomorphism by Lemma A.7, the bottom one because the underlying diagram commutes, and the bottom-right square by Lemma A.9.

Associativity finally follows from:

Lemma B.4. *The left square commutes up to isomorphism.*

Proof. By the classical limit formula for right extensions, given any presheaf X on $\mathbb{P}_{A,B} + \mathbb{P}_{(B,C),(C,D)}$,

- following the left and bottom arrows we obtain a presheaf X' mapping any $u \in \mathbb{P}_{A,B,C,D}$ to the limit of

$$\begin{aligned} ((\mathbb{P}_{A,B} + \mathbb{P}_{(B,C),(C,D)})/u)^{op} \\ \downarrow \\ (\mathbb{P}_{A,B} + \mathbb{P}_{(B,C),(C,D)})^{op} \\ \downarrow X \\ \text{Set}; \end{aligned}$$

- following the top and right arrows, we obtain a presheaf X' mapping any u to a similar limit with $\mathbb{P}_{A,B} + \mathbb{P}_{(B,C),(C,D)}$ replaced by $\mathbb{P}_{A,B} + \mathbb{P}_{B,C,D}$.

Now, the inclusion functor

$$(\mathbb{P}_{A,B} + \mathbb{P}_{B,C,D})/u \hookrightarrow (\mathbb{P}_{A,B} + \mathbb{P}_{(B,C),(C,D)})/u$$

is readily checked to be final [25, IX.3], so its opposite is initial and both limits are isomorphic. $\square$

B.3 Unitality

Left and right unitality are entirely symmetric, so we only treat one. First, we observe that, because $\iota_0: \mathbb{P}_A \rightarrow \mathbb{P}_{A,A}$ is a discrete

fibration, so is $i_0 + \mathbb{P}_{A,B} : \mathbb{P}_A + \mathbb{P}_{A,B} \rightarrow \mathbb{P}_{A,A} + \mathbb{P}_{A,B}$. Consider the diagram:

$$\begin{array}{ccccc}
 \emptyset + \mathbb{P}_{A,B} & \xrightarrow{\Pi} & \mathbb{P}_A + \mathbb{P}_{A,B} & \xrightarrow{\Sigma_{i_0 + \mathbb{P}_{A,B}}} & \mathbb{P}_{A,A} + \mathbb{P}_{A,B} \\
 \searrow \Pi & & \downarrow \Pi & & \downarrow \Pi_{[\lambda, \rho]} \\
 & & \mathbb{C} & \xrightarrow{\Sigma} & \mathbb{P}_{(A,A), (A,B)} \\
 \uparrow \Delta & & \uparrow \Delta & & \\
 \mathbb{P}_{A,B} & \xrightarrow{\Sigma} & \mathbb{P}_{A,A,B} & \xrightarrow{\Sigma} & \mathbb{P}_{A,B}
 \end{array}$$

where the functor $\mathbb{C} \rightarrow \mathbb{P}_{(A,A), (A,B)}$ is defined as

$$\text{el}\left(\prod_{[\lambda, \rho]} (\text{sing}(i_0 + \mathbb{P}_{A,B}))\right).$$

We thus obtain a local pushforward square which commutes up to isomorphism by Lemma A.12.

Let us now show that the bottom square is a pullback. By construction, an object c of $\mathbb{C}$ over any $u \in \mathbb{P}_{A,A,B}$ consists of an object $c_1 \in \mathbb{P}_A$ over $\delta_2(u)$ and an object $c_2 \in \mathbb{P}_{A,B}$ over $\delta_0(u)$. The latter has to be $\delta_0(u)$, and the former has to be $\delta_0(\delta_2(u))$ or equivalently $\delta_1(\delta_0(u))$, but it exists only when $u \cong i_0(\delta_0(u))$, i.e., when u is in the image of $\mathbb{P}_{A,B} \rightarrow \mathbb{P}_{A,A,B}$. Thus, the bottom square is exact by Lemma A.9.

Finally, the bottom row is isomorphic the identity on $\mathbb{P}_{A,B}$, so it only remains to prove that the composite $\mathbb{P}_{A,B} \cong \emptyset + \mathbb{P}_{A,B} \xrightarrow{\Pi} \mathbb{C} \xleftarrow{\Delta} \mathbb{P}_{A,B}$ is also isomorphic to the identity, which follows from the standard end formula for $\prod$ and the explicit description of lax colimits.

