

THIN GAMES WITH SYMMETRY AND CONCURRENT HYLAND-ONG GAMES

SIMON CASTELLAN, PIERRE CLAIRAMBAULT, AND GLYNN WINSKEL

Univ Lyon, CNRS, ENS de Lyon, UCB Lyon 1, LIP
e-mail address: simon.castellan@ens-lyon.fr

Univ Lyon, CNRS, ENS de Lyon, UCB Lyon 1, LIP
e-mail address: pierre.clairambault@ens-lyon.fr

Computer Laboratory, University of Cambridge
e-mail address: glynn.winskel@cl.cam.ac.uk

ABSTRACT. We build a cartesian closed category, called **Cho**, based on event structures. It allows an interpretation of higher-order stateful concurrent programs that is refined and precise: on the one hand it is conservative with respect to standard Hyland-Ong games when interpreting purely functional programs as innocent strategies, while on the other hand it is much more expressive. The interpretation of programs constructs compositionally a representation of their execution that exhibits causal dependencies and remembers the points of non-deterministic branching.

The construction is in two stages. First, we build a compact closed category **Tcg**. It is a variant of Rideau and Winskel’s category CG, with the difference that games and strategies in **Tcg** are equipped with *symmetry* to express that certain events are essentially the same. This is analogous to the underlying category of AJM games enriching simple games with an equivalence relations on plays. Building on this category, we construct the cartesian closed category **Cho** as having as objects the standard arenas of Hyland-Ong games, with strategies, represented by certain events structures, playing on games with symmetry obtained as expanded forms of these arenas.

To illustrate and give an operational light on these constructions, we interpret (a close variant of) Idealized Parallel Algol in **Cho**.

1. INTRODUCTION

In game semantics, *computation* is represented within a two-player game played between the *program* and its *execution environment* – the program is often considered to be *Player* and the execution environment *Opponent*. The two players make *moves* corresponding to computational events: the program calling an external function is a Player move, and this function returning is an Opponent move. Originally motivated by the very foundational quest of understanding higher-order sequentiality [HO00, AJM00], game semantics developed into a rich subject, with a wide scope spanning logical aspects of computation, through the

Key words and phrases: Game semantics, event structures, concurrency.

Curry-Howard correspondence, to the conception of new decision algorithms for equivalence or verification problems on higher-order programs.

Game semantics often plays one of the two following roles in the literature.

(1) *A syntax-free, compositional operational semantics.* The strategy interpreting a program is a syntax-free object – in essence it is a representation of the *behaviour* of the program, with no information as to *how* this behaviour is written down in the syntax. In particular, it abstracts cleanly from the bureaucratic aspects of the syntax and reduction of the language under examination. It is by nature *compositional*, because the strategy for a term is calculated by induction on its syntax tree, following the methodology of denotational semantics. In particular, the application of one term to another is interpreted as the composition of the corresponding strategies.

A compositional fine-grained description of the execution of higher-order programs is a useful tool – for instance, it provides methodologies to study problems such as termination or complexity in the abstract, in a syntax-free manner [CH10, Cla15]. Such a representation is also key to further program analysis. It can provide an invariant for compilation [Ghi07, Sch14], or a compositional model construction on which to perform model-checking [AGMO04]. Even in the purely functional case, it was recently proposed by Jones *et al* [BJ16] as a convenient closure-free way to compute the partial evaluation of a term.

Although historically the focus on game semantics has often been on the second role mentioned below, a good part of its recent developments have been indeed as a syntax-free, compositional operational semantics. In this direction, it is to be compared with various similar frameworks. *Normal form bisimulations* [LL07] are close relatives, as recently emphasized by Levy and Staton [LS14]. Recent developments of the *geometry of interaction* also pursue similar methods and objectives [HMH14, LFVY15]. Finally, Hirschowitz and collaborators have provided a very general framework in which one can give syntax-free descriptions of different kinds of programs [HP12, EHS15].

(2) *An observational classification of effects.* Beyond the use of game semantics as an operational semantics, the separation of the *observable behaviour* of a term from its syntax allowed researchers to study computational features of programs in terms of the observations that they permitted. One of the most acclaimed achievements of early game semantics is the identification of conditions on strategies (visibility, innocence, well-bracketing) in the context of Hyland-Ong games, that characterize, not syntactically but *observationally*, the behaviour of programs having access to certain effects. Indeed, *innocent well-bracketed* strategies are essentially purely functional programs [HO00]: relaxing innocence to *visibility* captures the use of ground state [AM96] while removing well-bracketing captures control operators [Lai97]. Finally, removing visibility captures terms that have access to higher-order references [AHM98]. This is known as the *semantic cube* or *Abramsky cube*.

Each combination of these conditions corresponds to a certain programming language, for which the strategies have exactly the same observation power as programs. In many cases the resulting model is *fully abstract*, without the need for a quotient. This allowed researchers in game semantics, starting with Ghica and McCusker [GM03], to give decision procedures for observational equivalence of programs in certain programming languages where the fully abstract games model is algorithmically presented [MW08, CHMO15].

Despite this impressive flexibility, each game semantics model comes with its limitations. The notion of *play*, which is at the heart of any game semantics model, specifies the observational power of the execution environment. Whereas the capabilities of Player can be adjusted to a certain extent via conditions on strategies, this cannot be pushed further

than what is wired into the model construction. For instance in Hyland and Ong’s original model, plays are well-bracketed and visible for both players – it follows that we only observe parts of programs visible to an evaluation context with no access to higher-order state and control operators. Clearly that can be solved; for instance by allowing more general plays as in [AHM98]. Then again, the whole model has to be rethought if one wishes to allow Opponent to be concurrent. The same line of thought led Ghica and Tzevelekos to define *system-level game semantics* [GT12], in an effort to take as few assumptions as possible as to the power of the execution environment. We advocate here another option, namely to use a *causal* model.

Causality. In *causal* game semantics, a program is not represented by an enumeration of all its possible interactions with an Opponent of observational strength wired into the model. Instead, it is represented by an abstract structure displaying information about the *causal choices* behind the program’s actions. On the one hand, this means that the model is more intensional and most likely further away from quotient-free full abstraction results. On the other, the representation makes few assumption as to the computational features available to the execution environment. This makes the model more modular, and a finer representation: from the causal game semantics of a program, it is always possible to recover – in an effective manner – the set of plays corresponding to an observation by a certain kind of environment. The causal representation has other advantages. For instance, as long advocated by Melliès, it allows us to get rid of artificialities in the standard play-based composition mechanism for innocent strategies, making more explicit the fact that it is relational (this is key to the full completeness results for fragments of linear logic [AM99, Mel05]). Furthermore, the importance of causal representations for programs has been advocated in the past for various purposes, ranging from error diagnostics [BFHJ03] to the study of reversible aspects of computation [CKV13]. Last but not least, causal representations display the evolution of a concurrent system with partial orders rather than interleavings (it is “*truly concurrent*”). Such representations avoid the *state explosion* problem of interleaving-based ones [God96], leading to potential applications to the verification of concurrent systems.

Contributions. Giving causal representations of the execution of programs is not a new problem. Such models have been developed for various process languages, from CCS [Win86], up to (recently) the full π -calculus [EHS15, CKV15]. There seem to exist few developments on truly concurrent semantics of concurrent languages with shared state, with the notable exception of a Petri net semantics for a simple imperative programming language [HW06].

In this paper, we give a general framework in which one can define such truly concurrent models for higher-order concurrent languages, with various synchronization primitives. This has the form of a cartesian closed category of arenas and concurrent strategies, which are certain event structures. The approach is conservative over the category of Hyland-Ong innocent strategies for PCF (and over the more recent work [TO15] in the non-deterministic case), in the sense that a pure term is interpreted as its forest of *P-views*. In this paper, we develop this category in detail, and illustrate it by spelling out the interpretation of Idealized Parallel Algol (IPA), a higher-order concurrent programming language with shared state and semaphores as synchronization primitives. The methodology is that of game semantics, which provides a well-rooted hope that any of the many languages that one can model in game semantics could be given a truly concurrent representation in this framework as well.

To achieve that, we build on the category \mathbf{CG} of *concurrent games on event structures* introduced by Rideau and Winskel [RW11, CCRW17], which itself belongs to a family of game semantics focusing on positionality rather than sequentiality initiated by Abramsky and Melliès in [AM99] and subsequently pushed by Melliès and others [Mel05, MM07]. Relative to \mathbf{CG} , the main technical difficulty is to handle *replication* in programming languages: the same resource may be accessed many times, but subsequent behaviour cannot depend on the low-level details of how replication is implemented in the model (it is the *uniformity* problem solved with *equivalence relations* in AJM games). The first layer of our framework addresses this difficulty: we first build a compact closed category \mathbf{Tcg} , designed to play a similar role as \mathbf{CG} in that it offers the basic compositional mechanisms on which further semantic constructions rely, while handling uniformity. Then, relying in \mathbf{Tcg} we build the announced cartesian closed category \mathbf{Cho} , designed as a conservative extension of usual Hyland-Ong games. We consider the main contributions of this paper to be these two categories of strategies \mathbf{Tcg} and \mathbf{Cho} ; the interpretation of IPA in \mathbf{Cho} is mostly there for motivations and illustrations purposes. We insist that the intermediate \mathbf{Tcg} is a contribution in itself rather than a means to an end for the construction of \mathbf{Cho} ; and indeed in some further work we have found it convenient to build on \mathbf{Tcg} directly rather than through \mathbf{Cho} .

Other work. Here, we find it appropriate to give some further insight on the broader context of the developments presented in this paper. The bulk of the present paper is the detailed construction of the games model used in the conference paper [CCW15] to build a fully abstract games model of PCF based on parallel evaluation. The result of [CCW15] also necessitates the development of a concurrent notion of innocence [HO00], not covered here. The approach to uniformity adopted here is that given in [CCW15], rather than the earlier approach of [CCW14] (which we will however mention in the course of the paper).

Our aim for the current paper is that it serves as a reference for the construction of \mathbf{Tcg} and \mathbf{Cho} . In that, besides the illustrative interpretation of IPA, it contains no striking application of the framework. Hence, to help motivate the rather lengthy and technically involved model construction, we feel it is helpful to add some further perspective on the use of this model. This framework is a cornerstone for a number of subsequent developments, obtaining achievements that were not within reach with the previously existing tools of game semantics. In [CCW15], pairing \mathbf{Cho} with a notion of concurrent innocence we have given a fully abstract model of PCF based on parallel evaluation. This was extended to non-deterministic PCF in the first author's PhD thesis [Cas17]. The category \mathbf{Tcg} was found to accommodate transparently probabilities, supporting the first notion of probabilistic innocence. This was applied successfully to probabilistic PCF [CCPW18] and the probabilistic λ -calculus [CP18]. Besides those a number of developments are currently under way, in which the constructions detailed here play a crucial role. A common trend in all these developments is that further structure sits on \mathbf{Tcg} and \mathbf{Cho} in a high-level modular way, and do not interact with the details of the construction, making those convenient for further semantic constructions despite their intricacies.

Outline. In Section 2, we give some basic ideas behind the formalization of game semantics on top of event structures, and introduce the key issues that we will have to solve in order to push these ideas to a fully-fledged games model. In particular, we will show that we need to move to a setting of event structures *with symmetry*, in order to handle uniformity

of strategies with respect to replicated resources. In Section 3, we give the first main contribution of this paper: the compact closed category **Tcg**. In Section 4, we rely on it to construct our cartesian closed category **Cho**. Finally in Section 5, we illustrate **Cho** by describing the interpretation of IPA.

2. ARENAS, CONCURRENT STRATEGIES, AND UNIFORMITY

This first section has several purposes: firstly, it aims to introduce the basic ideas behind our concurrent formulation of Hyland-Ong games. Secondly, it recalls from [CCRW17] the required preliminaries on games on event structures. Finally, it introduces the main difficulties encountered in trying to obtain a cartesian closed category based on this, motivating the definitions of Section 3.

2.1. Preliminaries on Idealized Parallel Algol. Before presenting our game semantics, we fix a syntax (inspired by [GM08]) for Idealized Parallel Algol (IPA). It will not be exactly the same language as in [GM08] – notably, it lacks semaphores. We omit them in order to keep the language simple, but they can be interpreted with the same methodology than shared variables. Note that the language is mostly here to fix notations and for providing examples and illustrations. Indeed, the focus on the paper is on the model construction rather than its applications, which will come later in companion papers.

The **types** of IPA are the following.

$$A, B ::= \mathbf{com} \mid \mathbb{B} \mid \mathbb{N} \mid A \rightarrow B \mid \mathbf{ref}$$

The type **com** is a type of *commands*, which returns no useful value (if it returns at all, it returns **skip**), but may perform read/write operations on the memory. The types \mathbb{N} and \mathbb{B} are types for expressions that (if they return) return respectively a natural number or a boolean. Finally, **ref** is the type for integer variables. Note that we consider *active expressions*, *i.e.* the evaluation of a term of type \mathbb{B} or \mathbb{N} can trigger side effects.

Raw **terms** of IPA are described as follows.

$$\begin{aligned} M, N ::= & x \mid M N \mid \lambda x. M \mid Y \mid \\ & \mathbf{tt} \mid \mathbf{ff} \mid \mathbf{if} M N_1 N_2 \mid \\ & n \mid \mathbf{succ} M \mid \mathbf{pred} M \mid \mathbf{iszero} M \mid \\ & \mathbf{skip} \mid M; N \mid M \parallel N \mid \\ & \mathbf{newref} r \mathbf{in} M \mid M := N \mid !M \mid \mathbf{mkvar} \end{aligned}$$

The first three lines describe the syntax of PCF [Plo77]. The fourth line describes *commands* and combinators for them. Finally, the fifth line gives the primitives for manipulating variables. We include the so-called *bad variable* constructor [AM96], but it will only play a very minor role in our development.

These terms are subject to mostly standard typing rules. We give most of them in Figure 1, omitting the standard rules for the λ -calculus, the fixpoint combinator, and constants. By convention, we use \mathbb{X} to range over ground types: **com**, \mathbb{B} , \mathbb{N} . By abuse of notation, we will also use \mathbb{N} and \mathbb{B} respectively for the sets of (total) natural numbers and booleans.

Although some of our typing rules seem restricted to output ground types, the full rules can be derived as syntactic sugar. For instance, a version of **if** that eliminates to **ref** can be

$$\begin{array}{c}
\frac{\Gamma \vdash M : \mathbf{com} \quad \Gamma \vdash N : \mathbf{com}}{\Gamma \vdash M \parallel N : \mathbf{com}} \quad \frac{\Gamma \vdash M : \mathbf{com} \quad \Gamma \vdash N : \mathbb{X}}{\Gamma \vdash M ; N : \mathbb{X}} \quad \frac{\Gamma \vdash M : \mathbf{ref} \quad \Gamma \vdash N : \mathbb{N}}{\Gamma \vdash M := N : \mathbf{com}} \\
\\
\frac{\Gamma, x : \mathbf{ref} \vdash M : \mathbb{X}}{\Gamma \vdash \mathbf{newref} \ x \ \mathbf{in} \ M : \mathbb{X}} \quad \frac{\Gamma \vdash M : \mathbb{B} \quad \Gamma \vdash N_i : \mathbb{X}}{\Gamma \vdash \mathbf{if} \ M \ N_1 \ N_2 : \mathbb{X}} \quad \frac{\Gamma \vdash M : \mathbb{N} \rightarrow \mathbf{com} \quad \Gamma \vdash N : \mathbb{N}}{\Gamma \vdash \mathbf{mkvar} \ M \ N : \mathbf{ref}}
\end{array}$$

FIGURE 1. Typing rules of IPA

obtained as:


$$\frac{\Gamma \vdash M : \mathbb{B} \quad \Gamma \vdash N_i : \mathbf{ref}}{\Gamma \vdash \mathbf{mkvar} (\lambda x. \mathbf{if} \ M (N_1 := x) (N_2 := x)) (\mathbf{if} \ M !N_1 !N_2) : \mathbf{ref}}$$

It is an easy verification that the other rules can be generalized similarly.

The language can be equipped with standard (small-step) operational semantics, see [GM08] for details. We omit it here since it will play no role in the technical development.

2.2. Partial orders and conflicts for strategies. We now start introducing our semantics. In the remainder of this section we introduce gradually the main ideas behind our model, relying as much as possible on examples. Our starting point will be the standard Hyland-Ong innocent semantics for PCF, which we will use to motivate concurrent games on event structures. This section will culminate on the issues of replication and uniformity, which will prompt the developments of Section 3.

2.2.1. Dialogue games. *Hyland-Ong games* formalize the intuition that a program is a strategy having a dialogue with its execution environment. A possible dialogue on the type $\mathbb{B} \rightarrow \mathbb{B} \rightarrow \mathbb{B} \rightarrow \mathbb{B}$ appears in Figure 2. The diagram is read from top to bottom. Each move is

FIGURE 2. A dialogue on $\mathbb{B} \rightarrow \mathbb{B} \rightarrow \mathbb{B} \rightarrow \mathbb{B}$

either by Player or Opponent, and is either a Question or an Answer. Questions correspond to variable calls, whereas Answers indicate a call terminating. The dashed lines between moves (traditionally called *justification pointers*) convey information about thread indexing; in this example they are redundant but become required at higher types – we will see more on them later.

In natural language, this diagram would read: “The environment starts the evaluation of a term of type $\mathbb{B} \rightarrow \mathbb{B} \rightarrow \mathbb{B} \rightarrow \mathbb{B}$ by interrogating its return type $(-, \mathcal{Q})$. The evaluation requires information on the first argument, so the term triggers its evaluation by playing under the corresponding component of the type $(+, \mathcal{Q})$. The evaluation of the argument terminates with value \mathbf{tt} $(-, \mathcal{A})$. Knowing that its first argument is \mathbf{tt} , the term now needs information on its second argument $(+, \mathcal{Q})$. This argument returns \mathbf{tt} $(-, \mathcal{A})$, and now the term computes and returns \mathbf{tt} at toplevel $(+, \mathcal{A})$.” The reader should recognize here a description of an execution of **if**.

In Hyland-Ong games, **sequential innocent strategies** consist of sets of dialogues as above, where Opponent moves are *justified* by the preceding one – such dialogues are known as *P-views*. A strategy for a term of PCF contains several such dialogues, specifying the term entirely. The full strategy for **if** contains in total four maximal P-views, displayed in Figure 3. Such non-empty sets of P-views (satisfying further conditions: determinism and




FIGURE 3. Maximal *P*-views for **if**

well-bracketing) are usually called *well-bracketed innocent strategies*. Because of their correspondence with certain normal forms (*PCF Böhm trees* [Cur98]), they are the cornerstone of Hyland-Ong games and of the full abstraction results they allowed.

2.2.2. *Partially ordered dialogues.* In Hyland-Ong games, every P-view is a total order, meaning that the whole sequential innocent strategy is a tree. In our framework, we question this premise. For instance, informally, the intuitive behaviour of the parallel composition operation $\parallel: \mathbf{com} \rightarrow \mathbf{com} \rightarrow \mathbf{com}$ of IPA is most elegantly represented as in Figure 4.

The diagram of Figure 4 is analogous to the previous ones, but is now partially ordered rather than totally ordered. The relation \rightarrow denotes immediate causality; it was unnecessary before, as it coincided with chronological contiguity. The justification pointers remain – we will see more on their precise nature later. Note that in the standard game semantics for IPA [GM08], this partial order would be only implicit; and given by all the possible linear orderings of the partial order above. Here instead, the partial order will be the first-class object of interest. Strategies, in particular, will be partially ordered.




FIGURE 4. A partially ordered dialogue

2.2.3. *Non-determinism.* Concurrent languages are, in general, non-deterministic. Note that we do not mean non-deterministic in the sense that, as above, the execution admits multiple distinct linear orderings. For us, non-determinism means that execution takes irreconciliable routes, even up to permutation of independent events. We illustrate that by the following two examples.

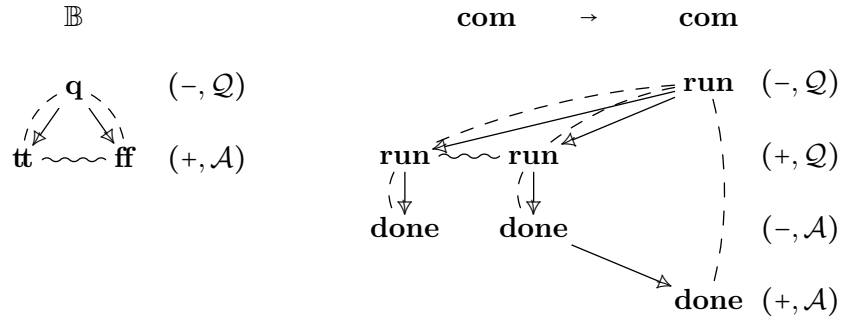


FIGURE 5. Non-deterministic dialogues

In the above two diagrams, the wiggly line \sim indicates *immediate conflict*. Two moves/events related by immediate conflict are incompatible, and can never occur together in an execution. Accordingly, the first example is a representation of our strategy for the non-deterministic boolean, which answers either true or false. The second example illustrates another key aspect of our model: we remember explicitly the point of non-deterministic choice. Here, Player silently flips a coin. If the result is heads, they evaluate the argument, then terminate. However, if the result is tails, they evaluate the argument, then diverge. Typical play-based game semantics would forget the halting branch, which is contained in the other. Instead, our model represents the two branches explicitly.

We now show how to make such diagrams formal.

2.3. Prestrategies on arenas.

2.3.1. *Event structures.* Such a combination of causality and non-determinism is elegantly expressed via Winskel's *event structures* [Win86].

Definition 2.1. An **event structure** (*es* for short) is a tuple $(E, \leq_E, \text{Con}_E)$ where E is a set of *events*, \leq_E is a partial order indicating *causality* and Con_E is a nonempty set of finite subsets of E , satisfying:

$$\begin{aligned} \forall e \in E, [e]_E &= \{e' \in E \mid e' \leq_E e\} \text{ is finite} \\ \forall e \in E, \{e\} &\in \text{Con}_E \\ \forall X \in \text{Con}_E, \forall Y \subseteq X, Y &\in \text{Con}_E \\ \forall X \in \text{Con}_E, \forall e \in E, \forall e' \in X, e &\leq_E e' \implies X \cup \{e\} \in \text{Con}_E \end{aligned}$$

The set Con_E of *consistent subsets* specifies which events can occur together in an execution of the system. The *states* of an event structure E , called the (finite) **configurations**, are those finite sets $x \subseteq E$ that are both consistent and **down-closed** (*i.e.* for all $e \in x$, for all $e' \leq e$, then $e' \in x$) – the set of configurations on E is written $\mathcal{C}(E)$, and is partially ordered by inclusion. Configurations with a maximal element are called **prime configurations**, they are those of the form $[e]$ for $e \in E$ (note that we drop the E in $[e]_E$ whenever, as above, this is clear from the context) – more generally, we will write $[X]_E$ for the down-closure of a set of events X in E . We will also use the notation $[e] = [e] \setminus \{e\}$. We write $x \xrightarrow{e} c$ to mean that $e \notin x$ and $x \cup \{e\} \in \mathcal{C}(E)$. Finally, when drawing event structures as above, we do not represent the full partial order \leq but the **immediate causality** generating it, defined as $e \rightarrow e'$ whenever $e < e'$ and for any $e \leq e'' \leq e'$, either $e = e''$ or $e'' = e'$. We will often omit the subscripts in \leq or \rightarrow when they are clear from the context.

In this paper, most of the event structures we consider (such as those in the previous subsection) have a simpler consistency structure.

Definition 2.2. An **event structure with binary conflict** is a triple $(E, \leq_E, \#_E)$, where \leq_E is a partial order and $\#_E$ is an irreflexive symmetric binary relation on E , such that:

$$\begin{aligned} \forall e \in E, [e]_E &\text{ is finite} \\ \forall e_1 \#_E e_2, \forall e_2 \leq_E e'_2, e_1 &\#_E e'_2 \end{aligned}$$

It is easy to check that an event structure with binary conflict is an event structure, with $\text{Con}_E = \{X \in \mathcal{P}_f(E) \mid \forall e_1, e_2 \in X, \neg(e_1 \#_E e_2)\}$. On the other hand, not every event structure can be described via a binary conflict (take *e.g.* three events with any subset of cardinal less than two being consistent). The strategies in the cartesian closed category we aim to build will only have binary conflict, and accordingly in Section 4 we will restrict to event structures with binary conflict. In the meantime, some aspects of the theoretical development are smoother when carried out with arbitrary consistency.

In an event structure with binary conflict, we can trace back conflicts to their original cause. For $e_1 \#_E e_2$ we say that the conflict is **minimal**, written $e \sim_E e'$, iff for all $e'' \leq e$ we have $\neg(e'' \#_E e')$ and for all $e'' \leq e'$ we have $\neg(e \#_E e'')$. As above, we will often drop the subscripts in $\#$ or \sim when they are clear from the context. Following this notation, all the diagrams of the previous subsection can be regarded as representations of event structures with binary conflict (ignoring the dashed lines).

2.3.2. *Games and arenas.* In game semantics, dialogues as in Subsection 2.2 obey the rules of a *game* inherited from the type. In order to define it, let us first recall the following notion from [CCRW17].

Definition 2.3. An **event structure with polarities (esp)** is an event structure A along with a **polarity function**

$$\text{pol}_A : A \rightarrow \{-, +\}$$

associating to any event a polarity, that is either $-$ for Opponent or $+$ for Player.

By a **game**, we simply mean an event structure with polarities.

Those games form the objects of the category CG of concurrent games of [CCRW17]. As one key aim of the paper is to reconstruct a version of Hyland-Ong games, certain games called *arenas* will play a particular role in our development (being the objects of our cartesian closed category **Cho**).


Definition 2.4. An **arena** is a **conflict-free** (all finite sets consistent) esp/game satisfying:

- *Forest.* if $a_1, a_2 \leq a \in A$, then either $a_1 \leq a_2$ or $a_2 \leq a_1$.
- *Alternation.* if $a_1 \rightarrow a_2$, then $\text{pol}(a_1) \neq \text{pol}(a_2)$.

An arena A is **negative** if all its minimal events have negative polarity.

Arenas are close representations of types. Although formulated a bit differently, our arenas are the same as in [HO00] (with the exception of the absence of the Question/Answer labeling, which we do not require: as we do not aim for a full abstraction result, no notion of well-bracketing is needed).

Example 2.5. Leaving for later the general interpretation of types, we have:



Throughout this paper, we will often omit the semantic brackets on types when this causes no confusion and simply refer to these arenas as \mathbb{B} , **com**, *etc.*

By convention, we represent immediate causality in arenas by dashed lines $- - -$ rather than \rightarrow . Events are ordered from top to bottom, and are annotated with their polarity. We observe in the third example – and it will be a general fact once we give the formal definitions – that each move in the arena $\llbracket \mathbf{com} \rightarrow \mathbf{com} \rrbracket$ *comes from* a well-defined occurrence of a base type **com** in $\mathbf{com} \rightarrow \mathbf{com}$: \mathbf{run}^- and \mathbf{done}^+ come from the output **com**, and \mathbf{run}^+ and \mathbf{done}^- come from the input **com**. As usual in game semantics, this is used in the representation of dialogues (as in Subsection 2.2): whenever possible, moves are placed under the corresponding base type occurrence.

2.3.3. *Prestrategies.* The dialogues of Subsection 2.2, and our notion of strategies (called *prestrategies* for now – more conditions are to come), will be event structures *labeled by a game*. In other words, a prestrategy will be an event structure S along with a labeling function $\sigma : S \rightarrow A$ associating to each event in S an image in the game. These labeling functions need to satisfy conditions corresponding to the notion of *map of event structures*.

Definition 2.6. Let E, F be event structures. A **morphism (map) of event structures** $f : E \rightarrow F$ is a function, satisfying:

- *Preservation of configurations.* For all $x \in \mathcal{C}(E)$, $fx \in \mathcal{C}(F)$,
- *Local injectivity.* For all $e, e' \in x \in \mathcal{C}(E)$, if $fe = fe'$ then $e = e'$.


Event structures and maps between them form a category \mathcal{E} .

A **prestrategy** on a game A is a map of event structures $\sigma : S \rightarrow A$.

So a prestrategy $\sigma : S \rightarrow A$ must only reach valid states of A , and behaves *linearly*: in a configuration, each event of the game appears at most once. We note in passing that for non-linear languages, this linearity assumption will be circumvented by creating duplicates of events – more on that later.

If $\sigma : S \rightarrow A$ is a prestrategy, then S automatically inherits from A a polarity function that we write $\text{pol}_S : S \rightarrow \{-, +\}$, leaving the dependency on σ implicit. Of course, it is equivalent to require S to be explicitly equipped with polarities, in a way preserved by σ – we then say that σ is a **map of esps**.

2.3.4. *Representations of prestrategies.* We will only draw prestrategies with binary conflict. When drawing such a $\sigma : S \rightarrow A$ as in Subsection 2.2, we only draw S (more precisely, with immediate causality \rightarrow and immediate conflict \rightsquigarrow), where each event is presented as its image through A , and placed under the corresponding ground type occurrence in the type. We use the dashed lines $- - -$ to represent the relation on S induced by immediate causality on A . For instance, the second diagram of Figure 5 is a representation of the map of event structures below.



As the reader can see, this explicit map notation is a bit cumbersome. Its representation as in Figure 5 conveys the relevant information – the only thing lost is the “name” (s_1, \dots, s_6) of moves in S . More formally, it should be clear to the reader that such a representation displays a finite prestrategy $\sigma : S \rightarrow A$ adequately up to **isomorphism** of prestrategies:

Definition 2.7. Let $\sigma : S \rightarrow A$ and $\tau : T \rightarrow A$ be prestrategies. A **morphism** from σ to τ is a map of event structures $f : S \rightarrow T$ such that $\tau \circ f = \sigma$.

Accordingly, an **isomorphism** between σ and τ is given by (f, g) , where $f : S \rightarrow T$ and $g : T \rightarrow S$ are maps between σ and τ such that $g \circ f = \text{id}_S$ and $f \circ g = \text{id}_T$. We write $\sigma \cong \tau$ to mean that σ and τ are isomorphic – in that case we might sometimes say that σ and τ are **strongly isomorphic** to emphasize the distinction with *weak isomorphisms*, to be defined in Definition 2.28.

2.4. Compositional structure. In order to obtain such representations of programs compositionally, the standard methodology of denotational semantics suggests to organize them as a category. Rideau and Winskel [RW11] give the basic ingredients for the construction of a (bi)category of games on event structures. We give here the main ideas and definitions, but refer the reader to [CCRW17] for a more in-depth construction with proofs.

We start with the following simple definition.

Definition 2.8. The **simple parallel composition** $E_1 \parallel E_2$ of two event structures E_1 and E_2 has:

- *Events.* The disjoint union $\{1\} \times E_1 \cup \{2\} \times E_2$,
- *Causality.* We have $(i, e) \leq_{E_1 \parallel E_2} (j, e')$ iff $i = j$ and $e \leq_{E_i} e'$,
- *Consistency.* For $X = \{1\} \times X_1 \uplus \{2\} \times X_2$ (often written simply $X_1 \parallel X_2$) a finite subset of $E_1 \parallel E_2$, we have $X \in \text{Con}_{E_1 \parallel E_2}$ iff $X_1 \in \text{Con}_{E_1}$ and $X_2 \in \text{Con}_{E_2}$.

The simple parallel composition of event structures with binary conflict still has binary conflict. Simple parallel composition can also be applied to maps $f : E \rightarrow E'$ and $g : F \rightarrow F'$ to form $f \parallel g : E \parallel F \rightarrow E' \parallel F'$ in the obvious way; as usual we also write $f \parallel F$ for $f \parallel \text{id}_F$. If A and B have polarities, so is $A \parallel B$, with $\text{pol}_{A \parallel B}((1, a)) = \text{pol}_A(a)$ and $\text{pol}_{A \parallel B}((2, b)) = \text{pol}_B(b)$.

In other words, the two event structures are put side by side, without any interaction. If A is a game, then there is also its **dual** A^\perp , defined as having the same events, causality, consistency as A , but reversed polarity: $\text{pol}_{A^\perp}(a) = -\text{pol}_A(a)$. Both operations $- \parallel -$ and $(-)^{\perp}$ are defined on all esps/games, but preserve arenas.

2.4.1. Morphisms. Given games A and B , a **prestrategy from A to B** is a prestrategy:

$$\sigma : S \rightarrow A^\perp \parallel B$$

We will sometimes write $\sigma : A \rightarrow B$ to keep the S anonymous. The basic example of a prestrategy from A to A is the *copycat strategy*.

Definition 2.9. Let A be a game. We define an event structure \mathbb{C}_A as having:

- *Events.* Those of $A^\perp \parallel A$,
- *Causality.* The transitive closure of the relation:


$$\begin{aligned} & \{((1, a), (1, a')) \mid a \leq_{A^\perp} a'\} \cup \{((2, a), (2, a')) \mid a \leq_A a'\} \cup \\ & \{((1, a), (2, a)) \mid \text{pol}_A(a) = +\} \cup \{((2, a), (1, a)) \mid \text{pol}_A(a) = -\} \end{aligned}$$
- *Consistency.* For X a finite subset of \mathbb{C}_A , we have $X \in \text{Con}_{\mathbb{C}_A}$ iff $[X]_{\mathbb{C}_A} \in \text{Con}_{A^\perp \parallel A}$.

In particular, if A is an arena, then \mathbb{C}_A is conflict-free.

The **copycat prestrategy** is the identity function, which is a map of es:

$$\mathbb{c}_A : \mathbb{C}_A \rightarrow A^\perp \parallel A$$

Example 2.10. The copycat prestrategy from $\llbracket \mathbf{com} \rightarrow \mathbf{com} \rrbracket$ to itself is:



Note that the partial order above is a tree, whose branches are exactly the P-views of the usual corresponding copycat strategy in Hyland-Ong games.

2.4.2. Interaction. As usual in game semantics, composition is obtained by a two-step process: parallel interaction, plus hiding. The main difficulty in defining the composition of prestrategies is parallel interaction – we first explain how it is done on a closed interaction between $\sigma : S \rightarrow A$ and $\tau : T \rightarrow A^\perp$. The *interaction* of σ and τ , written $\sigma \wedge \tau : S \wedge T \rightarrow A$, will be a labeled event structure (by which we mean an event structure $S \wedge T$ with a map $\sigma \wedge \tau$, the *labelling function*) describing the behaviours accepted by both σ and τ .

Its construction is done in several stages. Firstly, its *states* should correspond to certain pairings between matching states of σ and states of τ , *i.e.* pairs $(x, y) \in \mathcal{C}(S) \times \mathcal{C}(T)$ such that $\sigma x = \tau y$. Note that in such a case, the local injectivity assumption on σ and τ induces a bijection $\varphi_{x,y}$ between x and y – in fact matching pairs (x, y) are in one-to-one correspondence with bijections $\varphi : x \simeq y$ between configurations of S and T such that for all $s \in x$, $\tau(\varphi s) = \sigma s$, indicating which events *synchronise* with each other. However, not all such bijections represent valid states of the interaction, as σ and τ might not agree on the *order* in which to play events in x, y . This is addressed by requiring bijections to be *secured*, as below.

Definition 2.11. Let $\sigma : S \rightarrow A$ and $\tau : T \rightarrow A^\perp$ be prestrategies. A **secured bijection** between $x \in \mathcal{C}(S)$ and $y \in \mathcal{C}(T)$ is a bijection


$$\varphi : x \simeq y$$

such that for all $s \in x$ we have $\tau(\varphi s) = \sigma s$, and which is **secured**, in the sense that the reflexive transitive closure of

$$(s, t) \triangleleft (s', t') \Leftrightarrow s <_S s' \vee t <_T t'$$

is a partial order written \leq_φ on (the graph of) φ , making (φ, \leq_φ) a poset. We write $\mathcal{B}_{\sigma, \tau}^{\text{sec}}$ the set of secured bijections between σ and τ .

Example 2.12. In the diagram below are represented two prestrategies σ on $\mathbf{com} \parallel \mathbf{com}$, and τ on $(\mathbf{com} \parallel \mathbf{com})^\perp$.



The dotted line is the only pair in the unique maximal secured bijection in $\mathcal{B}_{\sigma,\tau}^{\text{sec}}$. The maximum configurations of σ and τ are matching, but not in a secured way.

This gives a notion of state of the interaction, but we expected to build a labeled event structure. Hence we wish to present $\mathcal{B}_{\sigma,\tau}^{\text{sec}}$ (up to isomorphism) as the set of configurations of an event structure $S \wedge T$. This is done via the *prime construction*. Say that a secured bijection (φ, \leq_φ) is a **prime** when it has exactly one maximal element (s_φ, t_φ) . In other words, a prime secured bijection is the data of one synchronisation (s_φ, t_φ) , plus a causally valid history for it. We now form:

Definition 2.13. The event structure $S \wedge T$ is obtained as follows:

- *Events.* Prime secured bijections $\varphi \in \mathcal{B}_{\sigma,\tau}^{\text{sec}}$.
- *Causal order.* Inclusion of secured bijections.
- *Consistency.* For X a finite subset of $\mathcal{B}_{\sigma,\tau}^{\text{sec}}$ of prime secured bijections, we have $X \in \text{Con}_{S \wedge T}$ iff $\cup X \in \mathcal{B}_{\sigma,\tau}^{\text{sec}}$.

There is a map of es $\sigma \wedge \tau : S \wedge T \rightarrow A$, given by $(\sigma \wedge \tau) \varphi = \sigma s_\varphi = \tau t_\varphi$.

In passing, we note that if $\sigma : S \rightarrow A$ and $\tau : T \rightarrow A^\perp$ have binary conflict (meaning S and T have binary conflict), then so does $\sigma \wedge \tau : S \wedge T \rightarrow A$, with $\neg(\varphi_1 \#_{S \wedge T} \varphi_2)$ iff $\varphi_1 \cup \varphi_2 \in \mathcal{B}_{\sigma,\tau}^{\text{sec}}$ – this easily boils down to the lemma below.

Lemma 2.14. Assume σ, τ have binary conflict, and let X be a finite subset of $\mathcal{B}_{\sigma,\tau}^{\text{sec}}$. Then the two following statements are equivalent.

- (1) We have $\cup X \in \mathcal{B}_{\sigma,\tau}^{\text{sec}}$,
- (2) For all $\varphi, \psi \in X$, $\neg(\varphi \# \psi)$.

Proof. (1) \Rightarrow (2). Obvious, since $\varphi \cup \psi$ is a down-closed subset of $\cup X$.

(2) \Rightarrow (1). First we show that $\cup X$ is a bijection from a configuration of S to a configuration of T . Indeed if $(s, t_1), (s, t_2) \in \cup X$ then $(s, t_1) \in \varphi$ and $(s, t_2) \in \psi$ for some $\varphi, \psi \in X$. But then $\varphi \cup \psi \in \mathcal{B}_{\sigma,\tau}^{\text{sec}}$ by hypothesis, so in particular it is a bijection, so $t_1 = t_2$. Hence $\cup X$ is a bijection. Furthermore its projections are configurations. Indeed, since X is pairwise compatible, $\pi_1 X$ is a pairwise compatible set of configurations of S , hence $\cup(\pi_1 X) \in \mathcal{C}(S)$ since S has binary conflict – and likewise for $\cup(\pi_2 X)$. It remains to prove that $\cup X$ is secured. But all $\varphi \in X$ are down-closed partial orders induced as subsets of a common preorder of synchronized events. Taking their union cannot introduce cycles, as these would already occur in the separate components by down-closure. Hence $\cup X \in \mathcal{B}_{\sigma,\tau}^{\text{sec}}$. \square

It will also be useful later to have a concrete understanding of how minimal conflict arises in an interaction; hence we prove the following lemma.

Lemma 2.15. *Let $\sigma : S \rightarrow A$ and $\tau : T \rightarrow A^\perp$ be prestrategies, and $\varphi \in \mathcal{B}_{\sigma,\tau}^{\text{sec}}$. Then if $\varphi : x \simeq y$ extends in $\mathcal{B}_{\sigma,\tau}^{\text{sec}}$ with (s_1, t_1) and (s_2, t_2) but $\varphi \cup \{(s_1, t_1), (s_2, t_2)\} \notin \mathcal{B}_{\sigma,\tau}^{\text{sec}}$. Then, either $x \cup \{s_1, s_2\} \notin \mathcal{C}(S)$ or $y \cup \{t_1, t_2\} \notin \mathcal{C}(T)$.*

In particular, if S, T have binary conflict and φ, ψ are events in $S \wedge T$ such that $\varphi \rightsquigarrow \psi$, then $s_\varphi \rightsquigarrow s_\psi$ or $t_\varphi \rightsquigarrow t_\psi$.

Proof. If $x \cup \{s_1, s_2\} \in \mathcal{C}(S)$, $y \cup \{t_1, t_2\} \in \mathcal{C}(T)$, and s_1, s_2 and t_1, t_2 are distinct, then clearly $\varphi \cup \{(s_1, t_1), (s_2, t_2)\} \in \mathcal{B}_{\sigma,\tau}^{\text{sec}}$. However if e.g. $s_1 = s_2$, then $\tau t_1 = \tau t_2$, hence $y \cup \{t_1, t_2\} \notin \mathcal{C}(T)$ by local injectivity – and similarly if $t_1 = t_2$.

The second part of the statement follows easily. \square

In fact, what we have described above is the *pullback* construction in \mathcal{E} . There are maps of event structures:

$$\begin{array}{ccc} \Pi_1 : S \wedge T & \rightarrow & S \\ \varphi & \mapsto & s_\varphi \end{array} \quad \begin{array}{ccc} \Pi_2 : S \wedge T & \rightarrow & T \\ \varphi & \mapsto & t_\varphi \end{array}$$

making the following square commute, and a pullback (from Lemma 2.11 of [CCRW17]):

$$\begin{array}{ccc} & S \wedge T & \\ \Pi_1 \swarrow & \downarrow \vee & \searrow \Pi_2 \\ S & & T \\ \sigma \searrow & & \swarrow \tau \\ & A & \end{array}$$

We motivated the pullback by asking for an es whose configurations are secured bijections. And indeed, those are in a very close correspondence.

Proposition 2.16. *For any $x \in \mathcal{C}(S \wedge T)$, we have $\varphi_x = \cup x : \Pi_1 x \simeq \Pi_2 x$. Moreover, the assignment:*

$$\begin{array}{ccc} \mathcal{C}(S \wedge T) & \rightarrow & \mathcal{B}_{\sigma,\tau}^{\text{sec}} \\ x & \mapsto & \varphi_x \end{array}$$

is an order-isomorphism (with both sets ordered by inclusion). Finally, there is a family of order-isomorphisms:

$$\begin{array}{ccc} \nu_x : x & \cong & \varphi_x \\ \psi & \mapsto & (s_\psi, t_\psi) \end{array}$$

that is natural in x , i.e. for $x \subseteq y$ and $\varphi \in x$ we have $\nu_x \varphi = \nu_y \varphi$.

Proof. Direct extension of Lemma 2.9 in [CCRW17]. \square

This allows us, when reasoning on configurations of a pullback, to manipulate directly secured bijections rather than compatible sets of prime secured bijections. Likewise, when reasoning on events of the pullback in an ambient configuration, we can directly apply ν and reason on synchronised pairs. In the proofs, we will often use this proposition implicitly and transfer silently between the different representations.


Finally, we are in position to define the parallel interaction of two prestrategies $\sigma : S \rightarrow A^\perp \parallel B$ and $\tau : T \rightarrow B^\perp \parallel C$. We simply form the pullback:

$$\begin{array}{ccc}
 & (\sigma \parallel C) \wedge (A \parallel \tau) & \\
 & \Pi_1 \swarrow \quad \searrow \Pi_2 & \\
 S \parallel C & & A \parallel T \\
 \sigma \parallel C \searrow & & \swarrow A \parallel \tau \\
 & A \parallel B \parallel C &
 \end{array}$$

We write $T \otimes S = (\sigma \parallel C) \wedge (A \parallel \tau)$ for the interaction, and $\tau \otimes \sigma : T \otimes S \rightarrow A \parallel B \parallel C$ for either side of the pullback square. Hence we get the interaction of σ and τ , $\tau \otimes \sigma$, as a labeled event structure.

Example 2.17. Consider the following prestrategies $\sigma : S \rightarrow \llbracket \mathbf{com} \rightarrow \mathbf{com} \rrbracket$ and $\tau : T \rightarrow \llbracket \mathbf{com} \rightarrow \mathbf{com} \rrbracket^\perp \parallel \llbracket \mathbf{com} \rrbracket$ (note that to match the definition of interaction above, we can consider $\sigma : S \rightarrow \mathbf{1}^\perp \parallel \llbracket \mathbf{com} \rightarrow \mathbf{com} \rrbracket$, where $\mathbf{1}$ is the empty arena).


$$\sigma : \llbracket \mathbf{com} \rightarrow \mathbf{com} \rrbracket \quad \tau : \llbracket \mathbf{com} \rightarrow \mathbf{com} \rrbracket^\perp \parallel \llbracket \mathbf{com} \rrbracket$$



We display below a representation of the interaction:

$$\tau \otimes \sigma : T \otimes S \rightarrow (\mathbf{com} \rightarrow \mathbf{com}) \parallel \mathbf{com}$$

We only display polarities for the moves in the right hand side \mathbf{com} . Indeed events on the left hand side (synchronised) part of the interaction have no well-defined polarity, as the two strategies disagree on them.



We leave it to the reader to check that each event in this diagram corresponds uniquely to a configuration in $S \parallel \mathbf{com}$ and a matching configuration in T such that the induced bijection is secured.

2.4.3. Hiding. Once we have performed the interaction, it is fairly simple to obtain the composition by ignoring the *synchronised events*, *i.e.* those that map to B . This is an instance of the following *projection* operation.

Definition 2.18. Let E be an event structure, and $V \subseteq E$ a set of events of E . The **projection of E to V** , written $E \downarrow V$, has components:

- *Events.* V .
- *Causality.* The order \leq_E restricted to V .
- *Consistency.* The sets $X \in \text{Con}_E$ such that $X \subseteq V$.


This gives an event structure – it is clear that a hiding of an event structure with binary conflict still has binary conflict. Note as well the *unique witness* property reminiscent of that used in studying the composition of deterministic strategies in standard game semantics: for any $x \in \mathcal{C}(E \downarrow V)$, there exists a *unique* $y \in \mathcal{C}(E)$ such that $y \cap V = x$, and whose maximal events are those of x ; obtained as $y = [x]_E = \{e \in E \mid \exists e' \in x, e \leq_E e'\} \in \mathcal{C}(E)$.

Finally, we define composition. From $\sigma : S \rightarrow A^\perp \parallel B$ and $\tau : T \rightarrow B^\perp \parallel C$, first compute the interaction $\tau \otimes \sigma : T \otimes S \rightarrow A \parallel B \parallel C$. Then, set $V \subseteq T \otimes S$ to comprise all $\varphi \in T \otimes S$ such that $(\tau \otimes \sigma) \varphi \notin B$. Writing $T \odot S = T \otimes S \downarrow V$, the composition of σ and τ is:

$$\begin{aligned} \tau \odot \sigma & : T \odot S \rightarrow A^\perp \parallel C \\ & \varphi \mapsto (\tau \otimes \sigma) \varphi \end{aligned}$$

From the fact that interaction and hiding preserve binary conflict, it follows that for $\sigma : S \rightarrow A^\perp \parallel B$ and $\tau : T \rightarrow B^\perp \parallel C$, if S, T, A, C have binary conflict, then so does $T \odot S$.

Example 2.19. Consider the interaction of Example 2.17. After hiding, the resulting composition is:



Note that the conflict between the two maximal events, although it was inherited in the interaction, becomes minimal after projection as its original cause has been hidden away.

Composition is associative up to isomorphism [CCRW17]. However, copycat is not neutral for composition with respect to prestrategies – it is only the case for *strategies* (see [CCRW17] for details):

Definition 2.20. A prestrategy $\sigma : S \rightarrow A$ on a game A is a **strategy** if it is:

- *Receptive.* For all $x \in \mathcal{C}(S)$, if $\sigma x \xrightarrow{a^-} c$, then there exists a unique $s \in S$ such that $x \xrightarrow{s} c$ and $\sigma s = a$.
- *Courteous.* If $s_1 \rightarrow_S s_2$ and $\text{pol}(s_1) = +$ or $\text{pol}(s_2) = -$, then $\sigma s_1 \rightarrow_A \sigma s_2$.

Putting everything together, we get [CCRW17]:

Theorem 2.21. *There is a compact closed category CG of games and strategies up to isomorphism.*

2.5. Interpreting programs and replication. The category CG is a general framework for composing concurrent strategies. We can rely on it to build a model of an affine variant of IPA: that involves restricting to negative arenas, and interpreting function space in IPA via the usual arrow arena construction. We refrain from giving the details here, since we will give them in the non-affine case later on. However, before going on to handling replication, we will give examples in the affine case and try to convey further intuition as to what the


model computes. Then we will present the expanded arenas used to handle replication, and we will introduce the issue of uniformity.

2.5.1. *Concurrent strategies and view functions.* As the reader familiar with Hyland-Ong games may have noticed, our examples before showed how to represent as concurrent strategies, *view functions* rather than expanded strategies – or, in Curien’s terminology [Cur98], *meager* rather than *fat* innocent strategies. And indeed, in our framework it is the case that a pure program will be interpreted directly as its view function, never constructing the full set of plays. For illustration, the interpretation of the affine pure program:

$$\lambda f^{\mathbb{B} \rightarrow \mathbf{com}}. f \mathbf{tt} : (\mathbb{B} \rightarrow \mathbf{com}) \rightarrow \mathbf{com}$$

will be the strategy:

$$(\mathbb{B} \rightarrow \mathbf{com}) \rightarrow \mathbf{com}$$



which the reader can match against the tree of P -views for the corresponding Hyland-Ong innocent strategy. The composition of such strategies is computed directly using pullbacks in \mathcal{E} , never constructing the expanded plays. In other words, we never work with full Hyland-Ong strategies, but always with their causal representations: the view functions.


But the usual strategies for stateful programs [AM96] are not generalizations of *meager* innocent strategies, but of *fat* ones: the behaviour of programs must be observed not only on P -views but on general plays. Hence, the reader may wonder if evading them causes us to lose that ability. Fortunately it is not the case, and strategies for stateful programs can be represented causally just as innocent strategies. For instance, consider the following example.

Example 2.22. Consider the following term of IA.

$$\mathbf{newref} \ b \ \mathbf{in} \ \lambda f^{\mathbf{com} \rightarrow \mathbf{com}}. f (b := \mathbf{tt}); !b : (\mathbf{com} \rightarrow \mathbf{com}) \rightarrow \mathbb{B}$$

Following (the affine variant of) the interpretation of Section 5, it yields the strategy:

$$(\mathbf{com} \rightarrow \mathbf{com}) \rightarrow \mathbb{B}$$



The \mathbf{done}^+ to the left is duplicated, witnessing the two outcomes of the race in the memory that happens if the argument does not respect the evaluation stack, and concurrently returns \mathbf{done}^- and asks for its argument. The reader familiar with the game semantics for

Idealized Algor [AM96] can check that taking the set of (well-bracketed) alternating linear orderings of configurations of this event structure yields the expected set of plays.

2.5.2. *Replication.* But so far, we have only seen *affine* programs and strategies, *i.e.* that call each resource at most once. As it stands, the local injectivity condition in Definition 2.6 forbids us from having two compatible events corresponding to the same move in the game. It is natural to consider dropping it, but then we lose access to the nice structural properties of \mathcal{E} (such as pullbacks). It is unclear to us how one would go about defining composition of strategies in such a setting, let alone proving that it forms a category; in particular if one insists on remembering the point of non-deterministic branching.

Instead, the solution behind our category **Cho** takes inspiration from AJM games [AJM00] and from the reconstruction of HO games in [HHM07]: we explicitly duplicate moves in arenas. Rather than playing directly on an arena A , our strategies play in $!A$, a variant of A with events duplicated countably many times, in depth. More formally:

Definition 2.23. Let A be an arena, and $a \in A$. An **indexing function** for a is a function:

$$\alpha : [a] \rightarrow \omega$$

which associates, to a and its dependencies, a **copy index**. From $\alpha : [a] \rightarrow \omega$, we write $\text{lbl } \alpha = a$ for its **label**, and $\text{ind } \alpha = \alpha(\text{lbl } \alpha)$ for the copy index of a .

Indexing functions will be the *events* of $!A$. Its full structure will be:

Definition 2.24. From an arena A , we build a new arena $!A$, comprising:

- *Events.* indexing functions $\alpha : [a] \rightarrow \omega$,
- *Causal order.* for $\alpha : [a] \rightarrow \omega$ and $\beta : [b] \rightarrow \omega$, we have $\alpha \leq_{!A} \beta$ iff $a \leq_A b$ and for all $a' \leq_A a$, $\alpha(a') = \beta(a')$.
- *Polarity.* For $\alpha \in !A$, $\text{pol}_{!A}(\alpha) = \text{pol}_A(\text{lbl } \alpha)$.


Moves in $!A$ have a rather complex structure. However, note that just as A – which was required to be an arena rather than a general game – $!A$ is a forest. For each $\alpha : [a] \rightarrow \omega$, either a is minimal and then so is α , or there is a unique $a' \rightarrow_A a$ – in which case the restriction of α to $[a']$ gives a unique α' such that $\alpha' \rightarrow_{!A} \alpha$. In other words, α is entirely determined by the data of $\text{lbl } \alpha = a$, $\text{ind } \alpha = \alpha(a)$, and its immediate predecessor α' , called its **justifier** $\text{just}(\alpha)$. Using this decomposition inductively, we can unambiguously draw configurations of $!A$ by annotating each event by its copy index, and its justifier.

Example 2.25. The following is a representation of a configuration of $!\mathbb{B}$:



where, *e.g.* the events labeled \mathbf{tt}^1 respectively denote $\{\mathbf{q} \mapsto 0, \mathbf{tt} \mapsto 1\}$ and $\{\mathbf{q} \mapsto 3, \mathbf{tt} \mapsto 1\}$.


Using the additional space granted by $!A$, we can now represent programs evaluating their arguments multiple times. For instance, a valid strategy $\sigma : S \rightarrow ![(\mathbb{B} \rightarrow \mathbb{B}) \rightarrow \mathbb{B} \rightarrow \mathbb{B}]$ for the term $\lambda f^{\mathbb{B} \rightarrow \mathbb{B}} x^{\mathbb{B}}. f(fx)$ could contain, for $i, j \in \omega$ and injective function $\langle -, - \rangle : \omega^2 \rightarrow \omega$, a configuration:



This diagram exploits the representation introduced just above for configurations of $!A$: each event is specified through its label and copy index. The full strategy σ would comprise such configurations for all $i, j \in \omega$. For each positive event, σ must provide a copy index – this choice must be made globally, in a way that avoids collisions to maintain local injectivity of σ .

2.5.3. *Strategies up to copy indices.* Using the composition mechanism introduced before, one may define an interpretation of terms of IPA as concurrent strategies on expanded arenas. However, as observed above, such strategies not only carry information about the events they play and their causal history, but also the data of specific *copy indices* that seem largely irrelevant – *e.g.*, as above, the choice of an injection $\langle -, - \rangle$. In fact, for reasons familiar from AJM games [AJM00], strategies will not satisfy the laws of cartesian closed categories unless we consider them *up to their specific choice of copy indices*. Let us observe that on an example.

Example 2.26. Consider the term M to be $f : \mathbb{B} \rightarrow \mathbb{B} \rightarrow \mathbb{B} \vdash \lambda x^{\mathbb{B}}. f x x : \mathbb{B} \rightarrow \mathbb{B}$. Its interpretation could contain:



Because of the contraction, the Opponent events of indices j_1 and j_2 corresponding to different events in the arena trigger Player events corresponding to the same event in the arena. Local injectivity is ensured by the functions $2n$ and $2n + 1$ having disjoint codomain.

Likewise, consider two terms, with chosen configurations of their strategies:

$$[[\lambda x^{\mathbb{B}} y^{\mathbb{B}}. x]] : ![[\mathbb{B} \rightarrow \mathbb{B} \rightarrow \mathbb{B}]] \quad [[\lambda x^{\mathbb{B}} y^{\mathbb{B}}. y]] : ![[\mathbb{B} \rightarrow \mathbb{B} \rightarrow \mathbb{B}]]$$



Then we have:

$$[[M]] \circ [[\lambda x^{\mathbb{B}} y^{\mathbb{B}}. x]] \quad : \quad ![[\mathbb{B} \rightarrow \mathbb{B}]] \qquad [[M]] \circ [[\lambda x^{\mathbb{B}} y^{\mathbb{B}}. y]] \quad : \quad ![[\mathbb{B} \rightarrow \mathbb{B}]]$$



These are required to be the same by the laws of cartesian closed categories, since $(\lambda f. M)(\lambda xy. x) =_{\beta} (\lambda f. M)(\lambda xy. y)$. However, they are not isomorphic as strategies.

In order to solve this mismatch, we first need to formalize what it means for two configurations of $!A$ to be the same up to the choice of copy indices.

Definition 2.27. Let $x, y \in \mathcal{C}(!A)$. A **reindexing iso** between x and y is an order-isomorphism:

$$\theta : x \cong y$$

which *preserves labels*: for all $\alpha \in x$, $\text{lbl } \alpha = \text{lbl } (\theta \alpha)$.

A reindexing iso $\theta : x \cong y$ is **positive** iff it preserves the copy index of negative events, *i.e.* for all $\alpha^- \in x$, $\text{ind } \alpha = \text{ind } (\theta \alpha)$. **Negative** reindexing isos are defined dually.