B.4 The boolean case

In this section, we show that our results in the presheaf case transfer to the boolean case. This will follow from showing that all the polynomial functors that we used commute with booleanisation. This is easy for left extensions and restrictions:

Proposition B.5. *For all functors $F : \mathbb{C} \rightarrow \mathbb{D}$, the following squares commute up to isomorphism.*

$$\begin{array}{ccc}
 [\mathbb{C}^{op}, \text{Set}] & \xrightarrow{\Sigma_F} & [\mathbb{D}^{op}, \text{Set}] \\
 \downarrow l_i & & \downarrow l_i \\
 [\mathbb{C}^{op}, 2] & \xrightarrow{\Sigma_F} & [\mathbb{D}^{op}, 2]
 \end{array}
 \quad
 \begin{array}{ccc}
 [\mathbb{C}^{op}, \text{Set}] & \xleftarrow{\Delta_F} & [\mathbb{D}^{op}, \text{Set}] \\
 \downarrow l_i & & \downarrow l_i \\
 [\mathbb{C}^{op}, 2] & \xleftarrow{\Delta_F} & [\mathbb{D}^{op}, 2]
 \end{array}$$

Proof. Commutation with restriction is obvious. For left extension, the reflection l_i being a left adjoint, preserves colimits, which is precisely what Σ_F computes. $\square$

Things do not work out so well with right extensions in general. In order to show that our polynomial functors commute with booleanisation, it is thus useful to delineate a sufficiently large class of limits that are preserved by l_i :

Lemma B.6. *The left adjoint l_i preserves products and the terminal object.*

Proof. That l_i preserves 1 is obvious. Now, a product $\prod_i X_i$ is non-empty just when each X_i is, hence just when $l(X_i) = 1$ for all i , i.e., when $\prod_i l(X_i) = 1$. Thus $l(\prod_i X_i) = 1$ iff $\prod_i l(X_i) = 1$, hence $l(\prod_i X_i) = \prod_i l(X_i)$. $\square$

Proposition B.7. *Booleanisation commutes with $\prod_F$, for any $F : \mathbb{C} \rightarrow \mathbb{D}$ such that for all $d \in \text{ob}(\mathbb{D})$ the comma category F/d is a coproduct of categories with a terminal object.*

Proof. Indeed, consider any such F . For any $X \in \widehat{\mathbb{C}}$ and $d \in \mathbb{D}$, letting $F/d \cong \sum_i \mathbb{D}_i^d$ with $\phi_i^d : F(c_i^d) \rightarrow d$ denoting the terminal object in $\mathbb{D}_i^d$, we have

$$\begin{aligned}
 l((\prod_F(X))(d)) &\cong l(\lim((F/d)^{op} \rightarrow \mathbb{C}^{op} \xrightarrow{X} \text{Set})) \\
 &\cong l(\prod_i X(c_i^d)) \\
 &\cong \prod_i l(X(c_i^d)) \quad (\text{by Lemma B.6}) \\
 &\cong (\lim((F/d)^{op} \rightarrow \mathbb{C}^{op} \xrightarrow{X} \text{Set} \xrightarrow{l} 2)) \\
 &\cong \prod_F(l \circ X)(d),
 \end{aligned}$$

as desired. $\square$

Proposition B.8. *The class of functors F such that $\prod_F$ commutes with booleanisation contains all functors $\nabla_{\mathbb{C}} : \mathbb{C} + \mathbb{C} \rightarrow \mathbb{C}$ and $! : \emptyset \rightarrow \mathbb{C}$, and it is stable under composition and coproduct (i.e., $F + G : \mathbb{C} + \mathbb{C}' \rightarrow \mathbb{D} + \mathbb{D}'$ is in it if F and G are).*

Proof. Easy consequences of the previous proposition. $\square$

Proof of Proposition 2.18. Both results state that two polynomial functors, say $\overline{P}_1$ and $\overline{P}_2$ between categories of the form $[\mathbb{C}^{op}, 2]$ are naturally isomorphic. But knowing that their set-versions, say P_1 and P_2 , are isomorphic, we may form

$$\begin{array}{ccc}
 [\mathbb{C}^{op}, \text{Set}] & \xrightleftharpoons[P_2]{P_1} & [\mathbb{D}^{op}, \text{Set}] \\
 \downarrow l_i & & \downarrow l_i \\
 [\mathbb{C}^{op}, 2] & \xrightleftharpoons[\overline{P}_2]{\overline{P}_1} & [\mathbb{D}^{op}, 2]
 \end{array}$$

In both cases, this diagram commutes serially by Propositions B.5 and B.8, and we have proved that the top parallel functors are isomorphic. But $[\mathbb{C}^{op}, l_i]$ is epi, which entails that the bottom parallel functors are also isomorphic, as desired. $\square$

B.5 Constraining strategies

The goal of this section is to prove Proposition 3.18. Before attacking this, let us reformulate the locality condition on c , using the category $\mathbb{P}_{(A,B), (B,C)}$ of Definition B.1.