Intuitively, two configurations of $\mathcal{C}(!A)$ related by a reindexing iso are distinct representations of one *thick subtree* of A in the sense of Boudes [Bou09], *i.e.* a subtree of the arena with duplicated sub-arenas. Two strategies are to be identified iff they are isomorphic, with the commuting triangle to $!A$ being weakened to a commutation *up to reindexing iso* – in fact, it turns out to be simpler to strengthen that to *positive reindexing isos*. Altogether:

Definition 2.28. Let $\sigma : S \rightarrow !A$, $\tau : T \rightarrow !A$ be two strategies. A **weak morphism from σ to τ** is $f : S \rightarrow T$, such that the triangle

$$\begin{array}{ccc} S & \xrightarrow{f} & T \\ \sigma \searrow & & \swarrow \tau \\ & !A & \end{array}$$

commutes *up to positive symmetry*, in the sense that for all $x \in \mathcal{C}(S)$, the set:


$$\{(\sigma s, \tau (f s)) \mid s \in x\}$$

is a positive reindexing iso. If $f : S \rightarrow T$, $g : T \rightarrow S$ are two weak morphisms such that $g \circ f = \text{id}_S$ and $f \circ g = \text{id}_T$, we say that (f, g) is a **weak isomorphism**, and write $\sigma \approx \tau$ to mean that σ and τ are weakly isomorphic.

The two strategies of Example 2.26 are weakly isomorphic. And in fact, a consequence of the developments of this paper is that the natural interpretation of terms $\vdash M : A$ of PCF as strategies on $!A$ hinted at here *is* sound and computationally adequate (it is reasonable to expect the same statement for IPA to be true as well, but it does not follow from the results in this paper). However, proving it bumps into a significant difficulty: without further constraints on strategies, weak isomorphism is *not* a congruence. Indeed, strategies can

behave differently depending on Opponent’s choice of copy indices. For instance, composing the two weakly isomorphic $\llbracket M \rrbracket \odot \llbracket \lambda x^{\mathbb{B}} y^{\mathbb{B}}. x \rrbracket$ and $\llbracket M \rrbracket \odot \llbracket \lambda x^{\mathbb{B}} y^{\mathbb{B}}. y \rrbracket$ of Example 2.26 with

$$\llbracket \mathbb{B} \rightarrow \mathbb{B} \rrbracket \dagger \llbracket \mathbb{B} \rightarrow \mathbb{B} \rrbracket$$



yields, in the one hand, a strategy that calls its argument, and on the other, one that does not. Clearly, they are not weakly isomorphic. This is because the strategy above is not *uniform*: its behaviour not only depends on Opponent’s moves, but also on their copy index. A useful analogy is that of a program that looks up the address where it is loaded in memory, and uses that information to specify its behaviour.

In AJM games [AJM00], uniformity is ensured by equipping games with an equivalence relation on plays not unlike our reindexing isomorphisms, and then requiring strategies to satisfy closure properties with respect to it. Our approach to uniformity bears some resemblance with that, but the richer structure of our plays prevents us from following the AJM recipe directly. To conclude this section and transition to the next, we give a few ideas motivating our approach to uniformity as detailed in the next section.

In traditional game semantics, plays are sequences of moves, *i.e.* total orders. Hence, the mere fact that two plays s_1 and s_2 (necessarily of the same length) are in relation in AJM games informs a one-to-one correspondence between the moves appearing in s_1 and those appearing in s_2 . In contrast, in our present setting the states (in the game) are partially ordered *configurations* $x \in \mathcal{C}(!A)$, so an equivalence $x \sim_{!A} y$ between $x, y \in \mathcal{C}(!A)$ no longer informs such a one-to-one correspondence. This correspondence will then have to become primitive: the equivalence between two configurations must be *witnessed* by a specific isomorphism, such as a reindexing iso. We like indeed to think of the set of reindexing isomorphisms as forming a “proof-relevant” equivalence relation between configurations of $!A$ – it expresses not only which configurations are in relation, but also *how*.

How, then, to express that a strategy $\sigma : S \rightarrow !A$ is uniform? It is tempting, for $x, y \in \mathcal{C}(S)$, to simply set $x \sim_S y$ when $\sigma x \sim_{!A} \sigma y$ – the witness isomorphism $\sigma x \cong \sigma y$ then informing an isomorphism $x \cong y$. This yields a set of isomorphisms between configurations of S , which then may be asked to satisfy closure properties ensuring uniformity. However, it turns out that this is too much to ask: such a simple definition rejects perfectly uniform strategies, some of which are definable through IPA. Intuitively (and unlike for the deterministic sequential strategies of AJM games), a term/strategy may *play symmetric moves for non-symmetric reasons*. So we opt instead for a more intensional option: we endow strategies with their own “proof-relevant equivalence relation” expressing which configurations they *considers* equivalent, and how. This “uniformity witness” is part of strategies with symmetry, and propagated through composition and other constructions on strategies. It cannot be recovered uniquely (except for *innocent* strategies [CCW15], though this is out of scope for this paper) – Appendix A.1 contains a more detailed discussion along the lines of this paragraph, with examples.

We now go on to the formalization of these ideas.

3. THIN CONCURRENT GAMES

Dealing with *uniformity* requires us to replicate the construction of concurrent games in a more expressive setting, capable to express that certain events or configurations might be *symmetric*, *i.e.* interchangeable; making their indistinguishability part of the structure. *Event structures with symmetry* [Win07] were designed precisely to cope with such situations. In this section, we construct a replacement for CG based on those, enforcing uniformity of strategies over symmetric events, and hence supporting uniform replication.

In the previous section, a game was an event structure structure A , while a strategy was an event structure S labeled by A , *i.e.* a map $\sigma : S \rightarrow A$ in the category \mathcal{E} of event structures and maps between them – furthermore, composition of strategies was obtained by leveraging universal constructions in \mathcal{E} , *e.g.* pullback for interaction. In order to build games with symmetry, it is mathematically appealing to replicate the definitions and constructions, but this time based on the category \mathcal{E}_\sim of event structures with symmetry and maps between them (to be defined in Definition 3.6). It is also as economical as we can do with the generality we wish to give the model (and in particular, so that it supports IPA) – we have explored a number of simpler alternatives, which fail in various ways.

In Section 3.1, we first recall event structures with symmetry, and expand on the above methodology for enriching concurrent games with symmetry. Though mathematically appealing, pushing these guidelines bumps against significant obstacles, allowing for two main solutions: the *saturated* (or *fat*) and *thin* approaches (respectively appearing in conference papers [CCW14] and [CCW15]). After briefly reviewing the fat case, the rest of the section commits to the thin. In Section 3.2, we develop thin concurrent games, focusing on the problem of uniformity, which imposes the most constraints on their design. Finally in Section 3.3, we show that thin games with symmetry form a compact closed category. Then, before going further along the main narrative of the paper and constructing the cartesian closed category of Concurrent Hyland-Ong games in Section 4, we show that \mathbf{Tcg} also supports the construction of an AJM-style exponential modality.

3.1. Symmetry on event structures and games. In this first part, we review the main technical tool – event structures with symmetry – and introduce the main challenges in constructing games based on those.

3.1.1. *Event structures with symmetry.* Intuitively, symmetry on event structures should behave as an equivalence relation, but also satisfy bisimulation-like properties in order to ensure that symmetric states have the same futures (up to symmetry). Looking for a notion of symmetry on event structures satisfying these two aspects, a natural methodology is to instantiate known categorical constructions.

In their seminar paper [JNW93], Joyal, Nielsen and Winskel gave a categorical notion of bisimulation between objects E and F as a *span*

$$\begin{array}{ccc} & B & \\ \swarrow l & & \searrow r \\ E & & F \end{array}$$

where l and r are *open maps*, *i.e.* satisfy a path lifting condition formulated as a factorisation property. Winskel later defined an *event structure with symmetry* in [Win07] as an event structure E with a span of open maps as above (with $E = F$), additionally satisfying categorical formulations of the laws of equivalence relations – we omit details, opting below for a more concrete equivalent definition. Besides their algebraic genesis, event structures with symmetry have already proved adequate as a modeling framework: most notably, Hayman and Winskel proved in [HW08] that the universal characterization of the unfolding of *safe* Petri nets as event structures could be extended to *general* Petri nets, provided one unfolds to *event structures with symmetry* – the multiple tokens in one place yielding distinct yet *symmetric* copies of enabled transitions.

Symmetry on event structures can be also defined via *isomorphism families* [Win07].


Definition 3.1 (Isomorphism families and event structures with symmetry). Let A be an event structure and \tilde{A} be a set of bijections between configurations of A . Then, \tilde{A} is an **isomorphism family** on A if it satisfies:

- (*Groupoid*) The set \tilde{A} contains all identity bijections, and is stable under composition and inverse of bijections.
- (*Restriction*) For every bijection $\theta : x \simeq y \in \tilde{A}$ and $x' \in \mathcal{C}(A)$ such that $x' \subseteq x$, then the restriction $\theta \upharpoonright x'$ of θ to x' is in \tilde{A} . In particular, $\theta x' \in \mathcal{C}(A)$.
- (*Extension*) For every $\theta : x \simeq y \in \tilde{A}$ and extension $x \subseteq x' \in \mathcal{C}(A)$, there exists a (non-necessarily unique) $y \subseteq y' \in \mathcal{C}(A)$ and an extension $\theta \subseteq \theta'$ such that $\theta' : x' \simeq y' \in \tilde{A}$.


In this case the pair $\mathcal{A} = (A, \tilde{A})$ is called an **event structure with symmetry (ess)**. We will use $\mathcal{S}, \mathcal{T}, \mathcal{A}, \mathcal{B}, \dots$ to range over event structures with symmetry. If A additionally has polarities and bijections in \tilde{A} preserve them, we say that \mathcal{A} is an **event structure with symmetry and polarities (essp)**.

One natural way to think of an isomorphism family is as a *proof-relevant history-preserving bisimulation and equivalence relation* between configurations: the equivalence between configurations is witnessed by precise bijections between their events.

The following represents a rather trivial example of an event structure with symmetry:



with events $\{1, 2\}$, no non-trivial causality and all finite sets consistent, and symmetry comprising all bijections between configurations meaning that all events are interchangeable. In this diagram we use an arrow \longleftrightarrow to convey the information that events are symmetric. In general however, the information of symmetry is more contextual and cannot be represented that easily. For instance, one may consider an event structure with symmetry with events $\{1, 2, 1', 2'\}$, trivial causality and all finite sets consistent, and symmetry comprising all bijections between configurations that are subsets of two maximal bijections:



It is a good exercise to verify that this yields an event structure with symmetry. There one can observe that from the empty bijection one can put 2 and 2' in correspondance,

however if $(1, 1)$ is already in the bijection, then its extension with $(2, 2')$ forms the bijection between configurations $\{(1, 1), (2, 2')\}$ which is *not* in the isomorphism family, not being a subset of either of the bijections above. In that sense symmetry is *contextual*: whether two events can be interchanged depends on which events have already been exchanged. Being sets of bijections, symmetries on event structures are rather hard to picture. In this paper, most diagrams representing event structures with symmetry will only display the event structure (the observable actions) and keep the symmetry (witnessing uniformity) implicit.

The following lemma, easy consequence of the (*Restriction*) axiom, is important to keep in mind when manipulating symmetry:

Lemma 3.2. *Let \mathcal{A} be an ess and $\theta : x \simeq y \in \tilde{\mathcal{A}}$. Then, θ is an order-isomorphism.*

Hence, if $\theta \in \tilde{\mathcal{A}}$, we write $\theta : x \cong y$ (rather than just $x \simeq y$) to indicate that θ preserves and reflects the (implicit, inherited from $\leq_{\mathcal{A}}$) ordering on x and y . Instead of $\theta : x \cong y \in \tilde{\mathcal{A}}$, we will also often use the more compact notation $\theta : x \cong_{\tilde{\mathcal{A}}} y$; and we will refer to θ as a **symmetry** between x and y . Given a symmetry θ , we write $\text{dom } \theta$ and $\text{codom } \theta$ for its domain and codomain respectively.

3.1.2. *Constructions on ess.* In the interpretation of games, as discussed in Section 2.5.3 all non-trivial symmetries come from *replication*. Accordingly our key construction of ess comes from the discussion of Section 2.5.2 – the following is established by a direct verification.

Proposition 3.3. *Let A be an arena. Recall $!A$ (Definition 2.24) whose events are functions $\alpha : [a] \rightarrow \mathbb{N}$ with $a \in A$. The sets $\tilde{!A}$ of reindexing isos (see Definition 2.27), $\tilde{!A}_-$ of negative reindexing isos, and $\tilde{!A}_+$ of positive reindexing isos, are isomorphism families on $!A$.*

This will be at the heart of our construction of the cartesian closed category of concurrent Hyland-Ong games, in Section 4. We mention in passing another construction on event structures with symmetry, more reminiscent of *AJM games* [AJM00].

Definition 3.4. If \mathcal{A} is an ess, then $!_{\text{AJM}}\mathcal{A}$ has events, causality, consistency that of $\|\omega A$. Its isomorphism family comprises those bijections $\theta : \|_{i \in I} x_i \simeq \|_{j \in J} y_j$ such that there exists a permutation $\pi : I \simeq J$, and for all $i \in I$ a symmetry $\theta_i : x_i \cong_{\mathcal{A}} y_{\pi(i)}$ such that for all $(i, a) \in \|_{i \in I} x_i$, we have $\theta(i, a) = (\pi(i), \theta_i(a))$.

The reader familiar with AJM games will recognize the similarity with the definition of the equivalence relation between plays in an exponential games in the AJM setting. In this paper we focus on HO-style games, but it has become crucial in further work that our setting with symmetry supports AJM-style games as well – we will come back to that in Section 3.3.4. Throughout the paper we reserve the notation $!$ for HO-style replication, *i.e.* the ess $!A$ of Proposition 3.3 from an arena A . In contrast, the AJM-style exponential will always be denoted by $!_{\text{AJM}}$ to reflect its lesser importance in the present development.

Besides the above, ess support all the basic constructions on event structures.

Definition 3.5. Let \mathcal{A} and \mathcal{B} be ess. We build their **simple parallel composition** as $(A \parallel B, \tilde{\mathcal{A}} \parallel \tilde{\mathcal{B}})$ where $\tilde{\mathcal{A}} \parallel \tilde{\mathcal{B}}$ is the set of bijections of the form $\theta_1 \parallel \theta_2 : x \parallel y \simeq x' \parallel y'$ where $x, x' \in \mathcal{C}(A), y, y' \in \mathcal{C}(B), \theta_1 \in \tilde{\mathcal{A}}, \theta_2 \in \tilde{\mathcal{B}}$ and $\theta_1 \parallel \theta_2$ is defined as $(i, a) \mapsto (i, \theta_i(a))$.

If \mathcal{A} is a game with symmetry, its **dual** \mathcal{A}^\perp has the same isomorphism family $\tilde{\mathcal{A}}$ on A^\perp .

3.1.3. *Categorical structure.* Just as maps of event structures were at the heart of the construction of concurrent games, games with symmetry will make use of maps between ess.

Definition 3.6. Let \mathcal{A}, \mathcal{B} be event structures with symmetry. A map of event structures $f : \mathcal{A} \rightarrow \mathcal{B}$ **preserves symmetry** iff for all $\theta : x \cong_{\tilde{\mathcal{A}}} y$, the bijection $f\theta = \{(fa, fa') \mid (a, a') \in \theta\}$ is in $\tilde{\mathcal{B}}$. In that case, f is a **map of event structures with symmetry**, written $f : \mathcal{A} \rightarrow \mathcal{B}$. Event structures with symmetry and their maps form a category written \mathcal{E}_{\sim} .

Finally, if \mathcal{A}, \mathcal{B} have polarities, maps are required to preserve those as well.

In \mathcal{E}_{\sim} , morphisms can be compared *up to symmetry*, abstracting away from the comparison of morphisms *up to the choice of copy indices* of the previous section.

Definition 3.7. Let $f, g : \mathcal{A} \rightarrow \mathcal{B}$ be maps of event structures with symmetry. They are **symmetric** (written $f \sim_{\tilde{\mathcal{B}}} g$) when for all $x \in \mathcal{C}(\mathcal{A})$, the bijection $\{(fs, gs) \mid s \in x\}$ is in $\tilde{\mathcal{B}}$.

As usual, we write $f \sim g$ instead of $f \sim_{\tilde{\mathcal{B}}} g$ when $\tilde{\mathcal{B}}$ is clear from the context.

With this definition, we will be able to reformulate Definition 2.28 by requiring the triangle to commute up to symmetry, *i.e.* $\tau \circ f \sim \sigma$ (postponed for now).

3.1.4. *Pullbacks and pseudo-pullbacks.* The construction of CG relies crucially on properties of the category \mathcal{E} of event structures and their maps. In order to construct games with symmetry, it is appealing to attempt replicating the same constructions, but building on \mathcal{E}_{\sim} rather than \mathcal{E} . In other words, a *game with symmetry* would be an ess \mathcal{A} with polarities, and a strategy on \mathcal{A} would be a map $\sigma : \mathcal{S} \rightarrow \mathcal{A}$ between ess. A strategy *from \mathcal{A} to \mathcal{B}* would be $\sigma : \mathcal{S} \rightarrow \mathcal{A}^{\perp} \parallel \mathcal{B}$, and those would be composed by pullback as in Section 2, and so on.

A first obstacle in replicating those constructions is that unlike \mathcal{E} , the category \mathcal{E}_{\sim} does not have all pullbacks (a proof of that appears in Appendix A.2) – one may understand that by the fact that \mathcal{E}_{\sim} really may be more adequately regarded as enriched over equivalence relations: indeed, maps should be compared up to symmetry rather than on the nose. Hence, one is tempted to have the role of pullbacks played by a universal construction taking symmetry into account, such as *pseudo-pullbacks* or *bi-pullbacks*.

Definition 3.8. Let $f : \mathcal{A} \rightarrow \mathcal{C}$ and $g : \mathcal{B} \rightarrow \mathcal{C}$ in \mathcal{E}_{\sim} . A **pseudo-pullback** of f and g is

$$\begin{array}{ccc} & \mathcal{P} & \\ \Pi_1 \swarrow & & \searrow \Pi_2 \\ \mathcal{A} & \begin{array}{c} \sphericalangle \\ \sim \\ \sphericalangle \end{array} & \mathcal{B} \\ f \searrow & & \swarrow g \\ & \mathcal{C} & \end{array}$$

commuting *up to symmetry*, and such that for all $f' : \mathcal{X} \rightarrow \mathcal{A}$, $g' : \mathcal{X} \rightarrow \mathcal{B}$ such that $f \circ f' \sim g \circ g'$ there exists a *unique* map $h : \mathcal{X} \rightarrow \mathcal{P}$ such that $\Pi_1 \circ h = f'$ and $\Pi_2 \circ h = g'$.

It is a **bi-pullback** iff for all $f' : \mathcal{X} \rightarrow \mathcal{A}$, $g' : \mathcal{X} \rightarrow \mathcal{B}$ such that $f \circ f' \sim g \circ g'$ there is $h : \mathcal{X} \rightarrow \mathcal{P}$, *unique up to symmetry*, such that $\Pi_1 \circ h \sim f'$ and $\Pi_2 \circ h \sim g'$.

In particular, any pseudo-pullback is a bi-pullback. It turns out that \mathcal{E}_{\sim} has all pseudo-pullbacks [Win07], so one may opt to use those for composing strategies. Though sensible, this choice has technically heavy consequences – we will come back to them later. However, there is another possibility to compose strategies with symmetry: it may be that though \mathcal{E}_{\sim} does not have *all* pullbacks, it has all those required to compute the interaction of strategies. And indeed, using polarity it is easy to capture maps that interact well in the presence


of symmetry: just as receptivity prevents a strategy from refusing an Opponent move, \sim -receptivity prevents it from refusing to consider two Opponent moves to be symmetric.

Definition 3.9. If \mathcal{A} is a essp, a map of ess $\sigma : \mathcal{S} \rightarrow \mathcal{A}$ is \sim -receptive iff for all $\theta : x_1 \cong_{\tilde{\mathcal{S}}} x_2$, for all $x_1 \xrightarrow{s_1^-} c$ and $\sigma x_2 \xrightarrow{a_2^-} c$ such that $\sigma \theta \cup \{(\sigma s_1, a_2)\} \in \tilde{\mathcal{A}}$, there is a unique s_2 such that $\sigma s_2 = a_2$, and we have $\theta \cup \{(s_1, s_2)\} \in \tilde{\mathcal{S}}$.

Moreover, σ as above is **strong-receptive** if it is both receptive and \sim -receptive.

Dual \sim -receptive maps always have pullbacks in \mathcal{E}_{\sim} : intuitively, the polarity helps determining whose responsibility it is to set two events as symmetric. However, before we prove that, let us include an example illustrating the following fact: besides ensuring pullbacks, \sim -receptivity is key to ensure that strategies based on ess are indeed uniform.

Example 3.10. Recall the non-uniform strategy from Section 2:



Assume now there is an isomorphism family $\tilde{\mathcal{S}}$ on this event structure S such that the labelling map $\sigma : \mathcal{S} \rightarrow ![[\mathbb{B} \rightarrow \mathbb{B}]^{\perp} \parallel ![[\mathbb{B} \rightarrow \mathbb{B}]]$ is \sim -receptive.

By \sim -receptivity (since the identity on $\{q_4^{-,0}, q_2^{+,0}\}$ must be in $\tilde{\mathcal{S}}$), we must have that the bijection $\{q_4^{-,0}, q_2^{+,0}, q_1^{-,0}\} \cong \{q_4^{-,0}, q_2^{+,0}, q_1^{-,1}\}$ is in $\tilde{\mathcal{S}}$. However, only the left hand side part can be extended by $q_3^{+,0}$, absurd.

As \sim -receptivity is crucial to ensure uniformity, it will be required whether we wish to compute interaction via pullback or pseudo-pullback.

We now prove that as claimed above, pullbacks exist along dual \sim -receptive maps of ess. To define the isomorphism family for the pullback, we first notice that bijections on configurations of the (plain) pullback induce bijections on their projections:

Lemma 3.11. *Let $\sigma : S \rightarrow A$ and $\tau : T \rightarrow A$ be maps of event structures. Let $\theta : w \simeq z$ be a bijection, where $w, z \in \mathcal{C}(S \wedge T)$. There are (unique) bijections $\theta_S : \Pi_1 w \simeq \Pi_1 z$ and $\theta_T : \Pi_2 w \simeq \Pi_2 z$ satisfying $\Pi_1 \circ \theta = \theta_S \circ \Pi_1$ and $\Pi_2 \circ \theta = \theta_T \circ \Pi_2$. Moreover, the mapping $\theta \mapsto (\theta_S, \theta_T)$ is monotonic w.r.t. inclusion.*

Proof. By local injectivity, Π_1 defines a bijection $w \simeq \Pi_1 w$ and $z \simeq \Pi_1 z$. With this remark, θ_S is simply defined as $\Pi_1 \circ \theta \circ \Pi_1^{-1}$. The equation and uniqueness are by definition, and monotonicity is obvious. The definition of θ_T is symmetric. \square

For $\sigma : \mathcal{S} \rightarrow \mathcal{A}$ and $\tau : \mathcal{T} \rightarrow \mathcal{A}$ maps of ess, define $\tilde{\mathcal{S}} \wedge \tilde{\mathcal{T}}$ to contain those bijections $\theta : w \cong z$ such that $\theta_S : \Pi_1 w \cong \Pi_1 z \in \tilde{\mathcal{S}}$ and $\theta_T : \Pi_2 w \cong \Pi_2 z \in \tilde{\mathcal{T}}$. Bearing in mind the correspondence between configurations of $S \wedge T$ and secured bijections $x \simeq y$, there is an order-isomorphism between those bijections $\theta \in \tilde{\mathcal{S}} \wedge \tilde{\mathcal{T}}$ and commutative squares between secured bijections $x \simeq y$ and $x' \simeq y'$ (ordered by componentwise union):

$$\begin{array}{ccc}
x & \xrightarrow{\sigma} & \sigma x & = & \tau y & \xrightarrow{\tau} & y \\
\theta_S \circ \llbracket \sigma \rrbracket & & & & & & \theta_T \circ \llbracket \tau \rrbracket \\
x' & \xrightarrow{\sigma} & \sigma x' & = & \tau y' & \xrightarrow{\tau} & y'
\end{array}$$

This definition indeed yields a pullback in \mathcal{E}_{\sim} :

Let $\sigma : \mathcal{S} \rightarrow \mathcal{A}$ and $\tau : \mathcal{T} \rightarrow \mathcal{A}^\perp$ be \sim -receptive maps of ess. The set $\widetilde{S} \wedge \widetilde{T}$ is an isomorphism family on $S \wedge T$ and the ess $(S \wedge T, \widetilde{S} \wedge \widetilde{T})$ is a pullback in \mathcal{E}_{\sim} of σ and τ , written $\mathcal{S} \wedge \mathcal{T}$.

Proof idea. The only difficulty is in proving the (*Extension*) condition. The argument exploits that an extension is positive for one of σ, τ , negative for the other – apply (*Extension*) for the positive, and \sim -receptivity for the other. The details are in Appendix B.1. \square

3.1.5. *Equivalences between strategies.* So, should we base our composition of strategies with symmetry on pullbacks or pseudo-pullbacks? To make up our mind another aspect must weight in: the universal property used to compose strategies impacts the equivalence up to which strategies may be considered. Indeed in CG strategies are naturally considered up to isomorphism (Definition 2.7), and that is of course preserved by pullbacks (pullbacks along isomorphic maps being isomorphic). In contrast, as pointed out in Section 2.5.3 in the presence of replication it is crucial to consider strategies *up to symmetry*, and there is no reason for an equivalence such as that of Definition 2.28 to be preserved by pullback. However, *pseudo-pullbacks* do preserve notions of *equivalence*:

Definition 3.12. Two maps of ess $\sigma : \mathcal{S} \rightarrow \mathcal{A}$ and $\sigma' : \mathcal{S}' \rightarrow \mathcal{A}$ are **weak-equivalent** iff there are maps of ess $f : \mathcal{S} \rightarrow \mathcal{S}'$, $g : \mathcal{S}' \rightarrow \mathcal{S}$ such that $g \circ f \sim \text{id}_{\mathcal{S}}$, $f \circ g \sim \text{id}_{\mathcal{S}'}$, and the two triangles below commute *up to symmetry*:

$$\begin{array}{ccc}
\mathcal{S} & \xrightarrow{f} & \mathcal{S}' \\
\sigma \searrow & \sim & \swarrow \sigma' \\
& \mathcal{A} &
\end{array}
\qquad
\begin{array}{ccc}
\mathcal{S}' & \xrightarrow{g} & \mathcal{S} \\
\sigma' \searrow & \sim & \swarrow \sigma \\
& \mathcal{A} &
\end{array}$$

They are **strong-equivalent** iff the triangles further commute on the nose.

It follows from their definition that pseudo-pullbacks preserve strong equivalence, and that bi-pullbacks preserve weak equivalence. Being a particular case of bi-pullbacks, pseudo-pullbacks preserve both. Weak equivalence looks like the notion of Definition 2.28 and indeed can serve to compare strategies up to their choice of copy indices (we will come back later to the subtle differences between weak equivalence and Definition 2.28), while the requirement that the triangles commute on the nose makes strong equivalence look less obviously relevant for that purpose. In any case, the take-away message for this discussion seems to be: *just use pseudo-pullbacks for interaction!*

Mathematically, this seems an elegant approach to concurrent games with symmetry: as \mathcal{E}_{\sim} is enriched over equivalence relations, in building concurrent games with symmetry it is natural to mimic the constructions of CG from \mathcal{E} using the corresponding operations on ess that take account of the enrichment. Indeed, this natural solution was the basis for our first account of symmetry in concurrent games [CCW14]. It works as a general framework (indeed [CCW14] contains a construction of HO and AJM-style games), but suffers from significant drawbacks concerning its applicability for further semantic purposes.




FIGURE 6. Saturated interpretation of the strategies of Example 2.26

3.1.6. *Saturated uniformity.* Let us investigate some consequences of computing interactions via pseudo-pullbacks. Consider maps of $\text{essp } \sigma : \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$ and $\tau : \mathcal{T} \rightarrow \mathcal{B}^\perp \parallel \mathcal{C}$ (regarded as strategies with symmetry – though we leave for later the precise definition of those), and say that mimicing the construction in Section 2.4.2, we compute the pseudo-pullback:

$$\begin{array}{ccc}
 & \mathcal{T} \otimes \mathcal{S} & \\
 \Pi_1 \swarrow & \downarrow & \searrow \Pi_2 \\
 \mathcal{S} \parallel \mathcal{C} & \sim & \mathcal{A} \parallel \mathcal{T} \\
 \sigma \parallel \mathcal{C} \searrow & & \swarrow \mathcal{A} \parallel \tau \\
 & \mathcal{A} \parallel \mathcal{B} \parallel \mathcal{C} &
 \end{array}$$

Unlike pullbacks, pseudo-pullbacks commute *up to symmetry*. Consequently, interaction via pseudo-pullback allows one to synchronize configurations that do not match on the nose, but *up to symmetry*. Accordingly, the pseudo-pullback analogue of Proposition 2.16 says that configurations in $\mathcal{C}(\mathcal{T} \otimes \mathcal{S})$ correspond to the data of $x_S \parallel x_C \in \mathcal{C}(\mathcal{S} \parallel \mathcal{C})$, $x_A \parallel x_T \in \mathcal{C}(\mathcal{A} \parallel \mathcal{T})$ and $\theta \in \tilde{\mathcal{A}} \parallel \tilde{\mathcal{B}} \parallel \tilde{\mathcal{C}}$ such that:

$$x_S \parallel x_C \xrightarrow{\sigma} x_A \parallel x_B \parallel x_C \cong_{\tilde{\mathcal{A}} \parallel \tilde{\mathcal{B}} \parallel \tilde{\mathcal{C}}}^{\theta} y_A \parallel y_B \parallel y_C \xleftarrow{\tau} y_A \parallel y_T$$

It carries a symmetry $\theta \in \tilde{\mathcal{A}} \parallel \tilde{\mathcal{B}} \parallel \tilde{\mathcal{C}}$ which mediates between configurations of S and T not quite matching on the game, but can also apply a symmetry on the visible output in the game. Everything becomes “up to symmetry”, including the visible actions of the interaction. Concretely, because of that, composition based on pseudo-pullbacks has a “saturation” effect: for $x \in \mathcal{C}(S)$ in a saturated $\sigma : S \rightarrow A$, for all $\theta : \sigma x \cong_{\tilde{\mathcal{A}}} y'$ there must be $\varphi : x \cong_S y$ such that $\sigma y = y'$. In other words, if a saturated strategy is prepared to play a move, by necessity it is also prepared to play non-deterministically all symmetric moves. To behave well with pseudo-pullback-based composition, all strategies need to be saturated, even copycat [CCW14]. We display in Figure 6 the saturated strategy corresponding to the two diagrams at the end of Example 2.26 – to alleviate the diagram we only represent a few binary conflicts, but in reality all positive events are pairwise incompatible and symmetric: the strategy chooses non-deterministically one copy index and plays it.

The reader with a background in game semantics will recognize here a phenomenon familiar from AJM games: in [BDER97] was introduced a variant of AJM games where strategies were similarly required to be saturated under the action of the equivalence relation. Independently of the specific purposes of [BDER97] (which required it for technical reasons), saturation provides a different approach to uniformity in AJM games than the traditional one [AJM00]. While it has remained rather marginal in game semantics, some recent works

based on AJM games are actually based on saturation [AJ09, VJA18], as it provides a slightly simpler mathematical foundation for uniformity.

While mathematically natural, saturation in concurrent games has some significant drawbacks. First of all, it makes strategies more “abstract”: even the underlying event structure becomes impossible to represent faithfully. As copy indices are chosen non-deterministically, all (non-trivial) strategies fail determinism in the usual sense of concurrent games [Win12] and it is tricky to recover a working notion of determinism *up to symmetry*. More generally, notions in CG that rely on a concrete analysis of the shape of conflict (such as *single-threadedness* – see Definition 4.8 – or all our further work on *innocence* [CCW15]) are hard to accommodate with saturated symmetry as the non-deterministic choice of symmetric events pollutes the genuine dynamic behaviour of the strategy. Last but not least, the saturated framework is not conservative over CG: one cannot “forget symmetry” functorially as composition mechanisms inherently involve it. This means that developments in CG are less likely to adapt transparently to saturated symmetry.

For these reasons, we have found it necessary to develop an alternative approach to uniformity (first presented in [CCW15]), called *thin concurrent games* as opposed to *saturated*, which is conservative over CG: in particular, the core components of the categorical structure (composition, copycat) match those in CG. In that respect, it is analogous to the traditional AJM way of handling uniformity [AJM00]. This has been a subtle endeavour, but the effort pays off: conservativity over CG means that many developments on CG extend transparently with thin symmetry, undisturbed by the uniformity requirements. Such successes include of course concurrent innocence [CCW15], but also the more recent probabilistic extension [CCPW18], and other works as of now unpublished. So the following construction of **Tcg** is hard work, but experience tells that the subtleties involved can be mostly abstracted away when relying on **Tcg**, and one can work like in CG simply with the additional proof obligations that further structure should be invariant under symmetry.

3.2. Thin concurrent games. Rather than changing the compositional machinery to preserve weak equivalence, we instead seek further restrictions on strategies for their standard interaction pullback to preserve it. Interestingly the answer turns out to be a minimality condition mirroring saturation, suggesting a sort of duality that is still under investigation.

3.2.1. Thinness. Our starting point is the fact that, unfortunately, weak equivalence of \sim -receptive maps is *not* preserved under composition or pullback – a counter-example appears in Appendix A.3. Our interpretation of that failure is that the (*Extension*) axiom is too permissive for strategies. Indeed, consider $\sigma : \mathcal{S} \rightarrow \mathcal{A}$ a \sim -receptive map of essps, and $\theta : x \cong_{\mathcal{S}}^{\theta} y$, suppose further that x extends by some positive event $s^+ \in S$ (strategies will be strong-receptive, as such extensions by negative events are inconsequential and entirely governed by the game). By (*Extension*), we then have some s' with:

$$\begin{array}{ccc} x \cup \{s\} & \cong_{\mathcal{S}}^{\theta'} & y \cup \{s'\} \\ \cup & & \cup \\ x & \cong_{\mathcal{S}}^{\theta} & y \end{array}$$

The crux of the issue is in how this s' is chosen. For saturated strategies the set of available choices for s' is *canonical*: there are as many s' as there are symmetric events in the game. In contrast, in the counter-example of Appendix A.3, there is no canonical choice

(more specifically, the counter-example exploits a situation where no valid assignment from the choices for s to the choices for s' can be described globally as a map of event structures).

This need of a canonical choice for positive extensions motivates the definition:

Definition 3.13. An essp \mathcal{A} is **thin** if for $\theta : x \cong_{\tilde{\mathcal{A}}} y$ and $x \xrightarrow{a_1^+} c$, there is a *unique* $y \xrightarrow{a_2^+} c$ s.t.

$$x \cup \{a_1\} \stackrel{\theta \cup \{(a_1, a_2)\}}{\cong_{\tilde{\mathcal{A}}}} y \cup \{a_2\}$$

The canonicity of the choice of positive extensions is ensured by uniqueness. *Thinness* is the main ingredient of our forthcoming definition of \sim -strategies; the next objective in our construction being to show that indeed (under mild further conditions), the interaction between thin maps preserves weak equivalence (a phenomenon we will from now on refer to as *congruence*, for brevity). For that we first mention a slightly simpler equivalent formulation of thinness – the equivalence proof is an interesting exercise.

Lemma 3.14. *An essp \mathcal{A} is thin iff for all $x \in \mathcal{C}(\mathcal{A})$, for all $\text{id}_x \sqsubseteq^+ \theta \in \tilde{\mathcal{A}}$, then $\theta = \text{id}_y$ for some $x \sqsubseteq y \in \mathcal{C}(\mathcal{A})$.*

The \sim -strategies will be certain \sim -receptive $\sigma : \mathcal{S} \rightarrow \mathcal{A}$, with \mathcal{S} thin (we will also say that σ is thin). An intuitive reading of thinness prompted by Lemma 3.14 is that σ will not take the initiative of declaring two positive events symmetric. As long as negative extensions of the symmetry remain identity bijections, positive extensions will too. As a consequence:

Lemma 3.15. *Let $\sigma : \mathcal{S} \rightarrow \mathcal{A}$ and $\tau : \mathcal{T} \rightarrow \mathcal{A}^\perp$ be thin \sim -receptive maps of essp.*

Then, $\tilde{\mathcal{S}} \wedge \tilde{\mathcal{T}}$ is trivial (reduced to identities).

Proof. We prove by induction that all bijections in $\tilde{\mathcal{S}} \wedge \tilde{\mathcal{T}}$ are identities. Let $z \in \mathcal{C}(\mathcal{S} \wedge \mathcal{T})$ and assume id_z extends by (e, e') to $\theta \in \tilde{\mathcal{S}} \wedge \tilde{\mathcal{T}}$. Assume for instance $\Pi_2 e$ is positive in \mathcal{T} (the other case is similar). By construction $\text{id}_{\Pi_2 z}$ extends to $\theta_T = \Pi_2 \theta \in \tilde{\mathcal{T}}$ by positive events $(\Pi_2 e, \Pi_2 e')$, hence $\Pi_2 e = \Pi_2 e'$ and θ_T is the identity because $\tilde{\mathcal{T}}$ is thin. By local injectivity of Π_2 it follows that e and e' must be equal, or incompatible extensions of z . But if they are incompatible, by Lemma 2.15 (and Proposition 2.16) it means that $\Pi_1 e$ and $\Pi_1 e'$ are incompatible extensions of $\Pi_1 z$ mapping to the same event in the game, contradicting the \sim -receptivity of σ . Hence $e = e'$ and θ is the identity. \square

In other words, a closed interaction between thin \sim -receptive maps has a trivial symmetry. Of course that will not be the case for an *open* interaction between $\sigma : \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$ and $\tau : \mathcal{T} \rightarrow \mathcal{B}^\perp \parallel \mathcal{C}$, where the external Opponent may contribute new symmetric pairs of negative events on $\mathcal{A}^\perp \parallel \mathcal{C}$. Indeed, computing this interaction involves the pullback

$$\begin{array}{ccc} & \mathcal{T} \otimes \mathcal{S} & \\ \Pi_1 \swarrow & \downarrow & \searrow \Pi_2 \\ \mathcal{S} \parallel \mathcal{C} & & \mathcal{A} \parallel \mathcal{T} \\ \sigma \parallel \mathcal{C} \searrow & & \swarrow \mathcal{A} \parallel \tau \\ & \mathcal{A} \parallel \mathcal{B} \parallel \mathcal{C} & \end{array}$$

on which Lemma 3.15 does not apply, because $\text{id}_{\mathcal{C}^\perp} : \mathcal{C}^\perp \rightarrow \mathcal{C}^\perp$ and $\text{id}_{\mathcal{A}} : \mathcal{A} \rightarrow \mathcal{A}$ are *not* thin – indeed, \mathcal{A} could be *e.g.* the essp of Proposition 3.3, whose isomorphism family includes all reindexing isomorphisms. In our proof of congruence, it will be of great importance that we are able to consider a sub-structure of the open interaction above where the external

Opponent *does not contribute new symmetric pairs*, *i.e.* restrict symmetries so as arrive in the realm of Lemma 3.15. Concretely, we need a thin replacement $\mathcal{A}_+ \rightarrow \mathcal{A}$ for $\mathcal{A} \rightarrow \mathcal{A}$, and likewise for \mathcal{C} . But no such replacement exist without further conditions on games¹.

3.2.2. Thin concurrent games and \sim -strategies. The requirement above suggests that besides being essps, games with symmetry \mathcal{A} should feature a sub-symmetry $\tilde{\mathcal{A}}_-$ such that $\text{id}_{\mathcal{A}} : \mathcal{A}_- \rightarrow \mathcal{A}$ satisfies the conditions of Lemma 3.15, *i.e.* is thin and \sim -receptive. By duality, we also need $\tilde{\mathcal{A}}_+$ such that $\mathcal{A}_+^\perp \rightarrow \mathcal{A}^\perp$ is thin \sim -receptive. These polarised sub-symmetries, along with their desired properties, are nicely captured by the following definition².

Definition 3.16. A **thin concurrent game (tcg)** is an essp \mathcal{A} with two additional isomorphism families $\tilde{\mathcal{A}}_-$ and $\tilde{\mathcal{A}}_+$ on A such that:

- (a) The families $\tilde{\mathcal{A}}_+$ and $\tilde{\mathcal{A}}_-$ are subsets of $\tilde{\mathcal{A}}$,
- (b) If $\theta \in \tilde{\mathcal{A}}_+ \cap \tilde{\mathcal{A}}_-$ then θ is an identity bijection,
- (c) If $\theta \in \tilde{\mathcal{A}}_-$ and $\theta \sqsubseteq^- \theta' \in \tilde{\mathcal{A}}$ then $\theta' \in \tilde{\mathcal{A}}_-$,
- (d) If $\theta \in \tilde{\mathcal{A}}_+$ and $\theta \sqsubseteq^+ \theta' \in \tilde{\mathcal{A}}$ then $\theta' \in \tilde{\mathcal{A}}_+$.

where $\theta \sqsubseteq^- \theta'$ (resp. $\theta \sqsubseteq^+ \theta'$) means that $\theta \subseteq \theta'$ such that $\theta' \setminus \theta$ only contains pairs of negative (resp. positive) events. The triple $(\mathcal{A}, \tilde{\mathcal{A}}_-, \tilde{\mathcal{A}}_+)$ will be often written simply \mathcal{A} .

Thin concurrent games support the basic operations on essps: the **dual** is $(\mathcal{A}, \tilde{\mathcal{A}}_-, \tilde{\mathcal{A}}_+)^\perp = (\mathcal{A}^\perp, \tilde{\mathcal{A}}_+, \tilde{\mathcal{A}}_-)$, and the **simple parallel composition** is performed componentwise:

$$(\mathcal{A}, \tilde{\mathcal{A}}_-, \tilde{\mathcal{A}}_+) \parallel (\mathcal{B}, \tilde{\mathcal{B}}_-, \tilde{\mathcal{B}}_+) = (\mathcal{A} \parallel \mathcal{B}, \tilde{\mathcal{A}}_- \parallel \tilde{\mathcal{B}}_-, \tilde{\mathcal{A}}_+ \parallel \tilde{\mathcal{B}}_+).$$

Of course, thin concurrent games include our guiding example:

Proposition 3.17. *Let A be an arena. Then, $(!A, !\tilde{\mathcal{A}}_-, !\tilde{\mathcal{A}}_+)$ is a thin concurrent game.*

Through this example one may get a more concrete understanding of the sub-symmetries: $\tilde{\mathcal{A}}_+$ comprises the symmetries where only Player has performed non-trivial exchanges between moves (so “copy indices” of negative events are preserved), and dually for $\tilde{\mathcal{A}}_-$.

We mention right away two important properties of tcgs. Firstly, it follows as required that \mathcal{A}_- is thin (and likewise for \mathcal{A}_+^\perp). Indeed, using the characterisation of Lemma 3.14, consider $x \in \mathcal{C}(A)$, and $\text{id}_x \sqsubseteq^+ \theta \in \tilde{\mathcal{A}}_-$. But id_x is in all three isomorphism families and in particular in $\tilde{\mathcal{A}}_+$, so $\theta \in \tilde{\mathcal{A}}_+$ as well by (d). Hence, $\theta = \text{id}_y$ for some $y \in \mathcal{C}(A)$ by (b).

Secondly, it follows from the axiom that any $\theta \in \mathcal{A}$ can be factored uniquely:

Lemma 3.18 (Decomposition lemma). *For \mathcal{A} a tcg, we have an order-isomorphism:*

$$\begin{aligned} \tilde{\mathcal{A}}_- \times_A \tilde{\mathcal{A}}_+ &\rightarrow \tilde{\mathcal{A}} \\ (\theta^-, \theta^+) &\mapsto \theta^- \circ \theta^+ \end{aligned}$$

where $\tilde{\mathcal{A}}_- \times_A \tilde{\mathcal{A}}_+ = \{(\theta^-, \theta^+) \in \tilde{\mathcal{A}}_- \times \tilde{\mathcal{A}}_+ \mid \text{codom } \theta^+ = \text{dom } \theta^-\}$ is ordered by pairwise inclusion and $\tilde{\mathcal{A}}$ is ordered by inclusion.

¹In fact the counter-example of Appendix A.3 can be adapted to show that without further hypothesis on games, congruence fails even for thin strategies.

²This formulation was a later improvement on the original definition of [CCW15].

Proof. The map is clearly well defined because \tilde{A}_- and \tilde{A}_+ are included in \tilde{A} .

Injectivity. Assume we have $\theta = \theta_1^- \circ \theta_1^+ = \theta_2^- \circ \theta_2^+ : x \cong_{\tilde{A}} y$. By using groupoid laws we get that $\theta_1^+ \circ (\theta_2^+)^{-1} = (\theta_1^-)^{-1} \circ \theta_2^- : z_2 \cong z_1 \in \tilde{A}_- \cap \tilde{A}_+$ hence both are equal to the identity.

Surjectivity. By induction on $\theta \in \tilde{A}$ we build a preimage. Assume we have the decomposition of $\theta = \theta^- \circ \theta^+$, and that θ extends to $\theta' : x' \cong y'$ by a pair of fixed polarity, say positive. We use the extension axiom on θ^- to get $\theta^- \subseteq \theta'^- : z' \cong_{\tilde{A}_-} y'$. It follows that $\theta^+ \subseteq (\theta'^-)^{-1} \circ \theta' : x' \cong_{\tilde{A}_+} z'$ is a positive extension of θ^+ so it must belong to \tilde{A}_+ by (d). Hence $\theta' = \theta'^- \circ ((\theta'^-)^{-1} \circ \theta')$ provides the required decomposition.

Monotonicity. Monotonicity of the decomposition follows from uniqueness. \square

So the full \tilde{A} is actually redundant, and can be uniquely recovered from \tilde{A}_- and \tilde{A}_+ . This structure of a game with two isomorphism families is strongly reminiscent of Mellès' earlier approach to uniformity by bi-invariance under *two* group actions [Mel03]. This suggests that there is something intrinsic in this decomposition of symmetry as compositions of Opponent's reindexings and Player's reindexing. This structure, that alternation and sequentiality allows to keep hidden in AJM's equivalence relations, seems to become inevitable in a thin treatment of symmetry in games where sequentiality is not hard-wired.

Before going on to the proof of congruence, as we are finally in a position to formalize what we mean by *strategy with symmetry*, we start by doing that.

Definition 3.19. Let \mathcal{A} be a tgc. A **pre- \sim -strategy** on \mathcal{A} is a thin, \sim -receptive map of essp $\sigma : \mathcal{S} \rightarrow \mathcal{A}$.

It is a **\sim -strategy** if it is also receptive and courteous (*i.e.* $\sigma : S \rightarrow A$ is a strategy).

We are mainly interested in the interaction and composition of \sim -strategies, but we perform some of the developments with the slightly more permissive conditions of pre- \sim -strategies as it is occasionally useful (in particular when we model state in Section 5.3) to be able to compose “strategies” that are not courteous or receptive.

3.2.3. Congruence. We start by stating the property ensuring congruence: the *bi-pullback property*. Our next objective is to prove it for all interactions between pre- \sim -strategies.

Definition 3.20 (Bi-pullback property). Let $\sigma : \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$ and $\tau : \mathcal{T} \rightarrow \mathcal{B}^\perp \parallel \mathcal{C}$ be \sim -receptive maps of essps. Their interaction has the **bi-pullback property** iff the pullback

$$\begin{array}{ccc}
 & \mathcal{T} \otimes \mathcal{S} & \\
 \swarrow \Pi_1 & \downarrow & \searrow \Pi_2 \\
 \mathcal{S} \parallel \mathcal{C} & & \mathcal{A} \parallel \mathcal{T} \\
 \swarrow \sigma \parallel \mathcal{C} & & \swarrow \mathcal{A} \parallel \tau \\
 & \mathcal{A} \parallel \mathcal{B} \parallel \mathcal{C} &
 \end{array}$$

is a bi-pullback.

In essence, the definition of interaction of strategies via pullback expresses that given $x_S \in \mathcal{C}(S)$ and $x_T \in \mathcal{C}(T)$ that match on B (*i.e.* $\sigma x_S = x_A \parallel x_B$ and $\tau x_T = x_B \parallel x_C$) and are causally compatible (see Proposition 2.16), we can form their synchronization uniquely as some $z \in \mathcal{C}(T \otimes S)$ such that $\Pi_1 z = x_S \parallel x_C$ and $\Pi_2 z = x_A \parallel x_T$. The bipullback property extends this method to construct synchronized states: it asserts in essence that if such $x_S \in \mathcal{C}(S)$ and $x_T \in \mathcal{C}(T)$ only match up to symmetry (*e.g.* they fail to agree because of copy index mismatches), one can always find a “common ground”, both strategies changing

their choice of symmetric events, yielding some $x_S \cong_{\tilde{S}} x'_S$ and $x_T \cong_{\tilde{T}} x'_T$ where x'_S and x'_T now match on the nose and can be synchronized directly.

This process of tweaking “copy indices” to find a common ground can be done interactively, and – beyond the bi-pullback property – is at the heart of the congruence problem. To illustrate that and showcase the inductive process at play, we include the following example of transporting a weak equivalence through composition.

Example 3.21. Consider the two following strategies on $!\llbracket \mathbf{com} \rrbracket$

$$\begin{array}{ccc} \sigma_1 : !\llbracket \mathbf{com} \rrbracket & & \sigma_2 : !\llbracket \mathbf{com} \rrbracket \\ \mathbf{run}^{-,i} & & \mathbf{run}^{-,i} \\ \downarrow & & \downarrow \\ \mathbf{done}^{+,0} & & \mathbf{done}^{+,1} \end{array}$$

only differing by the choice of copy index for **done**. There is an obvious weak equivalence between them, call $\varphi : S_1 \rightarrow S_2$ the obvious invertible map. Consider now the following strategy τ (which, in IPA, represents $x : \mathbf{com} \vdash x ; x : \mathbf{com}$):

$$\begin{array}{ccc} !\llbracket \mathbf{com} \rrbracket & \vdash & !\llbracket \mathbf{com} \rrbracket \\ & & \mathbf{run}^{-,i} \\ & \swarrow & \vdots \\ \mathbf{run}^{+, (i,0)} & & \mathbf{done}^{+, (j,k)} \\ \downarrow & & \vdots \\ \mathbf{done}^{-, j} & & \vdots \\ \downarrow & & \vdots \\ \mathbf{run}^{+, (i, j+1)} & & \vdots \\ \downarrow & & \vdots \\ \mathbf{done}^{-, k} & & \vdots \end{array}$$

In order to build a weak equivalence between the resulting compositions $\tau \odot \sigma_1$ and $\tau \odot \sigma_2$, a reasonable first step is to build a weak equivalence between the interactions $\tau \otimes \sigma_1$ and $\tau \otimes \sigma_2$. In particular, given a configuration of $T \otimes S_1$, we should be able to build a corresponding configuration of $T \otimes S_2$. Consider *e.g.* the following configuration of $T \otimes S_1$.


$$\begin{array}{ccc} !\llbracket \mathbf{com} \rrbracket & \vdash & !\llbracket \mathbf{com} \rrbracket \\ & & \mathbf{run}^{-,i} \\ & \swarrow & \vdots \\ \mathbf{run}^{(i,0)} & & \mathbf{done}^{+, (0,0)} \\ \downarrow & & \vdots \\ \mathbf{done}^0 & & \vdots \\ \downarrow & & \vdots \\ \mathbf{run}^{(i,1)} & & \vdots \\ \downarrow & & \vdots \\ \mathbf{done}^0 & & \vdots \end{array}$$

where events on the left hand side are drawn without polarity, as they are synchronised between σ_1 and τ . By projections, we get configurations $x \in \mathcal{C}(S_1 \parallel !\llbracket \mathbf{com} \rrbracket)$ and $y \in \mathcal{C}(T)$ such that $(\sigma_1 \parallel !\llbracket \mathbf{com} \rrbracket) x = \tau y$ and such that the induced bijection is secured.


In order to construct a configuration in $T \otimes S_2$, it is natural to try and replace x with $\varphi(x)$ – and that would work out if φ was a *strong* equivalence. But as it is only a weak equivalence, we do not have $(\sigma_2 \parallel !\llbracket \mathbf{com} \rrbracket)(\varphi x) = \tau y$, only

$$(\sigma_2 \parallel !\llbracket \mathbf{com} \rrbracket)(\varphi x) \cong_{!\llbracket \mathbf{com} \rrbracket \parallel !\llbracket \mathbf{com} \rrbracket} \tau y$$

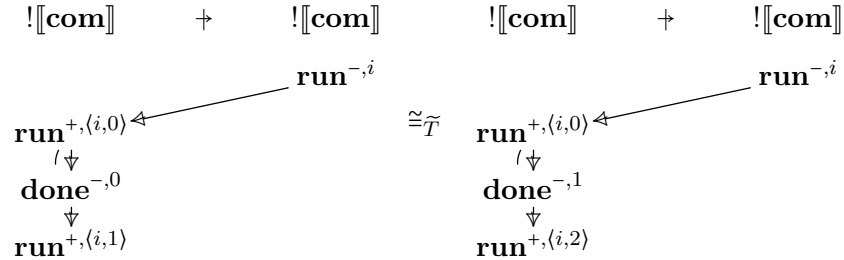
Here we observe the phenomenon hinted at above: the need to extract from $\varphi x \in \mathcal{C}(S_2)$ and $y \in \mathcal{C}(T)$, only matching up to symmetry, a valid configuration of $T \otimes S_2$. For our example, the only possibility is:



It appears that *both* φx and y had to change, in order to find an agreement as to the choice of copy indices. To compute it, we replay the interaction up to the first disagreement between σ_1 and σ_2 . By hypothesis, this disagreement yields symmetric configurations of the game. Hence by \sim -receptivity, \tilde{T} comprises a bijection:



By (Extension) in \tilde{T} , we know that this bijection can be extended to some:



Likewise, by \sim -receptivity of $\sigma_2 \parallel !\llbracket \mathbf{com} \rrbracket$ this extension is lifted to $\tilde{S}_2 \parallel !\llbracket \mathbf{com} \rrbracket$, and we then apply (Extension) on \tilde{S}_2 . And the process goes on, interactively between σ_2 and τ , until we get $x' \cong_{\tilde{S}_2 \parallel !\llbracket \mathbf{com} \rrbracket} \varphi x$ and $y' \cong_{\tilde{T}} y$ such that $(\sigma_2 \parallel !\llbracket \mathbf{com} \rrbracket)x' = \tau y'$ (which in our example, is the configuration of the interaction represented above).

Formalizing this interactive process of using \sim -receptivity on one strategy and extension on the other yields the following lemma:

Lemma 3.22 (Weak bipullback property). *Let $\sigma : \mathcal{S} \rightarrow \mathcal{A}$ and $\tau : \mathcal{T} \rightarrow \mathcal{A}^\perp$ be pre- \sim -strategies. Let $x \in \mathcal{C}(\mathcal{S})$ and $y \in \mathcal{C}(\mathcal{T})$ and $\theta : \sigma x \cong_{\tilde{\mathcal{A}}} \tau y$, such that the composite bijection*

$$x \stackrel{\sigma}{\cong} \sigma x \cong_{\tilde{\mathcal{A}}}^{\theta} \tau y \stackrel{\tau}{\cong} y$$

is secured. Then, there exists $z \in \mathcal{C}(\mathcal{S} \wedge \mathcal{T})$ along with $\theta_S : x \cong_{\tilde{\mathcal{S}}} \Pi_1 z$ and $\theta_T : \Pi_2 z \cong_{\tilde{\mathcal{T}}} y$, such that $\tau \theta_T \circ \sigma \theta_S = \theta$. Moreover, z is unique up to symmetry.

Proof. Uniqueness. Assume we have such (z, θ_S, θ_T) and $(z', \theta'_S, \theta'_T)$. Then it is easy to see that $\theta'_S \circ \theta_S^{-1} : \Pi_1 z \cong_{\tilde{\mathcal{S}}} \Pi_1 z'$ and similarly $\theta'_T \circ \theta_T^{-1} : \Pi_2 z \cong_{\tilde{\mathcal{T}}} \Pi_2 z'$. Those match on the game \mathcal{A} , so they induce a $z \cong z'$ in $\tilde{\mathcal{S}} \wedge \tilde{\mathcal{T}}$ as desired.

Existence. We proceed by induction on θ ; the base case is trivial. Assume θ extends by $(\sigma s, \tau t)$ to $\theta' : \sigma x' \cong \tau y'$. For instance, s is positive. We have $\theta_S : x \cong \Pi_1 z$ and x can be extended to x' by s , so by the extension property of the symmetry θ_S extends to $\theta'_S : x' \cong z'_S$. This means that $\tau \theta_T$ can be extended by symmetric negative (for T) events so by \sim -receptivity, θ_T can extend to $\theta'_T : z'_T \cong_{\tilde{\mathcal{T}}} y'$, with $\sigma z'_S = \tau z'_T$ by construction. Since the bijection $z'_S \cong z'_T$ is obviously secured, we get $z' \in \mathcal{C}(\mathcal{S} \wedge \mathcal{T})$ that satisfies our property. \square

Note that we did not need that σ and τ are thin – only \sim -receptivity. This statement is a step in the right direction, however the non-uniqueness of z (only up to symmetry) is problematic: it cannot be used to build maps, in particular it cannot be used to lift a weak equivalence $\sigma \rightarrow \sigma'$ to $\tau \otimes \sigma \rightarrow \tau \otimes \sigma'$. However, we will see now that if the interacting strategies are *thin*, we can use Lemma 3.15 to “tighten the screws” and show that then the choice of z is unique; from this observation the required map will follow.

We now prove the main technical result of this section, the *bi-pullback lemma*.

Lemma 3.23. *Let $\sigma : \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$ and $\tau : \mathcal{T} \rightarrow \mathcal{B}^\perp \parallel \mathcal{C}$ playing on tcgs. Then, their interaction has the bi-pullback property.*

Proof. Recall that for $f : \mathcal{X} \rightarrow \mathcal{S} \parallel \mathcal{C}$, $g : \mathcal{X} \rightarrow \mathcal{A} \parallel \mathcal{T}$ such that $\tau \circ g \sim_{\tilde{\mathcal{A}} \parallel \tilde{\mathcal{B}} \parallel \tilde{\mathcal{C}}} \sigma \circ f$, we need $h : \mathcal{X} \rightarrow \mathcal{T} \otimes \mathcal{S}$, unique up to symmetry, such that $\Pi_1 \circ h \sim_{\tilde{\mathcal{S}} \parallel \tilde{\mathcal{C}}} f$ and $\Pi_2 \circ h \sim_{\tilde{\mathcal{A}} \parallel \tilde{\mathcal{T}}} g$.