Lemma B.9. *An embedding $c : \mathbb{P}^c \hookrightarrow \mathbb{P}$ as in Proposition 3.18, i.e., between game settings sharing the same set of arenas and whose components are discrete fibrations, is a local constraint iff $\text{sing}(c_{(A,B), (B,C)})$, the presheaf induced by the embedding $\mathbb{P}_{(A,B), (B,C)}^c \hookrightarrow \mathbb{P}_{(A,B), (B,C)}$, is a sheaf for the Grothendieck topology induced by the embedding $\mathbb{P}_{A,B} + \mathbb{P}_{B,C} \hookrightarrow \mathbb{P}_{(A,B), (B,C)}$.*

Proof. A direct check. $\square$

Proof of Proposition 3.18. By Proposition B.2, the result reduces to the commutation of

$$\begin{array}{ccccc}
 \mathbb{P}_{A,B}^c + \mathbb{P}_{B,C}^c & \xrightarrow{\Pi} & \mathbb{P}_{(A,B), (B,C)}^c & \xleftarrow{\Delta} & \mathbb{P}_{A,B,C}^c & \xrightarrow{\Sigma} & \mathbb{P}_{A,C}^c \\
 \downarrow \Sigma_c \times \Sigma_c & & \downarrow \Sigma_c & & \downarrow \Sigma_c & & \downarrow \Sigma_c \\
 \mathbb{P}_{A,B} + \mathbb{P}_{B,C} & \xrightarrow{\Pi} & \mathbb{P}_{(A,B), (B,C)} & \xleftarrow{\Delta} & \mathbb{P}_{A,B,C} & \xrightarrow{\Sigma} & \mathbb{P}_{A,C}
 \end{array}$$

“ ”

☐

B.6 Concurrent innocence

In this section, we prove Theorem 4.13, in two lemmas. The first states stability of innocence under composition; the second says that copycat strategies are innocent.

Lemma B.10. *In any innocent game setting, the composite of two innocent strategies is again innocent.*

Proof. Because a strategy X , say on (A, B) , is innocent iff it is isomorphic to $\Pi_{i \in A, B}(\Delta_{i \in A, B}(X))$, it suffices to show that starting from any pair of behaviours $[B_1, B_2] \in \mathbb{V}_{A, B} + \mathbb{V}_{B, C}$, if we extend them to innocent strategies, compose, and eventually apply innocentisation (i.e., $\Pi_{i \in A, C} \circ \Delta_{i \in A, C}$), then the last step is redundant. In other words, we need to show that the perimeter of

$$\begin{array}{ccccccc}
& & \mathbb{P}_{A,B} + \mathbb{P}_{B,C} & \xrightarrow{\Pi} & \mathbb{P}_{(A,B),(B,C)} & \xleftarrow{\Delta} & \mathbb{P}_{A,B,C} & \xrightarrow{\Sigma} & \mathbb{P}_{A,C} \\
& \nearrow \Pi & & & \Delta \uparrow & & \Delta \uparrow & & \uparrow \Delta \\
\mathbb{V}_{A,B} + \mathbb{V}_{B,C} & \xrightarrow{\quad \Pi \quad} & \mathbb{V}_{(A,B),(B,C)} & \xleftarrow{\Delta} & \mathbb{V}_{A,B,C} & \xrightarrow{\Sigma} & \mathbb{V}_{A,C} & & (10) \\
& \searrow \Pi & & & \Pi \downarrow & & \Pi \downarrow & & \downarrow \Pi \\
& & \mathbb{P}_{A,B} + \mathbb{P}_{B,C} & \xrightarrow{\Pi} & \mathbb{P}_{(A,B),(B,C)} & \xleftarrow{\Delta} & \mathbb{P}_{A,B,C} & \xrightarrow{\Sigma} & \mathbb{P}_{A,C}
\end{array}$$

commutes up to isomorphism. We proceed by showing that inn-
 centisation is redundant at every intermediate step, but this re-
 quires us to define intermediate categories of views adequately:
 $\mathbb{V}_{ABC}$ is defined as the pullback $\mathbb{P}_{A,B,C} \times_{\mathbb{P}_{A,C}} \mathbb{V}_{A,C}$, i.e., the full
 subcategory of $\mathbb{P}_{A,B,C}$ spanning plays whose δ_1 -projection is a
 view; for $\mathbb{V}_{(A,B),(B,C)}$, consider first the lax limit $\mathbb{C}$ of $\mathbb{P}_{A,B} \leftarrow$
 $\mathbb{V}_{A,B,C} \rightarrow \mathbb{P}_{B,C}$; $\mathbb{V}_{(A,B),(B,C)}$ is its full subcategory spanning ob-
 jects from $\mathbb{V}_{A,B,C}$, $\mathbb{V}_{A,B}$ and $\mathbb{V}_{B,C}$.