Uniquess up to symmetry follows from $\Pi_1 \circ h \sim_{\tilde{\mathcal{S}} \parallel \tilde{\mathcal{C}}} f$ and $\Pi_2 \circ h \sim_{\tilde{\mathcal{A}} \parallel \tilde{\mathcal{T}}} g$ and definition of the symmetry on the pullback. The main difficulty is *existence*. As hinted above, the trick is to apply Lemma 3.22, not on the raw interaction pullback, but on that between:

$$\begin{aligned} (\sigma \parallel (\mathcal{C}^\perp)_-) : \mathcal{S} \parallel (\mathcal{C}^\perp)_- &\rightarrow \mathcal{A}^\perp \parallel \mathcal{B} \parallel \mathcal{C}^\perp \\ (\mathcal{A}_- \parallel \tau) : \mathcal{A}_- \parallel \mathcal{T} &\rightarrow \mathcal{A} \parallel \mathcal{B}^\perp \parallel \mathcal{C} \end{aligned}$$

Write $\mathcal{T} \otimes' \mathcal{S}$ for this pullback, with projections $\Pi_1 : \mathcal{T} \otimes' \mathcal{S} \rightarrow \mathcal{S} \parallel (\mathcal{C}^\perp)_-$ and $\Pi_2 : \mathcal{T} \otimes' \mathcal{S} \rightarrow \mathcal{A}_- \parallel \mathcal{T}$. The underlying event structure of $\mathcal{T} \otimes' \mathcal{S}$ is the same as for $\mathcal{T} \otimes \mathcal{S}$, but the symmetry is tighter: intuitively, it is that where the external Opponent *does not change their copy indices*. In fact, the two maps above are *thin*, so by Lemma 3.15, the symmetry of $\mathcal{T} \otimes' \mathcal{S}$ is very tight indeed: it is restricted to identities. Nevertheless, for $x \in \mathcal{C}(X)$ we can apply Lemma 3.22 to $f x \in \mathcal{C}(\mathcal{S} \parallel \mathcal{C})$ and $g x \in \mathcal{C}(\mathcal{A} \parallel \mathcal{T})$, and get $z \in \mathcal{C}(\mathcal{T} \otimes \mathcal{S})$ – but its uniqueness up to symmetry now holds in $\mathcal{T} \otimes' \mathcal{S}$ with trivial symmetry, so z is unique. By uniqueness, this association induces a function $\psi : \mathcal{C}(X) \rightarrow \mathcal{C}(\mathcal{T} \otimes \mathcal{S})$ such that $\Pi_1(\psi x) \cong_{\tilde{\mathcal{S}} \parallel \tilde{\mathcal{C}}} f x$ and $\Pi_2(\psi x) \cong_{\tilde{\mathcal{A}} \parallel \tilde{\mathcal{T}}} g x$. It is then routine to verify that this function is monotonic, preserves cardinality and unions (hence it is generated by a map of event structures); and that it preserves symmetry – details are omitted. \square

This concludes our proof of congruence for interactions, *i.e.* the following corollary, simply proved by applying the bi-pullback property.

Corollary 3.24. *Weak equivalence is preserved by interactions of pre- \sim -strategies on thin concurrent games.*

It will of course follow immediately, once composition is defined in Section 3.3.1, that it preserves weak equivalence as well.

3.2.4. *Weak isomorphism.* Weak equivalence is a natural lax version of the isomorphism of strategies in CG in the presence of symmetry: all equalities in the definition of isomorphism are replaced with symmetry. However, as the reader may have noticed, there is a mismatch between weak equivalence and the *weak isomorphisms* of Definition 2.28, used in Section 2 to compare strategies up to copy indices. Let us start this discussion by recasting weak isomorphism in the context of thin concurrent games.

Definition 3.25. Let $\sigma : \mathcal{S} \rightarrow \mathcal{A}$ and $\sigma' : \mathcal{S}' \rightarrow \mathcal{A}$ be maps of ess, where \mathcal{A} is a tcg. A **positive morphism** from σ to σ' is a map of ess $f : \mathcal{S} \rightarrow \mathcal{S}'$ such that $\sigma' \circ f \sim_{\mathcal{A}_+} \sigma$.

We say that f is a **weak isomorphism** if it is furthermore invertible *on the nose*, *i.e.* there is $g : \mathcal{S}' \rightarrow \mathcal{S}$ such that $g \circ f = \text{id}_{\mathcal{S}}$ and $f \circ g = \text{id}_{\mathcal{S}'}$ – we write $\sigma \approx \sigma'$ for the corresponding equivalence relation.

The thin concurrent games presented in [CCW15] relied solely on weak equivalence to compare strategies up to symmetry; its refinement with *weak isomorphism* presented here came later. At first sight it looks like in switching from weak equivalence to weak isomorphisms we are trading an arguably mathematically canonical notion for one that is more concrete, but also possibly more ad-hoc. Of course, if weak isomorphism is a congruence, then the change is convenient. Indeed, having mediating maps be inverses on the nose makes the equivalence more conservative over plain event structures: if $\sigma \approx \sigma'$ as above, then S and S' are isomorphic as plain event structures (though of course the projection to the game does not commute on the nose). The tighter the equivalence is, the easier it is to transport properties and structure across. But in fact, we will see in this section that it is not a compromise at all: if $\sigma : \mathcal{S} \rightarrow \mathcal{A}$ and $\sigma' : \mathcal{S}' \rightarrow \mathcal{A}$ are \sim -strategies, then they are weakly isomorphic if and only if they are weakly equivalent³!

In order to prove that, the first step is to observe that the mediating maps being inverses on the nose comes for free, provided one insists on using *positive* weak equivalences.

Lemma 3.26. *Let $\sigma : \mathcal{S} \rightarrow \mathcal{A}$ and $\sigma' : \mathcal{S}' \rightarrow \mathcal{A}$ be pre- \sim -strategies, and $f : \mathcal{S} \rightarrow \mathcal{S}'$, $g : \mathcal{S}' \rightarrow \mathcal{S}$ forming a positive weak equivalence, *i.e.* $\sigma' \circ f \sim_{\mathcal{A}_+} \sigma$ and $\sigma \circ g \sim_{\mathcal{A}_+} \sigma'$.*

Then, f and g are actually inverse on the nose (and so form a weak isomorphism).

The proof of that is obvious in the light of the following lemma, which shows plainly the phenomenon at play.

Lemma 3.27. *Let $\sigma : \mathcal{S} \rightarrow \mathcal{A}$ be a pre- \sim -strategy on a tcg \mathcal{A} , and let $\theta : x \cong_{\mathcal{S}} y$ such that $\sigma\theta \in \tilde{A}_+$. Then, $x = y$ and $\theta = \text{id}_x$.*

³For that, thinness plays a crucial role: if $\sigma : \mathcal{S} \rightarrow \mathcal{A}^+ \parallel \mathcal{B}$ and $\tau : \mathcal{T} \rightarrow \mathcal{B}^+ \parallel \mathcal{C}$ are pre- \sim -strategies, their composition via pullback and pseudo-pullback (as in Section 3.1.6) are weakly equivalent as both are given by a bi-pullback, but they are certainly not weakly isomorphic.

Proof. By induction on θ . If θ is empty, it is clear. For $\text{id}_x \xrightarrow{(s,s')}{\mathcal{C}} \theta \in \tilde{\mathcal{S}}$. If s and s' are positive, then by thin $s = s'$. If negative, then $\sigma\theta \in \tilde{\mathcal{A}}_+ \cap \tilde{\mathcal{A}}_-$, hence is an identity. So θ is a negative extension of id_x , whose image in \mathcal{A} is an identity; hence it is an identity by \sim -receptivity. \square

It is of course not the case that every weak equivalence is positive. However, every map between pre- \sim -strategies is symmetric to one which does preserve the projection on the game up to positive symmetry: intuitively, if f sends negative $s^- \in S$ to fs with a different ‘‘copy index’’, we set $f's$ to the unique matching move positively symmetric to fs .

Let $\sigma : \mathcal{S} \rightarrow \mathcal{A}$ and $\sigma' : \mathcal{S}' \rightarrow \mathcal{A}$ be pre- \sim -strategies, and $f : \mathcal{S} \rightarrow \mathcal{S}'$ such that $\sigma' \circ f \sim_{\tilde{\mathcal{A}}} \sigma$. Then, there exists a unique $f' : \mathcal{S} \rightarrow \mathcal{S}'$ such that $f \sim_{\tilde{\mathcal{S}}} f'$, and $\sigma' \circ f' \sim_{\tilde{\mathcal{A}}_+} f$.

Sketch. For $x \in \mathcal{C}(S)$, the hypotheses give us $\theta_x : \sigma'(fx) \cong_{\tilde{\mathcal{A}}} \sigma x$, which we need to make positive. For that, we first use Lemma 3.18 to decompose θ_x as

$$\sigma'(fx) \cong_{\tilde{\mathcal{A}}_-}^{\theta_x^-} y \cong_{\tilde{\mathcal{A}}_+}^{\theta_x^+} \sigma x$$

such that $\theta_x = \theta_x^+ \circ \theta_x^-$. The key idea is then to *transport* fx over this negative symmetry θ_x^- , yielding $x' \cong_{\tilde{\mathcal{S}}} fx$ such that $\sigma'x' \cong_{\tilde{\mathcal{A}}_+} y$ (which can be done by induction on fx and θ_x^-), so that $\sigma x \cong_{\tilde{\mathcal{A}}_+} \sigma'x'$ as well. We then set $f'x$ to be x' , and extract from this a map $f' : \mathcal{S} \rightarrow \mathcal{S}'$.

The details appear in Appendix B.4. \square

This is a rather powerful result: it entails that every symmetry class of such weak maps from σ to σ' has a canonical representative, namely the unique equivalent map for which the projection to the game commutes up to positive symmetry.

As an immediate corollary, we have:

Corollary 3.28. *Let $\sigma : \mathcal{S} \rightarrow \mathcal{A}$ and $\sigma' : \mathcal{S}' \rightarrow \mathcal{A}$ be pre- \sim -strategies.*

Then, they are weakly isomorphic if and only if they are weakly equivalent.

Proof. If σ and σ' are weakly equivalent, from Lemma 3.2.4 one easily gets a *positive* weak equivalence. By Lemma 3.26, it is a weak isomorphism. The other direction is trivial. \square

As mentioned earlier, from Lemma 3.23 it will be obvious once composition is defined that it preserves weak equivalence. But from the above, each weak equivalence is canonically represented by a weak isomorphism; the induced equivalence relation on \sim -strategies is the same. In particular, it will follow just as directly that weak isomorphism is a congruence.

3.3. Categorical structure. Since the beginning of Section 3 we have focused on the crucial problem of *congruence*, which imposes the most constraints on the design of games with symmetry. Now that this is solved, we unfold the rest of the work required to build a core setting for game semantics, *i.e.* a *compact closed category* of tcgs and \sim -strategies.

3.3.1. Composition. As the previous section focused on interaction, it makes sense to start the construction by completing it to get *composition*. Composition of pre- \sim -strategies will be defined by simply enriching the composition of Section 2.4.3 with symmetry.

Let $\sigma : \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$ and $\tau : \mathcal{T} \rightarrow \mathcal{B}^\perp \parallel \mathcal{C}$ be pre- \sim -strategies. Ignoring symmetry, recall from Section 2 that $\tau \circ \sigma : T \circ S \rightarrow \mathcal{A}^\perp \parallel \mathcal{C}$ is obtained using *projection* (Definition 2.18): given $V = \{p \in T \otimes S \mid (\tau \otimes \sigma)p \notin B\}$ we set $T \circ S = T \otimes S \downarrow V$, and $\tau \circ \sigma$ to be the corresponding restriction of $\tau \otimes \sigma$. We now extend this in the presence of symmetry.

Lemma 3.29. *Let \mathcal{E} be an ess and $V \subseteq E$ closed under symmetry, in the sense that for all $\theta : x \cong_{\tilde{E}} y$, for all $e \in V \cap x$, we have $\theta e \in V$ as well. Then, defining*

$$\tilde{E} \downarrow V = \{\theta : x \simeq y \mid x, y \in \mathcal{C}(E \downarrow V), \exists \theta \subseteq \theta' \in \tilde{E}, \theta' : [x]_E \cong_{\tilde{E}} [y]_E\}$$

we have that $\tilde{E} \downarrow V$ is an isomorphism family, making $\mathcal{E} \downarrow V = (E \downarrow V, \tilde{E} \downarrow V)$ into an event structure with symmetry.

Proof. As usual the axiom (Groupoid) is clear. In this proof we abbreviate $[x]_E$ to $[x]$ for $x \in \mathcal{C}(E \downarrow V)$ for clarity reasons.

(Restriction) Let $\theta : x \simeq y \in \tilde{E} \downarrow V$, and $x_0 \in \mathcal{C}(E \downarrow V)$ such that $x_0 \subseteq x$. By definition there is $\theta \subseteq \theta' : [x] \cong_{\tilde{E}} [y]$. We have $[x_0] \subseteq [x]$. Therefore, by (Restriction) on \tilde{E} we have $\theta'_0 \subseteq \theta'$ with $\theta'_0 : [x_0] \cong_{\tilde{E}} y'_0$. Since V is closed under symmetry, $\theta'_0 \cap V^2 : x_0 \simeq y'_0 \cap V$ is still a bijection, which by definition is in $\tilde{E} \downarrow V$. It is clear by construction that $\theta'_0 \cap V^2 \subseteq \theta$.

(Extension) Let $\theta : x \simeq y \in \tilde{E} \downarrow V$, and $x \subseteq x_0 \in \mathcal{C}(E \downarrow V)$. By definition there is $\theta \subseteq \theta' : [x] \cong_{\tilde{E}} [y]$. We have $[x] \subseteq [x_0] \in \mathcal{C}(E \downarrow V)$, therefore by (Extension) for \tilde{E} there is $\theta'_0 : [x_0] \cong_{\tilde{E}} y'_0$. Again since V is closed under symmetry, $\theta'_0 \cap V^2 : x_0 \simeq y'_0 \cap V$ is still a bijection. By definition it is in $\tilde{E} \downarrow V$, and by construction it contains θ . \square

Given \sim -receptive $\sigma : \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$ and $\tau : \mathcal{T} \rightarrow \mathcal{B}^\perp \parallel \mathcal{C}$ (where \mathcal{A}, \mathcal{B} and \mathcal{C} are tcgs), $V = \{p \in \mathcal{S} \otimes \mathcal{T} \mid (\tau \otimes \sigma)p \notin \mathcal{B}\}$ is closed under symmetry; we can thus apply Lemma 3.29. Accordingly we set $\tilde{T} \circ \tilde{S}$ as $\tilde{T} \otimes \tilde{S} \downarrow V$, *i.e.* comprising bijections $\theta : x \simeq y$ such that there is

$$\theta \subseteq \bar{\theta} : [x]_{T \otimes S} \cong_{\tilde{T} \otimes \tilde{S}} [y]_{T \otimes S}.$$

This makes $\mathcal{T} \circ \mathcal{S} = (T \circ S, \tilde{T} \circ \tilde{S})$ an event structure with symmetry. In fact, we will show in Lemma 3.31 that if \mathcal{S} and \mathcal{T} are *thin*, the witnessing symmetry $\bar{\theta}$ is unique.

Summing up, we state:

Lemma 3.30. *If $\sigma : \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$ and $\tau : \mathcal{T} \rightarrow \mathcal{B}^\perp \parallel \mathcal{C}$ are pre- \sim -strategies, then*

$$\tau \circ \sigma : \mathcal{T} \circ \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{C}$$

is a map of ess.

Proof. We prove that $\tau \circ \sigma$ preserves symmetry. Let $\theta : x \cong_{\tilde{T} \circ \tilde{S}} y$. By definition, there is $\theta \subseteq \bar{\theta} : [x] \cong_{\tilde{T} \otimes \tilde{S}} [y]$. Then, $(\tau \otimes \sigma)\bar{\theta}$ is some

$$\theta_A \parallel \theta_B \parallel \theta_C : x_A \parallel x_B \parallel x_C \cong_{\tilde{A} \parallel \tilde{B} \parallel \tilde{C}} y_A \parallel y_B \parallel y_C$$

since $(\tau \otimes \sigma)$ preserves symmetry. But then $(\tau \circ \sigma)\theta$ is

$$\theta_A \parallel \theta_C : x_A \parallel x_C \cong_{\tilde{A} \parallel \tilde{C}} y_A \parallel y_C$$

which is a valid symmetry in $\tilde{A} \parallel \tilde{C}$ as required. \square

In order to get a notion of composition for \sim -strategies, we need to show that composition preserves thinness, and \sim -receptivity. We will treat them in that order.

The preservation of thinness under composition boils down to one crucial property: the fact that any symmetry between configurations of the composition has a *unique* witness in the interaction. Indeed, recall from Lemma 3.15 that the closed interaction between dual thin maps has a trivial symmetry. Of course composition is obtained via an *open* interaction, which does *not* have trivial symmetry as the external Opponent can contribute new isomorphic pairs of negative events. Nevertheless, whenever the interaction stays within

B the phenomenon above applies, and the symmetry is fixed: only the external Opponent can put in relation two non-identical events first. As a result, a bijection in the symmetry of the interaction is fully determined by its restriction to visible events:

Lemma 3.31 (Unique witness). *Let $\sigma : \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$ and $\tau : \mathcal{T} \rightarrow \mathcal{B}^\perp \parallel \mathcal{C}$ be pre- \sim -strategies. Recall the set of visible events of the interaction:*

$$V = \{e \in T \otimes S \mid (\tau \otimes \sigma) e \notin B\}$$

Let $\theta : x \cong_{\tilde{T} \otimes \tilde{S}} y$ and $\theta' : x \cong_{\tilde{T} \otimes \tilde{S}} y'$ such that $\theta \cap V^2 = \theta' \cap V^2$. Then $\theta = \theta'$.

Proof. By hypothesis, we have that $y \cap V = y' \cap V$. Note that $\theta \circ \theta'^{-1} : y' \cong y \in (\tilde{S} \parallel \tilde{C}) \wedge (\tilde{A} \parallel \tilde{T})$ and contains $\text{id}_{y \cap V}$. So necessarily, the projection of $\theta \circ \theta'^{-1}$ to \mathcal{A} and \mathcal{C} is an identity bijection. As a result, the symmetry $\theta \circ \theta'^{-1}$ actually belongs to $(\tilde{S} \parallel (\mathcal{C}_+^\perp)^\perp) \wedge (\mathcal{A}_- \parallel \tilde{T})$. This is a pullback of pre- \sim -strategies, so $\theta \circ \theta'^{-1}$ is an identity by Lemma 3.15, so $\theta = \theta'$. \square

Using this, we can prove that thinness is stable under composition.

Lemma 3.32. *For $\sigma : \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$ and $\tau : \mathcal{T} \rightarrow \mathcal{B}^\perp \parallel \mathcal{C}$ pre- \sim -strategies, $\tau \circ \sigma$ is thin.*

Proof. Let $z \in \mathcal{C}(T \circ S)$ such that id_z extends by positives (e, e') to $\theta : x \cong y \in \tilde{T} \circ \tilde{S}$ with witness $\bar{\theta} : [x] \cong_{T \otimes S} [y]$. Write θ_0 for $\bar{\theta} \setminus \{(e, e')\} : x_0 \cong y_0$. By hypothesis, θ_0 behaves like the identity on the visible part of x_0 . Hence, by Lemma 3.31, θ_0 is the identity on x_0 .

Since $\text{id}_{x_0} = \theta_0$ can be extended by (e, e') to $\bar{\theta}$ which is positive in $T \circ S$ we can assume eg. $\Pi_2 e$ and $\Pi_2 e'$ are positive in T . Hence $\Pi_2 \theta_0$ (which is also an identity) extends by positive $(\Pi_2 e, \Pi_2 e')$. Since τ is thin, we have $\Pi_2 e = \Pi_2 e'$ from which $e = e'$ follows (e and e' are positive), as desired. \square

We now focus on \sim -receptivity. Unlike thinness, it turns out that \sim -receptivity is not preserved by composition without further hypotheses, so pre- \sim -strategies are not stable under composition. To ensure preservation of \sim -receptivity one needs courtesy, however it is sometimes necessary to consider “strategies” that are not quite \sim -strategies – in particular for the interpretation of state in Section 5.3. So we introduce a more restricted form of courtesy sufficient to ensure preservation of \sim -receptivity.

Definition 3.33. Let $\sigma : \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$ be a pre- \sim -strategy. We say that σ is (A, B) -courteous iff for all $s_1 \rightarrow s_2$ in \mathcal{S} , if $\text{pol}_{\mathcal{S}}(s_2) = -$ (i.e. $\text{pol}_{\mathcal{A}^\perp \parallel \mathcal{B}}(\sigma s_2) = -$), then s_1 and s_2 map to the same A/B component.

We will also say that $\sigma : \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$ is **componentwise courteous** to mean that it is (A, B) -courteous, when \mathcal{A} and \mathcal{B} are clear from the context.

So σ is not necessarily courteous, but is not allowed to influence negative moves across components. As announced, we have the following.

Let $\sigma : \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$ and $\tau : \mathcal{T} \rightarrow \mathcal{B}^\perp \parallel \mathcal{C}$ be pre- \sim -strategies, such that σ is (A, B) -courteous and τ is (B, C) -courteous. Then, $\tau \circ \sigma$ is \sim -receptive and (A, C) -courteous.

Proof. The key ingredient of the proof is that thanks to componentwise courtesy of σ and τ , the immediate dependency of a negative event has to be a visible event in the same component (and not a neutral event); hence availability of a negative extension is entirely determined by the visible part of the interaction, and \sim -receptivity follows. Technical details are fairly tedious, and relegated to Appendix B.2. \square

So, componentwise courteous pre- \sim -strategy are stable under composition. The \sim -strategies are precisely those that are furthermore courteous and receptive as plain strategies, and we know from [CCRW17] that those are stable under composition; so \sim -strategies are preserved under composition as well. We will see in Section 3.3.3 that composition is associative (up to *strong* isomorphism) – however, before then we now focus on a crucial element of the compositional structure: its identity, the \sim -strategy copycat.

3.3.2. *Copycat.* Recall that the copycat strategy on game A is a labeled event structure:

$$\mathfrak{C}_A : \mathbb{C}_A \rightarrow A^\perp \parallel A$$

where \mathbb{C}_A has the same events as $A^\perp \parallel A$, but additional immediate causal links from negative events on one side to matching positive events on the other side. Consequently, configurations $x \in \mathcal{C}(\mathbb{C}_A)$ decompose as $x = x_1 \parallel x_2 \in \mathcal{C}(A^\perp \parallel A)$.

The following definition is forced by the requirement that the map \mathfrak{C}_A should be a map of ess, and that each symmetry should be an order-iso.

Definition 3.34. Let \mathcal{A} be a tcg. Given $x = x_1 \parallel x_2 \in \mathcal{C}(\mathbb{C}_A)$, $y = y_1 \parallel y_2 \in \mathcal{C}(\mathbb{C}_A)$, the set of *symmetries* between x and y (written $\mathbb{C}_{\tilde{\mathcal{A}}}$) comprises any bijection $\theta = \theta_1 \parallel \theta_2$ such that $\theta_1, \theta_2 \in \tilde{\mathcal{A}}$, and which is an order-iso (for the order on x, y induced by $\leq_{\mathbb{C}_A}$).

This definition is forced by necessity. However, to reason on such symmetries, it will be convenient to rely on a more high-level characterisation that does not explicitly require an order-isomorphism. To introduce it, recall first from [CCRW17] that configurations $x \in \mathcal{C}(\mathbb{C}_A)$ are exactly those $x_1 \parallel x_2 \in \mathcal{C}(A \parallel A)$ such that (with polarity as in $A \parallel A$):

$$x_2 \supseteq^- x_1 \cap x_2 \subseteq^+ x_1$$

Furthermore, it is observed in [Win13, CCRW17] that this relation between x_2 and x_1 is a partial order called the “Scott order”, written $x_2 \sqsubseteq_A x_1$. This order is of crucial importance in the construction and study of the bicategory CG.

The set $\mathbb{C}_{\tilde{\mathcal{A}}}$ is equivalently defined as comprising the bijections

$$\theta_1 \parallel \theta_2 : x_1 \parallel x_2 \simeq_{\tilde{\mathcal{A}}^\perp \parallel \tilde{\mathcal{A}}} y_1 \parallel y_2$$

satisfying the further condition that for all $a \in x_1 \cap x_2$, we have $\theta_1(a) = \theta_2(a)$.

Proof. Fairly straightforward, details are in Appendix B.3. \square

In other words, $\mathbb{C}_{\tilde{\mathcal{A}}}$ comprises those $\theta_1 \parallel \theta_2 \in \tilde{\mathcal{A}}^\perp \parallel \tilde{\mathcal{A}}$ such that $\theta_2 \supseteq^- \theta_1 \cap \theta_2 \subseteq^+ \theta_1$, *i.e.* $\theta_2 \sqsubseteq_{\tilde{\mathcal{A}}} \theta_1$. This justifies the notation $\mathbb{C}_{\tilde{\mathcal{A}}}$, as this agrees with the description of configurations of copycat via the Scott order. Wrapping up this construction, we state:

Let \mathcal{A} be a tcg. Then, writing $\mathbb{C}_A = (\mathbb{C}_A, \mathbb{C}_{\tilde{\mathcal{A}}})$, the map

$$\mathfrak{C}_A : \mathbb{C}_A \rightarrow A^\perp \parallel A$$

is a \sim -strategy.

Sketch. The details are fairly tedious and relegated to Appendix B.3. Interestingly it relies on \mathcal{A} being a tcg: for arbitrary essp \mathcal{A} , the set $\mathbb{C}_{\tilde{\mathcal{A}}}$ fails (*Extension*) (see Appendix A.4). \square

3.3.3. Compact closed structure. We now describe the categorical structure of the constructions above. As all our constructions are conservative extensions of those in CG (for which categorical laws are proved in details in [CCRW17]), the proofs of categorical laws boil down to showing that all the isomorphisms involved preserve symmetry. In fact all laws can be actually *deduced* directly from those in CG established in [CCRW17], by exploiting the representation of event structures with symmetry as spans (Section 3.1.1). As the details are at the same time unsurprising and rather tedious, we chose to omit them.

Proposition 3.35. *We have a category \mathbf{Tcg} having tcgs as objects, and as morphisms from \mathcal{A} to \mathcal{B} the \sim -strategies $\sigma : \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$, up to weak isomorphism.*

Furthermore, composition of componentwise courteous pre- \sim -strategies is associative.

We also write $\sigma : \mathcal{A} \xrightarrow{\mathbf{Tcg}} \mathcal{B}$ if $\sigma : \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$ is a \sim -strategy from \mathcal{A} to \mathcal{B} , keeping the \mathcal{S} anonymous. Note that the associativity and unity laws actually hold up to *strong isomorphism*, i.e. the projection to the game is preserved on the nose. The situation will be the same for all laws relative to the compact closed structure. Of course, this implies that they hold up to weak isomorphism as well.

We now equip \mathbf{Tcg} with a monoidal structure. For that, observe first that the category \mathcal{EP}_\sim of essps and maps preserving symmetry and polarities already has a monoidal structure, with tensor \parallel extending to maps in the obvious way. The tensor for \mathbf{Tcg} is defined on tcgs \mathcal{A} and \mathcal{B} to be $\mathcal{A} \otimes \mathcal{B} = \mathcal{A} \parallel \mathcal{B}$ (below Definition 3.16). Likewise, the tensor of \sim -strategies $\sigma_1 : \mathcal{S}_1 \rightarrow \mathcal{A}_1^\perp \parallel \mathcal{B}_1$ and $\sigma_2 : \mathcal{S}_2 \rightarrow \mathcal{A}_2^\perp \parallel \mathcal{B}_2$ is the map $\sigma_1 \otimes \sigma_2$, defined as the composite

$$\mathcal{S}_1 \parallel \mathcal{S}_2 \xrightarrow{\sigma_1 \parallel \sigma_2} (\mathcal{A}_1^\perp \parallel \mathcal{B}_1) \parallel (\mathcal{A}_2^\perp \parallel \mathcal{B}_2) \xrightarrow{\gamma} (\mathcal{A}_1 \parallel \mathcal{A}_2)^\perp \parallel (\mathcal{B}_1 \parallel \mathcal{B}_2)$$

where γ is the obvious relabeling. Again, as this is compatible with the tensor of strategies in CG, establishing that we have a bifunctor $- \otimes - : \mathbf{Tcg} \times \mathbf{Tcg} \rightarrow \mathbf{Tcg}$ just amounts to the fact that the corresponding isomorphisms in CG additionally preserve symmetry.

In order to construct the compact closed structure of \mathbf{Tcg} , we need to define all the required structural isomorphisms, such as e.g. the strategy $(\mathcal{A} \parallel \mathcal{B}) \parallel \mathcal{C} \rightarrow \mathcal{A} \parallel (\mathcal{B} \parallel \mathcal{C})$ expressing that the monoidal product is associative. For that we follow [CCRW17] and simply *lift* them from the corresponding isomorphisms from the monoidal structure of the category \mathcal{EP}_\sim of essps and maps preserving both polarity and symmetry between them.

Definition 3.36. Let $f : \mathcal{A} \rightarrow \mathcal{B}$ be a strong-receptive courteous polarities-preserving map between tcgs. Then its **lifting** is the \sim -strategy

$$\bar{f} = (\mathcal{A}^\perp \parallel f) \circ \mathbb{C}_{\mathcal{A}} : \mathbb{C}_{\mathcal{A}} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$$

which is a \sim -strategy from \mathcal{A} to \mathcal{B} (in particular, it is thin).

We then have:

Theorem 3.37. *The category \mathbf{Tcg} is compact closed.*

Proof. For completeness, we list here all structural morphisms for the symmetric monoidal structure of \mathcal{EP}_\sim .

$$\begin{aligned} \rho_{\mathcal{A}} &: \mathcal{A} \parallel 1 &\rightarrow \mathcal{A} \\ \lambda_{\mathcal{A}} &: 1 \parallel \mathcal{A} &\rightarrow \mathcal{A} \\ s_{\mathcal{A}, \mathcal{B}} &: \mathcal{A} \parallel \mathcal{B} &\rightarrow \mathcal{B} \parallel \mathcal{A} \\ \alpha_{\mathcal{A}, \mathcal{B}, \mathcal{C}} &: (\mathcal{A} \parallel \mathcal{B}) \parallel \mathcal{C} &\rightarrow \mathcal{A} \parallel (\mathcal{B} \parallel \mathcal{C}) \end{aligned}$$

These isomorphisms are then lifted to \sim -strategies.

$$\begin{aligned} \overline{\rho_A} & : & \mathcal{A} \otimes 1 & \xrightarrow{\mathbf{Tcg}} & \mathcal{A} \\ \overline{\lambda_A} & : & 1 \otimes \mathcal{A} & \xrightarrow{\mathbf{Tcg}} & \mathcal{A} \\ \overline{s_{\mathcal{A},\mathcal{B}}} & : & \mathcal{A} \otimes \mathcal{B} & \xrightarrow{\mathbf{Tcg}} & \mathcal{B} \otimes \mathcal{A} \\ \overline{\alpha_{\mathcal{A},\mathcal{B},\mathcal{C}}} & : & (\mathcal{A} \otimes \mathcal{B}) \otimes \mathcal{C} & \xrightarrow{\mathbf{Tcg}} & \mathcal{A} \otimes (\mathcal{B} \otimes \mathcal{C}) \end{aligned}$$

As before, coherence and naturality laws are easily adapted from those in CG. Finally, there are (copycat) \sim -strategies

$$\eta_{\mathcal{A}} : 1 \xrightarrow{\mathbf{Tcg}} \mathcal{A}^\perp \otimes \mathcal{A} \qquad \epsilon_{\mathcal{A}} : \mathcal{A} \otimes \mathcal{A}^\perp \xrightarrow{\mathbf{Tcg}} 1$$

satisfying the necessary equations up to isomorphism of \sim -strategies. \square

Finally, we adapt from [CCRW17] the *lifting lemma*, which we will use later. It characterises the effect of composition of a \sim -strategy via a lifted map.

Lemma 3.38. *Let $f : \mathcal{B} \rightarrow \mathcal{C}$ be a strong-receptive courteous polarities-preserving map between tcgs, and $\sigma : \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$ a \sim -strategy. Then, those \sim -strategies are isomorphic:*

$$\begin{aligned} \overline{f} \circ \sigma & : \mathbb{C}_{\mathcal{B}} \circ \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{C} \\ (\mathcal{A}^\perp \parallel f) \circ \sigma & : \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{C} \end{aligned}$$

Note that for $f : \mathcal{B}^\perp \rightarrow \mathcal{A}^\perp$ strong-receptive and courteous, we have the dual lifting $\overline{f}^\perp = (f \parallel \mathcal{B}) \circ \mathbb{C}_{\mathcal{B}} : \mathcal{A} \xrightarrow{\mathbf{Tcg}} \mathcal{B}$; and, by duality, the symmetric lemma to the above holds: for $\sigma : \mathcal{B} \xrightarrow{\mathbf{Tcg}} \mathcal{C}$, $\sigma \circ \overline{f}^\perp \cong (f \parallel \mathcal{C}) \circ \sigma$. Finally we note:

Lemma 3.39. *Let $f : \mathcal{A} \rightarrow \mathcal{B}$ be an isomorphism of tcgs – so both f and f^{-1} are strong-receptive courteous. Then, $\overline{f} \cong \overline{f^{-1}}^\perp$.*

All this work to build just a compact closed category may feel a little bit anticlimactic to the reader, and understandably so since by itself, a compact closed category can only be used to interpret a simple logic such as Multiplicative Linear Logic. For us however, it represents a crucial achievement in the construction of our framework for game semantics: it covers all the basic compositional properties, while remaining completely agnostic as to the language under study, its features, evaluation strategy, *etc.*

The focus of the present paper is the construction of the category of Concurrent Hyland-Ong games in the next section; but we consider \mathbf{Tcg} as a crucial contribution of the paper in itself, and it can certainly be used in other ways than just through Concurrent Hyland-Ong games. For instance, just like in Hyland-Ong games, replication in \mathbf{Cho} is ubiquitous: one cannot speak anymore of linear resources. In contrast, building directly on \mathbf{Tcg} (as we did in [CCPW18]) with a AJM-style exponential allows for a more traditional approach to replication (*i.e.* a model of ILL), and allows mixing linear and non-linear resources.

3.3.4. AJM-style exponentials. For completeness, and to emphasize the different ways in which one can work with \mathbf{Tcg} , we conclude this section by building an AJM-style exponential – we only show that we have an exponential modality, and leave the construction of a full model of Intuitionistic Linear Logic as out of scope. The remaining sections of this paper are independent of this construction.

Of course, to construct an AJM-style exponential in **Tcg** the first step is to extend the construction of $!_{\text{AJM}}\mathcal{A}$ in Definition 3.4 to tcgs. On that front the first news is bad: $!_{\text{AJM}}\mathcal{A}$ cannot be made a tcg in general. For instance, write $\ominus\oplus$ for the tcg with two events, one negative and one positive, the rest of the structure being trivial. Then, $!_{\text{AJM}}(\ominus\oplus)$ cannot be decomposed into $!_{\text{AJM}}(\ominus\oplus)_-$ and $!_{\text{AJM}}(\ominus\oplus)_+$ satisfying the axioms of tcgs. Intuitively, that is because $!_{\text{AJM}}(\ominus\oplus)$ forces symmetry constraints across polarities, *e.g.* $\{(1, \ominus)\} \cong \{(1, \oplus)\}$ and $\{(1, \oplus)\} \cong \{(2, \oplus)\}$ are incompatible, although they have dual polarities and their projections are compatible. The issue runs deeper than a lack of generality of tcgs: $\mathbb{C}_{!_{\text{AJM}}(\ominus\oplus)}$ fails the (*Extension*) axiom of isomorphism families (see Appendix A.4).

Fortunately, this phenomenon is circumscribed to *non-polarized* games with minimal events belonging to both players. Say that a game \mathcal{A} is **negative** (resp. **positive**) if all its minimal events have negative (resp. positive) polarity. Then we have:

Proposition 3.40. *Let \mathcal{A} be a negative tcg. We define $!_{\text{AJM}}\widetilde{\mathcal{A}}_-$ to include all bijections $\theta : \prod_{i \in I} x_i \simeq \prod_{j \in J} y_j$ such that there exists a permutation $\pi : I \simeq J$, and for all $i \in I$ a symmetry $\theta_i : x_i \simeq_{\widetilde{\mathcal{A}}_-} y_{\pi(i)}$ such that for all $(i, a) \in \prod_{i \in I} x_i$, we have $\theta(i, a) = (\pi(i), \theta_i(a))$. Likewise, $!_{\text{AJM}}\widetilde{\mathcal{A}}_+$ is defined in the same way, with $\widetilde{\mathcal{A}}_+$ replacing $\widetilde{\mathcal{A}}_-$ and π restricted to be the identity. Then, $(!_{\text{AJM}}\mathcal{A}, !_{\text{AJM}}\widetilde{\mathcal{A}}_-, !_{\text{AJM}}\widetilde{\mathcal{A}}_+)$ is a tcg.*

Proof. Straightforward. □

It is possible to have an AJM-style exponential in a non-polarized setting in the *saturated* games of [CCW14] mentioned at the beginning of the section. Hence the inability to cover uniform replication on non-polarized games is a restriction of **Tcg**, however we believe the gains from **Tcg** far outweigh this cost, especially since there are very few situations in game semantics that require uniform replication in non-polarized games (in fact, the only example we are aware of is a games model of classical Linear Logic – interestingly, this is also the paper in which the saturated variant of AJM games was introduced [BDER97]).

We now prove that $!_{\text{AJM}}$ is an exponential modality in the *negative* subcategory of **Tcg**, that we define now.

Proposition 3.41. *There is a symmetric monoidal subcategory **Tcg**₋ of **Tcg** having as objects the negative tcgs, and as morphisms the **negative** \sim -strategies $\sigma : \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$, where the negativity of σ means that the minimal events of \mathcal{S} are negative.*

Proof. The only non-trivial thing to check is that negative (pre)strategies are stable under composition, which will be proved as Lemma 4.3 in the next section. □

We finally conclude this section:

Theorem 3.42. *The operation $!_{\text{AJM}}$ extends to a linear exponential comonad:*

$$!_{\text{AJM}} : \mathbf{Tcg}_- \rightarrow \mathbf{Tcg}_-$$

Proof. First we define the functorial action: for $\sigma : \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$ between negative tcgs, we have $!_{\text{AJM}}\sigma : !_{\text{AJM}}\mathcal{S} \rightarrow (!_{\text{AJM}}\mathcal{A})^\perp \parallel !_{\text{AJM}}\mathcal{B}$ defined in the obvious way; functoriality is a variant of that of \otimes . As for the monoidal structure, the components are lifted from maps in \mathcal{EP}_{\sim} . More precisely, for any negative tcgs \mathcal{A}, \mathcal{B} there are maps defined in the standard way:

$$\begin{array}{ll} (!_{\text{AJM}}\mathcal{A})^\perp \otimes (!_{\text{AJM}}\mathcal{A})^\perp & \rightarrow & (!_{\text{AJM}}\mathcal{A})^\perp & \mathcal{A}^\perp & \rightarrow & (!_{\text{AJM}}\mathcal{A})^\perp \\ (!_{\text{AJM}}!_{\text{AJM}}\mathcal{A})^\perp & \rightarrow & (!_{\text{AJM}}\mathcal{A})^\perp & 1 & \rightarrow & (!_{\text{AJM}}\mathcal{A})^\perp \\ (!_{\text{AJM}}(\mathcal{A} \otimes \mathcal{B}))^\perp & \rightarrow & (!_{\text{AJM}}\mathcal{A})^\perp \otimes (!_{\text{AJM}}\mathcal{B})^\perp & !_{\text{AJM}}1 & \cong & 1 \end{array}$$

All those maps are courteous and strong-receptive, so by the “dual lifting” operation below Lemma 3.38 they yield \sim -strategies in the other direction, *e.g.* $!_{\text{AJM}}\mathcal{A} \dashrightarrow !_{\text{AJM}}\mathcal{A} \otimes !_{\text{AJM}}\mathcal{A}$ and $!_{\text{AJM}}\mathcal{A} \dashrightarrow !_{\text{AJM}}!_{\text{AJM}}\mathcal{A}$. It follows from Lemma 3.38 that these \sim -strategies are natural and satisfy the required coherence laws up to positive symmetry. \square

4. CONCURRENT HYLAND-ONG GAMES

We have constructed a compact closed category \mathbf{Tcg} , which is equipped to deal with the problem evoked at the end of Section 2. Using it, we can revisit (more formally) the interpretation sketched in Subsection 2.5. Exploiting the developments of the previous section, and in particular the fact that weak isomorphism is a congruence, it will follow that the two terms of Example 2.26 cannot be distinguished by any strategy in the model. Indeed, we will get a cartesian closed category supporting *e.g.* the interpretation of IPA.

From now on, all event structures are assumed to have binary conflict. All the operations we will consider on them (simple parallel composition, composition, interaction, *etc.*) have been established throughout the development to preserve that property.

4.1. The cartesian category \mathbf{Cho} . We now construct the category \mathbf{Cho} proper, of *Concurrent Hyland-Ong games*; and prove that it is *cartesian*. The **objects** of \mathbf{Cho} will be negative arenas, as in Definition 2.4 – with the further restriction that arenas should have a *countable* set of events, assumed from now on. The **morphisms** from arena A to arena B will be certain \sim -strategies from $!A$ to $!B$ (up to weak isomorphism):

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

Just as in standard HO games, we restrict strategies in order to satisfy the laws of a cartesian category. We will now inspect the different requirements of a cartesian category, and introduce the additional conditions on strategies as they are required.

4.1.1. Terminal object and negativity. First of all, a cartesian category has a terminal object. In our case, this will be the **empty arena $\mathbf{1}$** , defined as having an empty set of events – note that $!\mathbf{1}$ also has an empty set of events. However, as it is, $\mathbf{1}$ is not a terminal object. For each negative arena A , it is easy to see that the unique labelling function

$$\mathbf{e}_A : \mathbf{1} \rightarrow (!A)^\perp \parallel (!\mathbf{1})$$

is a \sim -strategy. Crucially, it is receptive since, by negativity of A , the minimal events of $(!A)^\perp$ are all positive. However, \mathbf{e}_A might not be *unique* from $!A$ to $\mathbf{1}$, as illustrated below.

Example 4.1. The following diagram represents a \sim -strategy from $!\mathbf{com}$ to $!\mathbf{1}$.

$$\begin{array}{c} (!\mathbf{com})^\perp \parallel (!\mathbf{1}) \\ \mathbf{run}^{+,0} \end{array}$$

The answer to this issue is clear: we need to require morphisms in \mathbf{Cho} to be negative, just as arenas. As in Proposition 3.41, $\sigma : \mathcal{S} \rightarrow (!A)^\perp \parallel (!B)$ is **negative** whenever the underlying event structure \mathcal{S} is negative, *i.e.* its minimal events are negative – note that this definition makes sense in general without symmetry, for a prestrategy $\sigma : \mathcal{S} \rightarrow A$.

We then easily have:

Proposition 4.2. *For any negative arena A , the empty \sim -strategy:*

$$\mathbf{e}_A : \mathbf{1} \rightarrow (!A)^\perp \parallel (!\mathbf{1})$$

is the unique negative \sim -strategy from $!A$ to $!\mathbf{1}$. In more generality, the only negative prestrategy $\sigma : S \rightarrow A^\perp \parallel \mathbf{1}$ for a negative game A is the empty prestrategy.

Proof. Immediate, as in a negative prestrategy $\sigma : S \rightarrow A^\perp \parallel \mathbf{1}$, any hypothetical minimal events in S have nowhere to map to. \square

Thus, in order to get a category of \sim -strategies with a terminal object, we will require that all \sim -strategies are negative. Clearly copycat is negative, along with all \sim -strategies obtained by lifting. Moreover, negative \sim -strategies are stable under composition. Since negativity makes sense without symmetry, we prove that in slightly greater generality.

Lemma 4.3. *Let $\sigma : S \rightarrow A^\perp \parallel B$ and $\tau : T \rightarrow B^\perp \parallel C$ be negative prestrategies (with A, B, C negative). Then, $\tau \circ \sigma$ is still negative.*

Proof. First, maps of event structures preserve minimal events: for $f : A \rightarrow B$ and a minimal in A , it follows easily from the axioms that fa is minimal in B . Hence, minimal events of $T \otimes S$ are projected to minimal events of $S \parallel C$ and $A \parallel T$. Take $e \in T \otimes S$ a minimal event. If $(\tau \otimes \sigma)e$ is in A , then $\Pi_1 e$ is a minimal event of S projected to a (necessarily positive) minimal event of A – absurd because σ is negative. Likewise, if $(\tau \otimes \sigma)e$ is in B , this contradicts the negativity of τ . So minimal events of $T \otimes S$ are visible and are in C .

Now, take any minimal event $e \in T \circ S$. Since minimal events of $T \otimes S$ are visible, e is also minimal in $T \otimes S$. By the previous remark, $(\tau \circ \sigma)e$ is in C and is minimal. It is also negative because C is negative. \square

Therefore, the category having arenas as objects and as morphisms from A to B , *negative \sim -strategies* $\sigma : S \rightarrow (!A)^\perp \parallel (!B)$ up to weak isomorphism, has a terminal object $\mathbf{1}$. We now investigate the existence of products.

4.1.2. Binary products and single-threadedness. For two arenas A and B , their **product** $A \times B$ is defined as the parallel composition $A \parallel B$, which is still a negative arena.

Projections. Note that there is an injection map of event structures with symmetry:

$$\begin{array}{ccc} \mathbf{i}_A : & !A & \rightarrow & !(A \times B) \\ & (\alpha : [a] \rightarrow \omega) & \mapsto & \left(\begin{array}{ccc} \alpha' : & [(1, a)] & \rightarrow & \omega \\ & (1, a') & \mapsto & \alpha(a') \end{array} \right) \end{array}$$

Likewise, there is $\mathbf{i}_B : !B \rightarrow !(A \times B)$. Using those, we define the **projections**

$$\varpi_A : \mathbb{C}_{!A} \rightarrow !(A \times B)^\perp \parallel !A \quad \varpi_B : \mathbb{C}_{!B} \rightarrow !(A \times B)^\perp \parallel !B$$

by lifting the injections, *i.e.* $\varpi_A = \overline{\mathbf{i}_A}^\perp$ and $\varpi_B = \overline{\mathbf{i}_B}^\perp$ (see Definition 3.36).

Pairing. Now, for negative \sim -strategies $\sigma : \mathcal{S} \rightarrow !A^\perp \parallel !B$ and $\tau : \mathcal{T} \rightarrow !A^\perp \parallel !C$, we wish to define their pairing $\langle \sigma, \tau \rangle$, a \sim -strategy from $!A$ to $!(B \times C)$. This \sim -strategy will simply be obtained by relabeling the parallel composition of \mathcal{S} and \mathcal{T} . In simple cases, it suffices to take the co-pairing:

$$\begin{aligned} \langle \sigma, \tau \rangle & : \mathcal{S} \parallel \mathcal{T} \rightarrow !A^\perp \parallel !(B \times C) \\ & = [(!A^\perp \parallel i_B) \circ \sigma, (!A^\perp \parallel i_C) \circ \tau] \end{aligned}$$

However, this is not always well-defined as a \sim -strategy. Indeed, it might fail local injectivity if some events in S and T have the same image in $!A^\perp$. As a first step towards the general construction of pairing, let us prove that this gives a well-defined \sim -strategy when the images of σ and τ are disjoint.

Lemma 4.4. *If negative \sim -strategies $\sigma : \mathcal{S} \rightarrow !A^\perp \parallel !B$ and $\tau : \mathcal{T} \rightarrow !A^\perp \parallel !C$ have disjoint codomain on $!A^\perp$, then $\langle \sigma, \tau \rangle$ as above is a negative \sim -strategy.*

Proof. First, we prove that it is a \sim -strategy. That it is a map of essps along with courtesy and thinness are direct verifications. Strong-receptivity needs further attention. Take $\theta : x_S \parallel x_T \cong y_S \parallel y_T \in \tilde{\mathcal{S}} \parallel \tilde{\mathcal{T}}$, write $\theta = \theta_S \parallel \theta_T$. Its projection to the game is

$$\langle \sigma, \tau \rangle \theta = ((!A^\perp \parallel i_B) \circ \sigma \theta_S) \uplus ((!A^\perp \parallel i_C) \circ \tau \theta_T)$$

which is a valid symmetry in $G = !A^\perp \parallel !(B \times C)$. Assume it extends by a pair (c_1^-, c_2^-) . Since dependency in the game is forest-shaped, there are unique $d_1 \rightarrow_G c_1^-$ and $d_2 \rightarrow c_2^-$, and since symmetries are order-preserving, we have $(d_1, d_2) \in \langle \sigma, \tau \rangle \theta$. But that means that it must be either in $((!A^\perp \parallel i_B) \circ \sigma) \theta_S$, or in $((!A^\perp \parallel i_C) \circ \tau) \theta_T$. We can then apply strong-receptivity of σ , τ , and the injection maps, to produce the extension to $\theta = \theta_S \parallel \theta_T$. \square

Now, we prove that this simple pairing behaves well *w.r.t.* projections.

Proposition 4.5. *Assume negative \sim -strategies $\sigma : \mathcal{S} \rightarrow !A^\perp \parallel !B$ and $\tau : \mathcal{T} \rightarrow !A^\perp \parallel !C$ as in the previous lemma. Then, we have (strong) isomorphisms:*

$$\varpi_B \odot \langle \sigma, \tau \rangle \cong \sigma \quad \varpi_C \odot \langle \sigma, \tau \rangle \cong \tau$$

Proof. Let us prove the first. More precisely, we prove that the interactions $\varpi_B \otimes \langle \sigma, \tau \rangle$ and $\mathfrak{c}_{!B} \otimes \sigma$ are isomorphic. This will entail by restriction an isomorphism between the corresponding compositions, and the latter is isomorphic to σ as copycat is the identity.

We establish the isomorphism between $\varpi_B \otimes \langle \sigma, \tau \rangle$ and $\mathfrak{c}_{!B} \otimes \sigma$ first for the plain event structures – by Lemma 2.12 of [CCRW17] it suffices to prove that they have an isomorphic domain of configurations. Using Proposition 2.16, we know that configurations of the event structure for the former interaction correspond to secured bijections

$$(x_S \parallel x_T) \parallel x_B \simeq y_A \parallel (y_B^1 \parallel y_B^2)$$

where $x_S \in \mathcal{C}(S)$, $x_T \in \mathcal{C}(T)$, $\langle \sigma, \tau \rangle (x_S \parallel x_T) = y_A \parallel (i_B y_B^1)$, and $y_B^1 \parallel y_B^2 \in \mathcal{C}(\mathfrak{C}_{!B})$, and where the bijection is the unique such that image of events through the labelings $\langle \sigma, \tau \rangle \parallel !B$ and $!A \parallel \varpi_B$ match. In particular, $\langle \sigma, \tau \rangle (x_S \parallel x_T)$ does not reach $!C$. But any minimal events of x_T are negative by negativity of τ , and hence must be in $!C$ (since A is negative). Therefore, x_T is empty. Getting rid of x_T yields a secured bijection corresponding to a configuration of the event structure of $\mathfrak{c}_{!B} \otimes \sigma$. This association is bijective, and yields the required isomorphism between domains of configurations. By construction, it is clear that this isomorphism preserves symmetry. \square

So, we know how to construct a pairing behaving well with projections, when the paired strategies happen to have a disjoint codomain. However, for arbitrary $\sigma : \mathcal{S} \rightarrow !A^\perp \parallel !B$ and $\tau : \mathcal{T} \rightarrow !A^\perp \parallel !C$, there might in general be collisions: events $s \in \mathcal{S}$ and $t \in \mathcal{T}$ such that $\sigma s = \tau t$. In such a case, the co-pairing as above fails local injectivity, and therefore does not correspond to a strategy. Fortunately, we can *relabel* moves of \mathcal{S} and \mathcal{T} , not changing their weak isomorphism class, to ensure that there are no such collisions. For that, we note that there are maps of event structures with symmetry

$$\iota_e : !A^\perp \rightarrow !A^\perp \quad \iota_o : !A^\perp \rightarrow !A^\perp$$

such that $\iota_e \sim_+ \iota_o \sim_+ \text{id}_{!A^\perp}$, but such that ι_e and ι_o have disjoint codomain. For definiteness, say that ι_e sends (necessarily positive) minimal events with copy index i to the same events with copy index $2i$, and preserves the index of other events. Likewise, ι_o follows the injection $i \mapsto 2i + 1$. These maps preserve the index of negative events, so that $\iota_e \sim_+ \iota_o \sim_+ \text{id}_{!A^\perp}$.

Given arbitrary $\sigma : \mathcal{S} \rightarrow !A^\perp \parallel !B$ and $\tau : \mathcal{T} \rightarrow !A^\perp \parallel !C$, we define:


$$\sigma_e = (\iota_e \parallel !B) \circ \sigma \quad \tau_o = (\iota_o \parallel !C) \circ \tau$$

From $\iota_e \sim_+ \iota_o \sim_+ \text{id}_{!A^\perp}$ it is obvious that $\sigma \approx \sigma_e$ and $\tau \approx \tau_o$, but σ_e and τ_o now have disjoint codomains: σ_e (resp. τ_o) only reaches indexing functions in $!A$ whose index for minimal events is even (resp. odd). Therefore, using Proposition 4.5, we define:


$$\langle \sigma, \tau \rangle = \langle \langle \sigma_e, \tau_o \rangle \rangle$$

The **pairing** of arbitrary negative \sim -strategies $\sigma : \mathcal{S} \rightarrow !A^\perp \parallel !B$ and $\tau : \mathcal{T} \rightarrow !A^\perp \parallel !C$ is defined as $\langle \sigma, \tau \rangle$. We have, as required, $\varpi_B \odot \langle \sigma, \tau \rangle = \varpi_B \odot \langle \langle \sigma_e, \tau_o \rangle \rangle \cong \sigma_e \approx \sigma$, and for the same reason $\varpi_C \odot \langle \sigma, \tau \rangle \approx \tau$. It is an immediate verification that $\langle -, - \rangle$ preserves weak isomorphism, so it will still make sense as an operation on the quotient category.

Example 4.6. Consider the copycat strategy $\mathfrak{c}_{![[\text{com}]]}$ on $![[\text{com}]]$.



Following the definition above, $\langle \mathfrak{c}_{![[\text{com}]]}, \mathfrak{c}_{![[\text{com}]]} \rangle$ is the \sim -strategy illustrated below.



As prescribed by the construction, the positive moves on the left hand side had to be relabeled to avoid the collision in the case where $i = j$.

$$\begin{array}{ccc}
\sigma_1 & : & !(\llbracket \mathbf{com} \rrbracket \times \llbracket \mathbf{com} \rrbracket) \quad \langle \varpi_l \odot \sigma_1, \varpi_r \odot \sigma_1 \rangle & : & !(\llbracket \mathbf{com} \rrbracket \times \llbracket \mathbf{com} \rrbracket) \\
& & \begin{array}{ccc} \mathbf{run}^{-,i} & & \mathbf{run}^{-,j} \\ \downarrow & \swarrow & \\ \mathbf{done}^{+,j} & & \end{array} & & \begin{array}{cc} \mathbf{run}^{-,i} & \mathbf{run}^{-,j} \end{array} \\
\sigma_2 & : & !(\llbracket \mathbf{com} \rrbracket \times \llbracket \mathbf{com} \rrbracket) \quad \langle \varpi_l \odot \sigma_2, \varpi_r \odot \sigma_2 \rangle & : & !(\llbracket \mathbf{com} \rrbracket \times \llbracket \mathbf{com} \rrbracket) \\
& & \begin{array}{ccc} \mathbf{run}^{-,i} & & \mathbf{run}^{-,j} \\ \downarrow & & \downarrow \\ \mathbf{done}^{+,0} & \sim & \mathbf{done}^{+,0} \end{array} & & \begin{array}{cc} \mathbf{run}^{-,i} & \mathbf{run}^{-,j} \\ \downarrow & \downarrow \\ \mathbf{done}^{+,0} & \mathbf{done}^{+,0} \end{array}
\end{array}$$

FIGURE 7. Failures to surjective pairing

Note that the above only displays the event part of the \sim -strategy $\langle \mathbf{c}!_{\llbracket \mathbf{com} \rrbracket}, \mathbf{c}!_{\llbracket \mathbf{com} \rrbracket} \rangle$, but its construction also equips it with a symmetry ensuring its uniformity.

Surjective pairing. In order to obtain a product, we also need to prove surjective pairing, that is, that for all $\sigma : S \rightarrow !A^\perp \parallel !(B \times C)$, we have:

$$\sigma \approx \langle \varpi_B \odot \sigma, \varpi_C \odot \sigma \rangle$$

However, as it stands, this is in general not the case.

Example 4.7. In Figure 7, we display on the left hand side two \sim -strategies $\sigma_1, \sigma_2 : !\mathbf{1} \xrightarrow{\mathbf{Tcg}} !(\llbracket \mathbf{com} \rrbracket \times \llbracket \mathbf{com} \rrbracket)$, and on the right hand side the corresponding distinct \sim -strategies obtained by projection and pairing.

We observe that surjective pairing fails for these strategies, as behaviours that span both components get erased through composition with the projections.

This is analogous as in standard Hyland-Ong games [Har99], where *single-threadedness* ensures that strategies treat independently events hereditarily caused by distinct minimal events. The definition is independent from symmetry, so we state it first in more generality.