Returning to our claim, the top-left square commutes by Lemma A.7, the top square because the underlying diagram commutes, the top-right square by Lemma A.9, the bottom-left square because the underlying diagram commutes, the bottom-right one by locality. Finally, the bottom square commutes because it is exact by the zig-zag criterion using view-analyticity and observing that it is a pullback of fully faithful functors. $\square$

We finally prove innocence of identities:

Lemma B.11. *Copycat strategies are innocent.*

Proof. We proceed as for preservation of innocence by composition: by showing that copycat is the same as copycat followed by innocentisation. This yields the diagram

$$\begin{array}{ccccc}
& & \mathbb{P}_A & \xrightarrow{\Sigma} & \mathbb{P}_{A,A} \\
& \nearrow \Pi & \uparrow \Delta & & \uparrow \Delta \\
\emptyset & \xrightarrow{\Pi} & \mathbb{V}_A & \xrightarrow{\Sigma} & \mathbb{V}_{A,A} \\
& \searrow \Pi & \downarrow \Pi & & \downarrow \Pi \\
& & \mathbb{P}_A & \xrightarrow{\Sigma} & \mathbb{P}_{A,A}
\end{array}$$

The bottom left triangle commutes because underlying functors do; the top left triangle commutes by Lemma A.7; the bottom square commutes by locality; the top one by Lemma A.9. $\square$

B.7 Prefix-based innocence

This section is devoted to the proof of Proposition 4.16. We proceed as in the previous section: we first need to show that the following diagram commutes up to isomorphism

$$\begin{array}{ccccccc}
\mathbb{P}_{A,B} + \mathbb{P}_{B,C} & \xrightarrow{\Pi} & \mathbb{P}_{(A,B),(B,C)} & \xleftarrow{\Delta} & \mathbb{P}_{A,B,C} & \xrightarrow{\Sigma} & \mathbb{P}_{A,C} \\
\Delta \downarrow & & \Delta \downarrow & & \Delta \downarrow & & \Delta \downarrow \\
\mathbb{P}_{A,B}^+ + \mathbb{P}_{B,C}^+ & \xrightarrow{\Pi} & \mathbb{P}_{(A,B),(B,C)}^+ & \xleftarrow{\Delta} & \mathbb{P}_{A,B,C}^+ & \xrightarrow{\Sigma} & \mathbb{P}_{A,C}^+ \\
\Pi \nearrow & & & & & & \Delta \nearrow \\
\mathbb{V}_{A,B} + \mathbb{V}_{B,C} & & \text{(Theorem 4.13)} & & & & \mathbb{V}_{A,C} \\
\Pi \searrow & & & & & & \Pi \searrow \\
\mathbb{P}_{A,B}^+ + \mathbb{P}_{B,C}^+ & \xrightarrow{\Pi} & \mathbb{P}_{(A,B),(B,C)}^+ & \xleftarrow{\Delta} & \mathbb{P}_{A,B,C}^+ & \xrightarrow{\Sigma} & \mathbb{P}_{A,C}^+ \\
\Delta \uparrow & & \Delta \uparrow & & \Delta \uparrow & & \Delta \uparrow \\
\mathbb{P}_{A,B} + \mathbb{P}_{B,C} & \xrightarrow{\Pi} & \mathbb{P}_{(A,B),(B,C)} & \xleftarrow{\Delta} & \mathbb{P}_{A,B,C} & \xrightarrow{\Sigma} & \mathbb{P}_{A,C}
\end{array}
\quad (11)$$

By Theorem 4.13, this reduces to exactness of (I) and (II). The latter follows by hypothesis (discreteness of $\mathbb{P}_{A,B,C}^+ \rightarrow \mathbb{P}_{A,C}^+$ and (II) being a pullback). The former follows by construction of $\mathbb{P}_{(A,B),(B,C)}^+$, using the zig-zag criterion.

Innocence of copycat presheaves follows similarly, as does coincidence of copycats with the essential image of restriction. Coincidence of innocence with essential image of restriction is obvious.