Definition 4.8. Let $\sigma : S \rightarrow A$ be a prestrategy. We say that σ is **single-threaded** if it satisfies the following two conditions.

- (1) For any $s \in S$, $[s]$ has exactly one minimal event written $\text{init}(s)$.
- (2) Whenever $s_1 \# s_2$ in S , $\text{init}(s_1) = \text{init}(s_2)$.

Single-threaded \sim -strategies always satisfy surjective pairing.

Proposition 4.9. *Let $\sigma : S \rightarrow !A^\perp \parallel !(B \times C)$ be a single-threaded \sim -strategy. Then:*

$$\sigma \approx \langle \varpi_B \odot \sigma, \varpi_C \odot \sigma \rangle$$

Proof. First of all, we define two subsets of S as follows:

$$\begin{aligned}
S_B &= \{s \in S \mid \sigma(\text{init}(s)) \in B\} \\
S_C &= \{s \in S \mid \sigma(\text{init}(s)) \in C\}
\end{aligned}$$

(we abuse notations slightly with $\in B, \in C$).

By single-threadedness, S_B and S_C are disjoint and down-closed, with no immediate conflict spanning both components – in other words, $S = S_B \uplus S_C$. They are obviously still event structures. The restrictions of σ (along with a simple relabeling to $!B/!C$)

$$\sigma_B : S_B \rightarrow !A^\perp \parallel !B \quad \sigma_C : S_C \rightarrow !A^\perp \parallel !C$$

are receptive and courteous, *i.e.* are strategies.

This decomposition also works at the level of symmetries. Any $\theta \in \tilde{S}$ preserves S_B and S_C . Indeed if $(s_B, s_C) \in \theta$, then $(\text{init}(s_B), \text{init}(s_C)) \in \theta$ as well: absurd, since one maps to $!B$ and the other to $!C$. It follows that $\theta = \theta_B \uplus \theta_C$ where θ_B and θ_C are bijections between configurations of S_B and S_C respectively. The set of restrictions to S_B (resp. S_C) of symmetries in \tilde{S} yields a set of bijections between configurations of S_B (resp. S_C), which is easily checked to satisfy the conditions for an isomorphism family \tilde{S}_B (resp. \tilde{S}_C). The labeling functions σ_B and σ_C preserve symmetry. Strong-receptivity and thinness follow directly from those for σ , so σ_B and σ_C are \sim -strategies.

By construction, σ_B and σ_C have disjoint codomain; so we can form their pairing $\langle\langle \sigma_B, \sigma_C \rangle\rangle$ without relabeling. Then, the obvious bijection $S = S_B \uplus S_C \cong S_B \parallel S_C$ is an isomorphism of event structures, preserves symmetry, and preserves labeling so as to yield an isomorphism of \sim -strategies:

$$\sigma \cong \langle\langle \sigma_B, \sigma_C \rangle\rangle$$

By Proposition 4.5, it follows that $\varpi_B \odot \sigma \cong \sigma_B$ and $\varpi_C \odot \sigma \cong \sigma_C$. But clearly, $\langle\langle \sigma_B, \sigma_C \rangle\rangle \approx \langle \sigma_B, \sigma_C \rangle$, and $\langle -, - \rangle$ preserves weak isomorphism, so we have surjective pairing. \square

So, single-threadedness ensures surjective pairing. It is clear that copycat \sim -strategies – and lifted \sim -strategies in general – on (expanded) arenas are single-threaded, since $\mathbb{C}_{!A}$ has the shape of a conflict-free forest. In order to get a cartesian category, the last thing to check is that single-threaded strategies are stable under composition.

Single-threadedness and its stability under composition is independent from symmetry, so we state it and prove it below in greater generality.

Let $\sigma : S \rightarrow A^\perp \parallel B$ and $\tau : T \rightarrow B^\perp \parallel C$ be negative single-threaded prestrategies. Then, $\tau \odot \sigma$ is single-threaded.

Proof. The details are rather tedious; we postpone them to Appendix B.5. \square

We have finished constructing our basic category of Concurrent Hyland-Ong games. Let us call **Cho** the category having: as objects, negative arenas; and as morphisms from A to B , negative single-threaded \sim -strategies $\sigma : S \rightarrow !A^\perp \parallel !B$, up to weak isomorphism.

As for **Tcg** we will also sometimes write $\sigma : A \xrightarrow{\mathbf{Cho}} B$ to mean that σ is a morphism from A to B in **Cho** (or just $\sigma : A \dashrightarrow B$ when **Cho** is clear from the context). We get:

Proposition 4.10. *The category **Cho** has finite products.*

In particular, it follows as usual that \times is a bifunctor $\mathbf{Cho}^2 \rightarrow \mathbf{Cho}$, by setting $\sigma_1 \times \sigma_2 = \langle \sigma_1 \odot \varpi_{A_1}, \sigma_2 \odot \varpi_{A_2} \rangle$, for $\sigma_1 : A_1 \dashrightarrow B_1$ and $\sigma_2 : A_2 \dashrightarrow B_2$.

When constructing the cartesian closed structure, we will leverage the compact closed structure of **Tcg**. Therefore, it is useful to connect the cartesian structure of **Cho** with the monoidal structure of **Tcg**. For that, we note that there is an isomorphism of essps:

$$m_{A,B} : !(A \times B) \cong !A \otimes !B$$

Using Definition 3.36, it lifts to an iso $\overline{m}_{A,B}$ in **Tcg** between them. Consequently:

Lemma 4.11. *Let $\sigma_1 : A_1 \xrightarrow{\mathbf{Cho}} B_1$ and $\sigma_2 : A_2 \xrightarrow{\mathbf{Cho}} B_2$. Then,*

$$\sigma_1 \times \sigma_2 \approx \overline{m_{B_1, B_2}^{-1}} \circ (\sigma_1 \otimes \sigma_2) \circ \overline{m_{A_1, A_2}}$$

Proof. Write $\sigma_1 : \mathcal{S}_1 \rightarrow !A_1^\perp \parallel !B_1$ and $\sigma_2 : \mathcal{S}_2 \rightarrow !A_2^\perp \parallel !B_2$. By definition,

$$\sigma_1 \times \sigma_2 = \langle \sigma_1 \circ \varpi_{A_1}, \sigma_2 \circ \varpi_{A_2} \rangle$$

By Lemma 3.38, $\sigma_1 \circ \varpi_{A_1} \cong (i_{A_1^\perp} \parallel !B_1) \circ \sigma_1$ and $\sigma_2 \circ \varpi_{A_2} \cong (i_{A_2^\perp} \parallel !B_2) \circ \sigma_2$. These maps have disjoint codomain, so up to weak isomorphism their pairing is

$$\langle (i_{A_1^\perp} \parallel !B_1) \circ \sigma_1, (i_{A_2^\perp} \parallel !B_2) \circ \sigma_2 \rangle : \mathcal{S}_1 \parallel \mathcal{S}_2 \rightarrow !(A_1 \times A_2)^\perp \parallel !(B_1 \times B_2)$$

Likewise by Lemma 3.38, $\overline{m_{B_1, B_2}^{-1}} \circ (\sigma_1 \otimes \sigma_2) \circ \overline{m_{A_1, A_2}}$ has (up to isomorphism) esp $\mathcal{S}_1 \parallel \mathcal{S}_2$ and labeling function the obvious relabeling of $\sigma_1 \otimes \sigma_2$ by m_{A_1, A_2} and m_{B_1, B_2} . It is a simple verification that these two coincide. \square

4.2. Cartesian closure. We finish the construction of our cartesian closed category by describing the cartesian closure. We have constructed **Cho** as a subcategory of **Tcg** – which, as a compact closed category, is symmetric monoidal closed. We wish to leverage this closed structure of **Tcg** in order to transfer it to **Cho**.


4.2.1. Arrow arena. For two thin concurrent games \mathcal{A} and \mathcal{B} in **Tcg**, the corresponding exponential object (following the compact closed structure) is obtained as $\mathcal{A}^\perp \parallel \mathcal{B}$. In **Cho**, where objects are arenas, this hints at defining the exponential object of A, B as $A^\perp \parallel B$. Indeed, it is easy to check that $!(A^\perp \parallel B) \cong !A^\perp \parallel !B$, so this matches the closed structure of **Tcg**. However, objects in **Cho** are required to be *negative* arenas, and $A^\perp \parallel B$ is no longer negative. Therefore, we are brought to introduce a negative variant of $A^\perp \parallel B$, that would be an object of **Cho**. The natural choice, familiar from Hyland-Ong games, is to make events in A depend on minimal events of B . It would be incorrect to make events of A depend on *all* minimal events of B , so we will instead create as many copies of A as they are minimal events in B . Writing $\min(B)$ for the set of minimal events of B , we define:

Definition 4.12. Let A, B be two negative arenas. Their **arrow** is $A \Rightarrow B$, with the following components.

- *Events, and polarity.* Those of $(\parallel_{b \in \min(B)} A^\perp) \parallel B$.
- *Causality.* As follows:

$$\leq_{(\parallel_{b \in \min(B)} A^\perp) \parallel B} \cup \{((2, b), (1, (b, a))) \mid b \in \min(B) \ \& \ a \in A\}$$

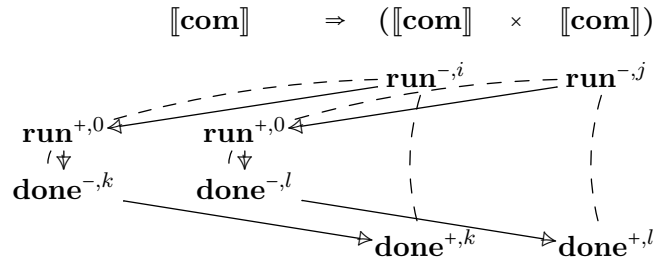
Example 4.13. The reader can check that $\llbracket \mathbf{com} \rrbracket \Rightarrow \llbracket \mathbf{com} \rrbracket$ is the arena presented as $\llbracket \mathbf{com} \rightarrow \mathbf{com} \rrbracket$ in Example 2.5. As $\llbracket \mathbf{com} \rrbracket$ has only one minimal event, there is no duplication of the left hand side. However, the arena $\llbracket \mathbf{com} \rrbracket \Rightarrow (\llbracket \mathbf{com} \rrbracket \times \llbracket \mathbf{com} \rrbracket)$ is displayed below.



This is exactly the arena construction of [HO00], where arenas are forests.

4.2.2. *Cartesian closed structure.* Our proof of cartesian closure will leverage the compact closed structure of **Tcg**. More precisely, we will show that there is a bijection (up to weak isomorphism) between negative single-threaded \sim -strategies playing respectively on $!A^\perp \parallel !(B \Rightarrow C)$ and $!A^\perp \parallel (!B^\perp \parallel !C)$. This bijection will leave the internal event structure of strategies unchanged, and will only operate through relabeling.

First, we describe the action of the bijection from $!A^\perp \parallel !(B \Rightarrow C)$ to $!A^\perp \parallel (!B^\perp \parallel !C)$. Let us first explain it on an example. Consider a \sim -strategy represented as below – which is, in essence, a curried version of the contraction on $\llbracket \mathbf{com} \rrbracket$ of Example 4.6.



Note that the positive moves on the left hand side have copy index 0, whereas in Example 4.6 they were carefully chosen so as to avoid collisions. This makes sense because the current arena has more causal links: the two positive moves on the left are already made distinct by their justification pointers, so there is no need to distinguish them further via their copy indices. As this example illustrates, we cannot simply relabel this \sim -strategy to $\llbracket \mathbf{com} \rrbracket^\perp \parallel (\llbracket \mathbf{com} \rrbracket \times \llbracket \mathbf{com} \rrbracket)$ without changing copy indices, as that would result in a collision, *i.e.* a failure of local injectivity of the labeling function.

Therefore, we use countability of the arena in order to do a collision-free relabeling.

Lemma 4.14. *There is a strong-receptive, courteous map of essps:*

$$\chi_{A,B} : !(A \Rightarrow B) \rightarrow !A^\perp \parallel !B$$

which, additionally, preserves the copy index of negative events.

Proof. For events $b \in B$ we use $\#b$ for the natural number associated to b by the countability of B . As in Section 2.5, we use $\langle - \rangle : \omega^3 \rightarrow \omega$ for any injective function; the collision with the pairing operation should not generate any confusion.

We set:

$$\begin{array}{lcl} \chi_{A,B} : & !(A \Rightarrow B) & \rightarrow !A^\perp \parallel !B \\ & (\alpha : [(1, (b, a))] \rightarrow \omega) & \mapsto (1, \alpha') \\ & (\beta : [(2, b)] \rightarrow \omega) & \mapsto (2, \beta') \end{array}$$

where:

$$\begin{array}{lcl} \alpha' : [a] & \rightarrow & \omega \\ a' & \mapsto & \langle \#b, \alpha((2, b)), \alpha((1, (b, a'))) \rangle \quad (\text{if } a' \in \min(A)) \\ a' & \mapsto & \alpha((1, (b, a'))) \quad (\text{otherwise}) \end{array}$$

and:

$$\begin{array}{lcl} \beta' : [b] & \rightarrow & \omega \\ b' & \mapsto & \beta((2, b')) \end{array}$$

This $\chi_{A,B}$ preserves symmetry, is strong-receptive (since minimal events of A^\perp are positive) and courteous (it only breaks immediate causal links from minimal events of B to minimal events of A^\perp , so from negative to positive). \square

This allows us, from $\sigma : \mathcal{S} \rightarrow !C^\perp \parallel !(A \Rightarrow B)$, to define its relabeling:

$$\begin{aligned} \Phi(\sigma) &: \mathcal{S} \rightarrow !C^\perp \parallel (!A^\perp \parallel !B) \\ &= (!C^\perp \parallel \chi_{A,B}) \circ \sigma \end{aligned}$$

For well-chosen hashing function \sharp and injection $\langle - \rangle$, this relabeling applied to the curried contraction above yields exactly the \sim -strategy of Example 4.6.

Before going on to the other direction, we note a further property of this relabeling.

Lemma 4.15. *Let $\sigma : \mathcal{S} \rightarrow !C^\perp \parallel !(A \Rightarrow B)$ be a negative single-threaded \sim -strategy. Take $s_1, s_2 \in \mathcal{S}$ such that σs_1 has the form $(2, \beta)$ with $\text{lbl } \beta = (2, b)$ ($b \in \min(B)$), and $\sigma s_2 = (2, \alpha)$ with $\text{lbl } \alpha = (1, (b', a))$. Then, $b = b'$ iff $s_1 = \text{init}(s_2)$.*

Proof. Straightforward consequence of single-threadedness. \square

Lemma 4.16. *Let $\sigma_1, \sigma_2 : \mathcal{S} \rightarrow !C^\perp \parallel !(A \Rightarrow B)$ be two negative single-threaded \sim -strategies sharing the same internal ess. Then, $\sigma_1 \sim_+ \sigma_2$ iff $\Phi(\sigma_1) \sim_+ \Phi(\sigma_2)$.*

Proof. if. Assume $\Phi(\sigma_1) \sim_+ \Phi(\sigma_2)$. Take $x \in \mathcal{C}(S)$, and form $\theta = \{(\sigma_1 s, \sigma_2 s) \mid s \in x\}$. We wish to prove that θ is a valid symmetry on $!C^\perp \parallel !(A \Rightarrow B)$. Firstly, we remark that the following diagram of bijections commutes.

$$\begin{array}{ccc} & x & \\ \sigma_1 \swarrow & & \searrow \sigma_2 \\ \sigma_1 x & \xrightarrow{\theta} & \sigma_2 x \\ \downarrow !C^\perp \parallel \chi_{A,B} & & \downarrow !C^\perp \parallel \chi_{A,B} \\ \Phi(\sigma_1) x & \xrightarrow{\theta'_{!C^\perp} \parallel (\theta'_{!A^\perp} \parallel \theta'_{!B})} & \Phi(\sigma_2) x \end{array}$$

It follows that θ decomposes as $\theta_{!C^\perp} \parallel \theta_{!(A \Rightarrow B)}$ with $\theta_{!C^\perp} \in \widetilde{!C^\perp}$, and we are left to prove that $\theta_{!(A \Rightarrow B)} \in \widetilde{!(A \Rightarrow B)}$. By construction it is a bijection, so we need to prove that it preserves and reflects causality, that it preserves labels, and that it preserves indices of negative events – which is clear, as they are preserved throughout this diagram.

We prove that it preserves immediate causality. The only nontrivial case concerns immediate causal links not preserved by $\chi_{A,B}$, *i.e.* those of the form:

$$\sigma_1 s_1 = (2, \{(2, b) \mapsto n\}) \rightarrow (2, \{(2, b) \mapsto n, (1, (b, a)) \mapsto p\}) = \sigma_1 s_2$$

But then, by Lemma 4.15, we have $s_1 = \text{init}(s_2)$. Since labels are preserved by $\theta'_{!A^\perp}$ and $\theta'_{!B}$, and using Lemma 4.15 again, we still have $\theta(\sigma_2 s_1) \rightarrow \theta(\sigma_2 s_2)$. The argument also applies to the θ^{-1} , which therefore is an order-isomorphism.

Preservation of labels also follows directly from Lemma 4.15. Finally, θ is a positive symmetry as all bijections involved preserve the copy index of negative events.

only if. By preservation of symmetry for $\chi_{A,B}$, and the fact that it preserves the copy index of negative events. \square

Relabeling from $!C^\perp \parallel (!A^\perp \parallel !B)$ to $!C^\perp \parallel !(A \Rightarrow B)$ is more subtle: we go from a game having one copy of A to one having as many as minimal moves in B . Thus, choosing the label for events formerly mapping to A requires identifying a copy of A corresponding

to some minimal event in B . Here condition (1) of single-threadedness is crucial: each s mapped to A has a unique minimal dependency $\text{init}(s)$ mapped to a minimal event of B , hence specifying the copy of A that s should be sent to. More formally, we prove:

Lemma 4.17. *For any single-threaded \sim -strategy $\sigma : \mathcal{S} \rightarrow !C \parallel (!A \parallel !B)$, there is $\sigma' : \mathcal{S} \rightarrow !C \parallel !(A \Rightarrow B)$, unique up to positive symmetry, such that $\sigma \sim_+ (!C \parallel \chi_{A,B}) \circ \sigma'$.*

Proof. We define $\sigma' : \mathcal{S} \rightarrow !C \parallel !(A \Rightarrow B)$. For $s \in \mathcal{S}$, if $\sigma(s) = (1, \gamma)$ we set $\sigma'(s) = (1, \gamma)$ still. If $\sigma(s) = (2, (2, \beta))$ with $\beta : [b] \rightarrow \omega$, then we set $\sigma'(s) = (2, \beta')$ with

$$\begin{array}{lcl} \beta' & : & [(2, b)] \rightarrow \omega \\ & & (2, b') \mapsto \beta(b') \end{array}$$

If $\sigma(s) = (2, (1, \alpha))$ with $\alpha : [a] \rightarrow \omega$, then by condition (1) of single-threadedness it has a unique minimal dependency $\text{init}(s) \leq s$. By hypothesis, $\sigma(\text{init}(s))$ has the form $(2, (2, \beta))$ with $\beta = \{b \mapsto n\}$. Therefore we set:

$$\begin{array}{lcl} \alpha' & : & [(1, (b, a))] \rightarrow \omega \\ & & (1, (b, a')) \mapsto \alpha(a') \\ & & (2, b) \mapsto n \end{array}$$

and we define $\sigma'(s) = (2, \alpha')$.

It is routine to check that this map is strong-receptive and courteous, and that its composition with $!C \parallel \chi_{A,B}$ is positively symmetric to σ . It follows from Lemma 4.16 that it preserves symmetry, and that it is unique up to positive symmetry. \square

From that, we deduce the following.

Proposition 4.18. *There is a bijection Φ up to weak isomorphism, preserving and reflecting weak isomorphism, between:*

- *Negative, single-threaded \sim -strategies $\sigma : \mathcal{S} \rightarrow !C^\perp \parallel !(A \Rightarrow B)$,*
- *Negative, single-threaded \sim -strategies $\sigma' : \mathcal{S} \rightarrow !C^\perp \parallel (!A^\perp \parallel !B)$.*

Moreover this bijection is compatible with pre-composition: for all $\tau : \mathcal{T} \rightarrow !D^\perp \parallel !C$,

$$\Phi(\sigma) \circ \tau \simeq \Phi(\sigma \circ \tau)$$

Proof. On the one hand $\Phi(\sigma)$ is obtained as $(!C^\perp \parallel \chi_{A,B}) \circ \sigma$, while $\Phi^{-1}(\sigma')$ is obtained by the unique factorisation of Lemma 4.17. The bijection up to weak isomorphism follows from Lemma 4.17 as well.

We now prove stability under composition. By definition, we have $\Phi(\sigma) = (!C^\perp \parallel \chi_{A,B}) \circ \sigma$. But by Lemma 3.38 this is the same (up to isomorphism) as $\overline{\chi_{A,B}} \circ \sigma$, so the action of Φ can be obtained by post-composition via a lifted map. Stability under composition follows immediately by associativity of composition. \square

And finally, we deduce:

Theorem 4.19. *The category **Cho** is cartesian closed.*

Proof. We already know that it is cartesian. Throughout this proof, in the construction of the components of the cartesian closed structure, we ignore the associativity and unity isomorphisms from the compact closed structure of **Tcg** – those can be easily and uniquely recovered from the context.

For any two arenas A, B , we first define the *evaluation* \sim -strategy:

$$\begin{aligned} \text{ev}_{A,B} &: A \times (A \Rightarrow B) \xrightarrow{\text{Cho}} B \\ &= (\epsilon_{!A} \otimes !B) \odot (!A \otimes \Phi(\mathfrak{c}_{!(A \Rightarrow B)})) \odot \overline{m_{A,A \Rightarrow B}} \end{aligned}$$

Likewise, for any $\sigma : A \times C \xrightarrow{\text{Cho}} B$, we define its *curryfication* as:

$$\begin{aligned} \Lambda(\sigma) &: C \xrightarrow{\text{Cho}} (A \Rightarrow B) \\ &= \Phi^{-1}(!A^\perp \otimes (\sigma \odot \overline{m_{A,C}^{-1}})) \odot (\eta_{!A} \otimes !C) \end{aligned}$$

It is then a straightforward equational reasoning to prove the two equations [LS88], for $\sigma : A \times C \xrightarrow{\text{Cho}} B$ and $\tau : C \xrightarrow{\text{Cho}} (A \Rightarrow B)$,

$$\begin{aligned} (\beta) \quad \text{ev}_{A,B} \odot (A \times \Lambda(\sigma)) &\approx \sigma \\ (\eta) \quad \Lambda(\text{ev}_{A,B} \odot (A \times \sigma)) &\approx \sigma \end{aligned}$$

using mainly Proposition 4.18 and the compact closed structure of **Tcg**, in combination with Lemma 4.11 to relate the cartesian structure of **Cho** and the monoidal structure of **Tcg** – all the structural isomorphisms involved in the definition cancel each other. \square

4.3. Recursion. As the final technical part of this paper, we prove that **Cho** supports the interpretation of a fixpoint combinator.

Usually in game semantics, the interpretation of the fixpoint combinator Y is obtained by showing that the category of games and strategies is enriched over a category of sufficiently complete partial orders. Here however it will not be the case: indeed, just as in AJM games [AJM00], our cartesian closed category is a quotient (its morphisms being weak isomorphism classes). It is not clear that the natural ordering on weak isomorphism classes is complete. However, this is not a real issue: although weak isomorphism classes of \sim -strategies might not form a complete partial order, concrete \sim -strategies do. Therefore, when solving recursive strategy equations, we will make sure to work with concrete \sim -strategies rather than weak isomorphism classes.

Our first step will be to order (concrete) \sim -strategies.

Definition 4.20. Let $\sigma : \mathcal{S} \rightarrow \mathcal{A}$, $\tau : \mathcal{T} \rightarrow \mathcal{A}$ be two \sim -strategies on a tcg \mathcal{A} . We write $\sigma \trianglelefteq \tau$ iff $\mathcal{S} \subseteq \mathcal{T}$, the inclusion map $\mathcal{S} \hookrightarrow \mathcal{T}$ is a map of essps, with all data in \mathcal{S} coinciding with the restriction of that in \mathcal{T} , and such that for all $s \in \mathcal{S}$, $\sigma s = \tau s$

The \sim -strategies on \mathcal{A} ordered by \trianglelefteq form a directed complete partial order (dcpo). It is not *pointed* though – it does not have a least element. Indeed, although a \trianglelefteq -minimal \sim -strategy only comprises (by receptivity) events matching minimal negative events of \mathcal{A} , their *name* in \mathcal{S} is arbitrary, so there is one \trianglelefteq -minimal \sim -strategy on \mathcal{A} for each renaming of the minimal negative events of \mathcal{A} . For each \mathcal{A} we distinguish one \trianglelefteq -minimal \sim -strategy

$$\perp_{\mathcal{A}} : \min^-(\mathcal{A}) \rightarrow \mathcal{A}$$

that has as events the negative minimal events of \mathcal{A} with induced symmetry, and as labeling function the identity. Not every \sim -strategy is above $\perp_{\mathcal{A}}$. However, for every \sim -strategy $\sigma : \mathcal{S} \rightarrow \mathcal{A}$, we pick one $\sigma \cong \sigma^\dagger$ such that $\perp_{\mathcal{A}} \trianglelefteq \sigma^\dagger$ obtained by renaming the minimal negative events of \mathcal{S} . We write $\mathcal{D}_{\mathcal{A}}$ for the pointed dcpo of \sim -strategies above $\perp_{\mathcal{A}}$.

Lemma 4.21. *For any tcg \mathcal{A} , $\mathcal{D}_{\mathcal{A}}$ is a pointed dcpo with $\perp_{\mathcal{A}}$ as minimal element.*

Proof. If $\Gamma = \{\gamma : \mathcal{S}_\gamma \rightarrow \mathcal{A}\} \subseteq \mathcal{D}_\mathcal{A}$ is a directed subset of $\mathcal{D}_\mathcal{A}$, we form

$$\vee \Gamma = \cup \gamma : \bigcup_{\gamma \in \Gamma} \mathcal{S}_\gamma \rightarrow \mathcal{A}$$

with all components defined as componentwise union.

It is direct that this defines a \sim -strategy, which is the least upper bound of Γ . \square

If all \sim -strategies in a directed set Γ are negative or single-threaded, so is $\vee \Gamma$. We now note that all the operations we defined on \sim -strategies in this section are continuous for \sqsubseteq .

Lemma 4.22. *Composition, tensor, pairing, curryfication and the $(-)^{\dagger}$ operation defined above are continuous for \sqsubseteq .*

Proof. Straightforward. \square

From the above, we deduce the following.

Corollary 4.23. *For any arena A there is a fixpoint combinator $\mathcal{Y}_A : (A \Rightarrow A) \xrightarrow{\mathbf{Cho}} A$, i.e. a single-threaded \sim -strategy such that:*

$$\mathcal{Y}_A \approx ev_{A,A} \odot (\mathcal{Y}_A, \mathfrak{C}_{!(A \Rightarrow A)})$$

Proof. First, we define the following operation, using the combinators on **Cho**.

$$\begin{aligned} F & : \mathcal{D}_{!(A \Rightarrow A)^+ \|!A} \rightarrow \mathcal{D}_{!(A \Rightarrow A)^+ \|!A} \\ & \sigma \mapsto (ev_{A,A} \odot (\sigma, \mathfrak{C}_{!(A \Rightarrow A)}))^{\dagger} \end{aligned}$$

By Lemma 4.22 it is continuous, and from the outermost dagger it has indeed value in $\mathcal{D}_{!(A \Rightarrow A)^+ \|!A}$. Thus, we can take its least fixpoint $\mathcal{Y}_A \in \mathcal{D}_{!(A \Rightarrow A)^+ \|!A}$. The weak isomorphism in the statement actually follows as an *equality*. \square

5. INTERPRETATION OF IPA

In this section, we illustrate our model by defining the interpretation of IPA, displaying the interpretation of programs of interest, and proving a few properties along the way.

We emphasise here that our purpose is *not* to prove full abstraction, nor to prove deep properties of the interpretation. We feel indeed that given the length of the paper, the specifics of such an endeavour are best left for later. Furthermore, it is our impression that it serves the purpose of this paper better (introducing and developing Concurrent Hyland-Ong games) to give the reader an understanding of what the model computes, what it can and cannot do, rather than delve into additional technical developments.

Throughout this section, by *strategy* we mean \sim -*strategy* (symmetries will be implicit).

5.1. Sequential innocent part. In this subsection, we focus on the interpretation of the (sequential) innocent part of IPA, *i.e.* essentially PCF, plus the combinators for commands. In other words, it lacks state and parallel composition of commands.

$$\begin{array}{cccc}
\llbracket \mathbf{tt} \rrbracket = !\llbracket \mathbb{B} \rrbracket & \llbracket \mathbf{ff} \rrbracket = !\llbracket \mathbb{B} \rrbracket & \llbracket n \rrbracket = !\llbracket \mathbb{N} \rrbracket & \llbracket \mathbf{skip} \rrbracket = !\llbracket \mathbf{com} \rrbracket \\
\begin{array}{c} \mathbf{q}^{-,i} \\ \downarrow \\ \mathbf{tt}^{+,0} \end{array} & \begin{array}{c} \mathbf{q}^{-,i} \\ \downarrow \\ \mathbf{ff}^{+,0} \end{array} & \begin{array}{c} \mathbf{q}^{-,i} \\ \downarrow \\ n^{+,0} \end{array} & \begin{array}{c} \mathbf{run}^{-,i} \\ \downarrow \\ \mathbf{done}^{+,0} \end{array}
\end{array}$$

FIGURE 8. Interpretation of constants of IPA

Interpretation of types. The arenas for the types **com** and \mathbb{B} were given in Example 2.5. The interpretation for \mathbb{N} is a countably infinite variant of the interpretation of \mathbb{B} :

$$\begin{array}{c}
\llbracket \mathbb{N} \rrbracket \\
\mathbf{q}^{-} \\
\downarrow \\
0^+ \quad 1^+ \quad \dots \quad n^+ \quad \dots
\end{array}$$

The interpretation extends to all types (not containing **ref**), with $\llbracket A \rightarrow B \rrbracket = \llbracket A \rrbracket \Rightarrow \llbracket B \rrbracket$.

Interpretation of terms. The interpretation follows the standard lines of the interpretation of the λ -calculus in a cartesian closed category. A *context* $\Gamma = x_1 : A_1, \dots, x_n : A_n$ is interpreted as the product $\prod_{1 \leq i \leq n} \llbracket A_i \rrbracket$ (which is just the parallel composition of the A_i s). A *typing sequent* $\Gamma \vdash M : A$ is interpreted as a **Cho**-morphism:

$$\llbracket \Gamma \vdash M : A \rrbracket : \llbracket \Gamma \rrbracket \xrightarrow{\mathbf{Cho}} \llbracket A \rrbracket$$

For the λ -calculus combinators – variables, application, abstraction λ , the interpretation is standard (and we do not detail it). For the fixpoint combinators, we use the combinator \mathcal{Y} of Section 4.3. The interpretation of constants is displayed in Figure 8. Note that we only display representations, treating multiple copies of Opponent moves symbolically. The reader should be able to expand them unambiguously to the full event structures, and to detail their isomorphism families. Note also that we give these interpretations over the empty context – they can easily be relabeled to any context Γ .

Likewise, the interpretation of function symbols is given in Figure 9. We have only one figure for a unary function $\text{op} : \mathbb{X} \rightarrow \mathbb{X}$, which covers (up to obvious relabeling) the cases of **succ** : $\mathbb{N} \rightarrow \mathbb{N}$, **pred** : $\mathbb{N} \rightarrow \mathbb{N}$ and **iszero** : $\mathbb{N} \rightarrow \mathbb{B}$. The interpretation of sequents involving those follows as usual, with *e.g.* the following composition in **Cho**:

$$\llbracket \Gamma \vdash \mathbf{if} M N_1 N_2 : \mathbb{X} \rrbracket = \llbracket \mathbf{if} \rrbracket \odot \langle \llbracket M \rrbracket, \llbracket N_1 \rrbracket, \llbracket N_2 \rrbracket \rangle : \llbracket \Gamma \rrbracket \xrightarrow{\mathbf{Cho}} \llbracket \mathbb{X} \rrbracket$$

At this point, we have defined the interpretation of the sequential innocent fragment of IPA. Using the cartesian closed structure and the definition of the fixpoint combinator, it would be straightforward to prove soundness and adequacy of the interpretation, *e.g.* using logical relations. We refrain from detailing this – rather standard – proof.

The paper already contains some examples of the interpretation of terms of the fragment of IPA currently under study, most notably in Section 2 – where for some, copy indices need to be adequately adjoined. The interpretation of such terms yields rather simple event structures, whose causal order is forest-shaped and without conflict. Modulo copy indices, and as it was noted in Section 2, these forests exactly coincide with the *view*




FIGURE 9. Interpretation of sequential function symbols of IPA

functions of standard Hyland-Ong games: their branches are exactly the P -views. Hence, our interpretation computes the composition of innocent strategies while staying within a causal representation corresponding to view functions, never resorting to expanded plays.

Non-determinism. Although the fragment of the language currently under study is deterministic, we find it interesting to study some examples given by its extension with a non-deterministic primitive. Therefore, we add to the language a new constant $\text{coin} : \mathbb{B}$ which returns a random boolean. Its interpretation is (an obvious extension with copy indices of) the strategy on the left hand side of Figure 5. For $\Gamma \vdash M, N : A$, we define as syntactic sugar a non-deterministic sum $\Gamma \vdash M + N : A$ as $\Gamma \vdash \text{if coin } M N : A$.

We give in Figure 10 representations of the interpretation of some well-chosen terms. Copy indices are not exactly as given by the interpretation function (though they are up to weak isomorphism): they have been relabeled for convenience of presentation.

As Figure 10a illustrates, the model represents non-determinism in a *non-idempotent way*: redundant non-deterministic choices are kept separate by the interpretation. In Figure 10c, \perp (which is syntactic sugar for $\mathcal{Y}(\lambda x^{\mathbb{B}}.x)$) is interpreted as the empty strategy. The interpretation of $\text{tt} + \perp$ illustrates that, despite displaying explicitly the point of non-deterministic branching, the hiding step of the interpretation removes some diverging branches of the interaction. Figures 10b and 10d display two strategies which have the same





FIGURE 10. Interpretation of some non-deterministic terms

branches (P-views), but differ in their branching points. This gives an interpretation of non-deterministic sequential programs that is similar to Tsukada and Ong’s recent presheaf-based model [TO15], although our composition mechanism is very different. It is fairly easy to capture exactly their category as a subcategory of **Cho**, whose morphisms are *sequential innocent* [CCW14, CCW15] but not deterministic.

5.2. Concurrent innocent part. Now, we go on to show how our model represents concurrent primitives. The only concurrent primitive of IPA is parallel composition, whose interpretation relies on the following strategy

$$[[\parallel]] = ![[\mathbf{com} \times \mathbf{com}]^\perp \parallel [\mathbf{com}]]$$



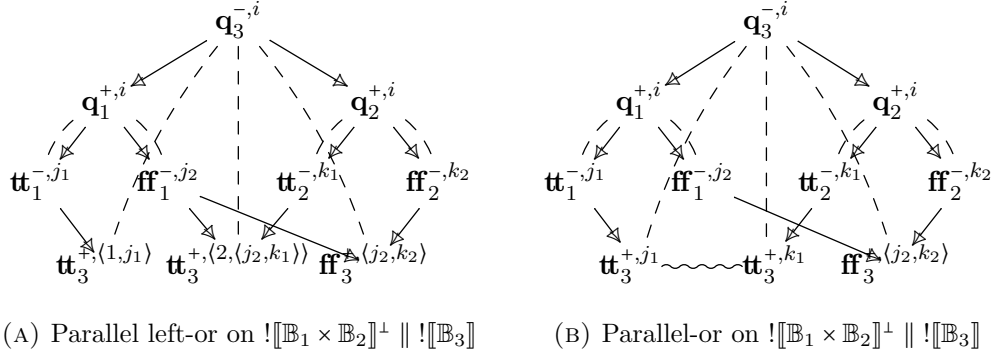


FIGURE 11. Two concurrent innocent strategies

Using this strategy we can define $[[\Gamma \vdash M \parallel N : \mathbf{com}]] = [[\parallel]] \odot \langle [[M]], [[N]] \rangle$.

This strategy is no longer a forest, but rather a directed acyclic graph. We also note that this is a *deterministic* strategy: there is no conflict in its event structure. As we shall see later, without any non-deterministic primitive, it is only in the presence of shared state that non-deterministic strategies will arise. In fact, a major advantage of our approach to modeling concurrent languages is that, not being based on interleavings, we represent the execution of such non-interfering terms deterministically.

In [CCW15], we exploit this property: we give a concurrent notion of innocence where strategies are directed acyclic graphs rather than forests, and using this notion we give an intensionally fully abstract model of a variant of PCF where independent computations are performed in parallel. The detailed construction is out of the scope of this particular paper, but let us illustrate it with two examples that are both concurrent innocent.

Figure 11 displays two concurrent innocent strategies (we associate moves to the corresponding sub-type using indices rather than location). In Figure 11a, we have a strategy for a parallel implementation of the left or, that is strict in its left argument. Indeed, although the strategy starts evaluating both its arguments in parallel, it can only return at toplevel if its first argument has returned. However, this is not true anymore for the strategy of Figure 11b. There, it suffices that *one* argument returns \mathbf{tt} for the overall computation to return \mathbf{tt} – indeed, this strategy computes the well-known *parallel-or* function [Plo77].

5.3. Stateful part. Finally, we finish the interpretation of IPA and describe how to interpret the primitives dealing with manipulations of state. For the simplicity of presentation, references only store booleans; however the method applies just as well to integers.

A variable can be interacted with in two ways: via reading and writing. As usual in game semantics, we follow this idea for the interpretation of variables, and take $[[\mathbf{ref}]]$ to be a product arena comprising actions for reading the reference or writing on the reference. More precisely, we define:

$$[[\mathbf{ref}]] = \begin{array}{c} \mathbf{R}^- \\ \swarrow \quad \searrow \\ \mathbf{tt}^+ \quad \mathbf{ff}^+ \end{array} \quad \begin{array}{cc} \mathbf{W}_{\mathbf{tt}}^- & \mathbf{W}_{\mathbf{ff}}^- \\ | & | \\ \mathbf{ok}^+ & \mathbf{ok}^+ \end{array}$$

We now describe the interpretation of term constructors for the manipulation of state. As usual, assignment and dereferencing are simply interpreted as (sequential innocent)




FIGURE 12. Strategies for assignment and dereferenciation

strategies that interact with the memory cell. We give in Figure 12 the strategies used in the interpretation of those. Using those, we can define:

$$\begin{aligned} \llbracket \Gamma \vdash M := N : \mathbf{com} \rrbracket &= \llbracket - := - \rrbracket \odot \langle \llbracket M \rrbracket, \llbracket N \rrbracket \rangle \\ \llbracket \Gamma \vdash !M : \mathbb{B} \rrbracket &= \llbracket !- \rrbracket \odot \llbracket M \rrbracket \end{aligned}$$

Before giving the interpretation of genuine references, we mention that the interpretation of **mkvar** exploits as usual the isomorphism between $\llbracket \mathbf{ref} \rrbracket$ and $\llbracket \mathbb{B} \rrbracket \times \llbracket \mathbf{com} \rrbracket^2$ [AM96].


New reference. As usual, the subtle part is the interpretation of **newref**. Indeed, whereas the strategies for assignment and dereferenciation only interact with the interface of the variable in an innocent way, it is **newref** that provides an implementation for the memory.

If $\Gamma, x : \mathbf{ref} \vdash M : A$ depends on a reference x , its interpretation plays on (up to iso) $\llbracket \mathbf{ref} \rrbracket^\perp \parallel \llbracket \Gamma \rrbracket^\perp \parallel \llbracket A \rrbracket$. Naively (we will see that this is a slight simplification), all we have to do is to build a strategy cell $\llbracket !\mathbf{ref} \rrbracket$, and compose $\llbracket M \rrbracket$ with it to obtain $\llbracket \mathbf{newref} \ x \ \mathbf{in} \ M \rrbracket$.

To define cell, we keep in mind the operational behaviour of a memory cell. In our (sequentially consistent) understanding of memory in a concurrent setting, although reads and writes are called concurrently, they are performed in some sequential order by the central memory. Thus the behaviour of a boolean memory cell is best described as the prefix language of the infinite traces:

$$\begin{aligned} \text{Cell}_{\mathbf{tt}} &::= \mathbf{R}^- \cdot \mathbf{tt}^+ \cdot \text{Cell}_{\mathbf{tt}} \mid \mathbf{W}_{\mathbf{tt}}^- \cdot \mathbf{ok}^+ \cdot \text{Cell}_{\mathbf{tt}} \mid \mathbf{W}_{\mathbf{ff}}^- \cdot \mathbf{ok}^+ \cdot \text{Cell}_{\mathbf{ff}} \\ \text{Cell}_{\mathbf{ff}} &::= \mathbf{R}^- \cdot \mathbf{ff}^+ \cdot \text{Cell}_{\mathbf{ff}} \mid \mathbf{W}_{\mathbf{tt}}^- \cdot \mathbf{ok}^+ \cdot \text{Cell}_{\mathbf{tt}} \mid \mathbf{W}_{\mathbf{ff}}^- \cdot \mathbf{ok}^+ \cdot \text{Cell}_{\mathbf{ff}} \end{aligned}$$

This language is ordered by prefix, so that $\text{Cell}_{\mathbf{ff}}$ is a forest. Setting all incomparable words to conflict with each other, we get an event structure whose events are words, and configurations are prefix-closed sets of prefixes of a word – so in one-to-one correspondence with words. This event structure, with the obvious labeling function, can be regarded as a prestrategy on $\llbracket \mathbf{ref} \rrbracket$ (not on $\llbracket !\mathbf{ref} \rrbracket$). But in order to fit in our framework, we need to equip it with copy indices (and symmetry). This calls for extra bookkeeping, as we need to make sure that the same copy index is not used twice in the same branch. We define


FIGURE 13. The pre- \sim -strategy newcell

$$\begin{aligned} \text{Cell}_{\mathbf{tt}}^{I_{\mathbf{R}}, I_{\mathbf{tt}}, I_{\mathbf{ff}}} &::= \mathbf{R}^{-,i} \cdot \mathbf{tt}^{+,0} \cdot \text{Cell}_{\mathbf{tt}}^{I_{\mathbf{R}} \cup \{i\}, I_{\mathbf{tt}}, I_{\mathbf{ff}}} & (i \notin I_{\mathbf{R}}) \\ &| \mathbf{W}_{\mathbf{tt}}^{-,i} \cdot \mathbf{ok}^{+,0} \cdot \text{Cell}_{\mathbf{tt}}^{I_{\mathbf{R}}, I_{\mathbf{tt}} \cup \{i\}, I_{\mathbf{ff}}} & (i \notin I_{\mathbf{tt}}) \\ &| \mathbf{W}_{\mathbf{ff}}^{-,i} \cdot \mathbf{ok}^{+,0} \cdot \text{Cell}_{\mathbf{ff}}^{I_{\mathbf{R}}, I_{\mathbf{tt}}, I_{\mathbf{ff}} \cup \{i\}} & (i \notin I_{\mathbf{ff}}) \end{aligned}$$

and similarly for $\text{Cell}_{\mathbf{ff}}^{I_{\mathbf{R}}, I_{\mathbf{tt}}, I_{\mathbf{ff}}}$. Then we define the event structure Cell via $\text{Cell}_{\mathbf{ff}}^{\emptyset, \emptyset, \emptyset}$, as we did above. It has an isomorphism family, that relates any two words differing only on their copy indices. Moreover the names of the events denote the labeling function to $!\llbracket \mathbf{ref} \rrbracket$ (with all positive moves pointing – that is, being immediately dependent in the game – to the previous move). Overall, we get a map of essp :

$$\text{cell} : \text{Cell} \rightarrow !\llbracket \mathbf{ref} \rrbracket$$

Example 5.1. The following diagram represents a sub-event structure of Cell .



We have constructed $\text{cell} : \text{Cell} \rightarrow !\llbracket \mathbf{ref} \rrbracket$ a map of essps . Unfortunately, cell is *not* a \sim -strategy: it is neither receptive (after playing $\mathbf{R}^{-,4}$ above one cannot play $\mathbf{R}^{-,4}$, although it is compatible in the game) nor courteous (we have $\mathbf{ff}^{+,0} \rightarrow \mathbf{R}^{-,4}$ which does not hold in $!\llbracket \mathbf{ref} \rrbracket$). However, cell is a *thin pre- \sim -strategy*, and as such can be composed with $\llbracket M \rrbracket : !\llbracket \mathbf{ref} \rrbracket^\perp \parallel !\llbracket A \rrbracket$ to obtain $\llbracket M \rrbracket \odot \text{cell}$ – and it turns out that $\llbracket M \rrbracket \odot \text{cell}$ is *always* a valid \sim -strategy.

Still, that is not quite what we want. The intended semantics for **newref** x in M is that each of its evaluations spawns a new, independent memory cell, whereas the operation above would have it spawned once and for all and shared over all copies of M . In other words, $\llbracket M \rrbracket \odot \text{cell}$ is a valid \sim -strategy indeed, but it might not be *single-threaded*. So finally, we build another pre- \sim -strategy displayed in Figure 13, where Cell means a copy of the pre- \sim -strategy above, with minimal events pointing as indicated.

Finally, from $\Gamma, x : \mathbf{ref} \vdash M : \mathbb{X}$, we define:

$$\llbracket \mathbf{newref} \ r \ \mathbf{in} \ M \rrbracket = \text{newcell} \odot \Lambda(\llbracket M \rrbracket) : !\llbracket \Gamma \rrbracket \xrightarrow{\mathbf{Tcg}} !\llbracket \mathbb{X} \rrbracket$$

Then, despite newcell being a pre- \sim -strategy rather than a \sim -strategy, we have:

Proposition 5.2. *For any $\Gamma, x : \mathbf{ref} \vdash M : \mathbb{X}$, the thin pre- \sim -strategy:*

$$\llbracket \mathbf{newref} \ r \ \mathbf{in} \ M \rrbracket = \mathbf{newcell} \odot \Lambda(\llbracket M \rrbracket)$$

is a single-threaded \sim -strategy.

Proof. The composition is well-defined (as a map of essps) since both compounds are \sim -receptive. Moreover, both compounds are also *componentwise courteous* (see Definition 3.33), so by Lemma 3.3.1 the composition $\mathbf{newcell} \odot \Lambda(\llbracket M \rrbracket)$ is a componentwise courteous pre- \sim -strategy. It is also thin, negative and single-threaded as these properties are stable under composition (respectively Lemmas 3.32, 4.3 and Proposition 4.1.2).

It remains to check that it is receptive and courteous. But that does not involve symmetry at all; and by the results of [RW11, CCRW17] it suffices to check that

$$\mathfrak{C}_{![\mathbb{X}]} \odot (\mathbf{newcell} \odot \Lambda(\llbracket M \rrbracket)) \odot \mathfrak{C}_{![\Gamma]} \cong \mathbf{newcell} \odot \Lambda(\llbracket M \rrbracket)$$

but that follows from the composition of componentwise courteous pre- \sim -strategies being associative, $\Lambda(\llbracket M \rrbracket)$ being a strategy, and the easy fact that $\mathfrak{C}_{![\mathbb{X}]} \odot \mathbf{newcell} \cong \mathbf{newcell}$. \square

This concludes the definition of the interpretation of IPA in **Cho**. As said before, we do not aim in this paper to prove properties of this interpretation, such as soundness or adequacy – those could be either proved directly as in [AM96], or more easily by constructing a functor to the interpretation of [AM96] linearizing the partial orders. In any case, the proof would take additional space without bringing much insight or taking advantage of the more refined representation offered by our event structures strategies, so we chose not to include it. However, we will now illustrate this interpretation by providing some examples.

5.4. Some examples.

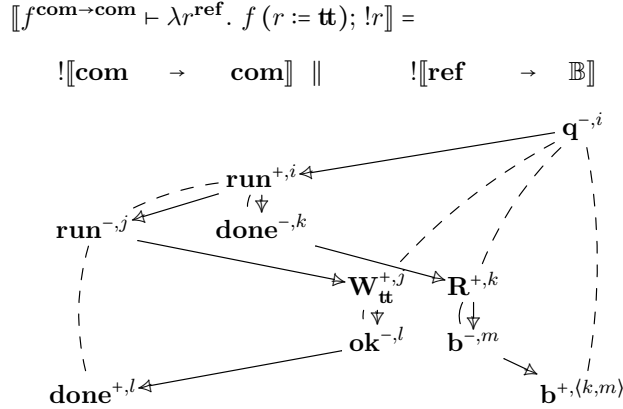
First example: strictness test. As a first example, we detail the interpretation of the term of Example 2.22. Recall that it was:

$$\mathbf{newref} \ b \ \mathbf{in} \ \lambda f^{\mathbf{com} \rightarrow \mathbf{com}}. f \ (b := \mathbf{tt}); !b : (\mathbf{com} \rightarrow \mathbf{com}) \rightarrow \mathbb{B}$$


As the constructor **new** is only defined on terms of ground type, this is just syntactic sugar for $\lambda f^{\mathbf{com} \rightarrow \mathbf{com}}. \mathbf{newref} \ b \ \in f \ (b := \mathbf{tt}); !b : (\mathbf{com} \rightarrow \mathbf{com}) \rightarrow \mathbb{B}$. In order to define its interpretation, the first step is to define:

$$\llbracket f : \mathbf{com} \rightarrow \mathbf{com}, b : \mathbf{ref} \vdash f \ (b := \mathbf{tt}); !b : \mathbb{B} \rrbracket : \llbracket \mathbf{com} \rightarrow \mathbf{com} \rrbracket \times \llbracket \mathbf{ref} \rrbracket \xrightarrow{\mathbf{Cho}} \llbracket \mathbb{B} \rrbracket$$

This is covered by the definitions above, using the cartesian closed structure and the strategies of Figure 12 for assignment and dereferenciation. Computing this yields the strategy represented below (again, the copy indices given by the actual interpretation function differ, but this is irrelevant up to weak isomorphism). Again, this deterministic event structure is forest-shaped and its branches are versions with explicit copy indices of the P-views of the corresponding innocent strategy in Hyland-Ong games.



Now, we compose it with `newcell`. We represent below the event structure resulting from their interaction. Events of the hidden/synchronised part of the interaction no longer have a well-defined polarity, hence we set it to 0. After hiding, the minimal conflict between the first two events in cell is inherited by the final positive events. The reader can check that hiding yields (up to the copy indices) the event structure of Example 2.22.



The reader familiar with Abramsky and McCusker's model for IA will see that taking the plays – *i.e.* alternating well-bracketed linear orderings of configurations, without copy indices – yields the expected sequential strategy. But our model says more, *e.g.* it specifies the behaviour of the strategy if Opponent *both* asks its argument *and* returns in parallel.

Second example: synchronization through state. We interpret the following term of IPA.

$$x : \text{com}, y : \text{com} \vdash \text{newref } r \text{ in}$$


$$\text{if } (!r) \perp (x; r := \text{tt}) \parallel$$

$$\text{if } (!r) y \perp$$


$$: \text{com}$$

This term simulates sequential composition through parallel composition and state up to may-equivalence: the only execution that survives divergence is the one where the first thread is executed before the second, so that y is run after x has terminated.

As before, we first compute the interpretation of the variant of this term where the variable has been abstracted away, obtaining the following strategy.




We now compute a part of the interaction with newcell, pictured below.



which, after hiding, yields a strategy with sub-event structures such as:

$$\llbracket \lambda x^{\text{com}}, y^{\text{com}}. \text{newref } r \text{ in } (\text{if } (!r) \perp (x; r := \mathbf{tt}) \parallel (\text{if } (!r) y \perp) \rrbracket =$$

$$\llbracket \mathbf{com} \rrbracket^\perp \parallel \llbracket \mathbf{com} \rrbracket^\perp \parallel \llbracket \mathbf{com} \rrbracket$$



Here, there are several observations to make.

Firstly, the copy index of the call to y does not depend on k . This might seem surprising: the diagram suggests that if Opponent plays two occurrences of $\mathbf{done}^{-,k}$, once with $k = 0$ and once with $k = 1$, Player will play the subsequent $\mathbf{run}^{+, (i,0)}$ twice, breaking local injectivity. In fact, recomputing the interaction with two occurrences of $\mathbf{done}^{-,k}$ one realizes that the two occurrences of $\mathbf{run}^{+, (i,0)}$ do exist, but they conflict with each other: this triggers new $\mathbf{W}_{\mathbf{tt}}^{0, (0,k)}$ events, which would be sequentialized in some order by the memory. The $\mathbf{R}^{0, 2i+1}$ would happen at some point during that sequentialization; each of these possible occurrences of $\mathbf{R}^{0, 2i+1}$ would lead to a call to y – so there would be multiple non-deterministic calls to y . This illustrates that our symbolic representation of strategies is incomplete, and does not specify in general their behaviour if Opponent replicates their moves (though it is crucial in [CCW15] that this representation becomes complete for *innocent* strategies).

Secondly, we note that this term, in the Ghica-Murawski model of IPA [GM08], would be interpreted by the same strategy than that for sequential composition. Unlike their model, we keep some information about non-deterministic branching; meaning that we do remember here that the term has a chance to diverge. In the interpretation presented in this paper, we do not remember *all* the information about divergences though. If one was to simplify the term above to $\text{newref } r \text{ in } \lambda x^{\text{com}}, y^{\text{com}}. x; r := \mathbf{tt} \parallel \text{if } !r \ y \ \perp$, the branch where the read arrives too early *w.r.t.* the write would be hidden away by composition. The sole purpose of the superfluous read in our example above is to create a race in memory before x , spawning two non-deterministic copies of the execution of x . In one of them the computation is doomed, as the second thread is stuck in a loop.

6. CONCLUSIONS

In this paper, we have given the detailed development leading to our cartesian closed category \mathbf{Cho} of *Concurrent Hyland-Ong* games, a setting that we illustrated with an interpretation of IPA. The cartesian closed category \mathbf{Cho} conservatively extends standard Hyland-Ong games, in the sense that in our setting purely functional programs are interpreted as (copy-index aware versions of) their tree of P-views – but our setting also supports stateful, non-deterministic, or concurrent languages, or any combination thereof.

The cornerstone of our construction is a compact closed category \mathbf{Tcg} of *thin concurrent games*, which extends Rideau and Winskel’s category CG of games and strategies as event structures [RW11, CCRW17]. Note the interest of \mathbf{Tcg} is not restricted to the construction of \mathbf{Cho} . It supports games that are much more general than those obtained from arenas. The future will tell how this mathematical space is best exploited, but in subsequent work we have already sometimes found it more convenient to build directly on \mathbf{Tcg} and on the AJM-style exponential $!_{\text{AJM}}$ rather than on \mathbf{Cho} .

Overall, we believe the framework is a very powerful setting for game semantics, whose ramifications will take some time to explore. Because it is conservative over traditional innocent game semantics but remembers the non-deterministic branching points, it natively supports a notion of non-deterministic innocence [CCW14, Cas17] (achieving this in traditional game semantics has long remained an open problem, only solved recently via reworking Hyland-Ong games using ideas from sheaf theory to remember the non-deterministic branching points [TO15]). For the same reason, it has been possible to extend the present framework to give models of non-deterministic languages adequate for any of *may*, *must* and *fair*-equivalence, whereas traditional game semantics are mostly confined to *angelic* non-determinism. The framework extends, also transparently, with quantitative information: in [CCPW18, CP18] it has been extended with probabilities, along with a notion of probabilistic innocence permitting a definability result and a collapse to the probabilistic relational model – such results were not within reach using the traditional toolbox of game semantics. A number of further extensions are under active development.

Beyond theoretical results, it is our hope that the truly concurrent nature of this model will prove useful as a basis for algorithmic analysis and verification of concurrent programs.

Acknowledgments. This work was partially supported by the LABEX MILYON (ANR-10-LABX-0070), and by the ERC Advanced Grant ECSYM.

REFERENCES

- [AGMO04] Samson Abramsky, Dan R. Ghica, Andrzej S. Murawski, and C.-H. Luke Ong. Applying game semantics to compositional software modeling and verification. In Kurt Jensen and Andreas Podelski, editors, *Tools and Algorithms for the Construction and Analysis of Systems, 10th International Conference, TACAS 2004, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2004, Barcelona, Spain, March 29 - April 2, 2004, Proceedings*, volume 2988 of *Lecture Notes in Computer Science*, pages 421–435. Springer, 2004.
- [AHM98] Samson Abramsky, Kohei Honda, and Guy McCusker. A fully abstract game semantics for general references. In *Thirteenth Annual IEEE Symposium on Logic in Computer Science, Indianapolis, Indiana, USA, June 21-24, 1998*, pages 334–344. IEEE Computer Society, 1998.
- [AJ09] Samson Abramsky and Radha Jagadeesan. Game semantics for access control. *Electr. Notes Theor. Comput. Sci.*, 249:135–156, 2009.
- [AJM00] Samson Abramsky, Radha Jagadeesan, and Pasquale Malacaria. Full abstraction for PCF. *Inf. Comput.*, 163(2):409–470, 2000.
- [AM96] Samson Abramsky and Guy McCusker. Linearity, sharing and state: a fully abstract game semantics for Idealized Algol with active expressions. *Electr. Notes Theor. Comput. Sci.*, 3:2–14, 1996.
- [AM99] Samson Abramsky and Paul-André Melliès. Concurrent games and full completeness. In *LICS*, pages 431–442. IEEE Computer Society, 1999.
- [BDER97] Patrick Baillot, Vincent Danos, Thomas Ehrhard, and Laurent Regnier. Believe it or not, AJM’s Games Model is a Model of Classical Linear Logic. In *Proceedings, 12th Annual IEEE Symposium on Logic in Computer Science, Warsaw, Poland, June 29 - July 2, 1997* [DBL97], pages 68–75.

- [BFHJ03] Albert Benveniste, Eric Fabre, Stefan Haar, and Claude Jard. Diagnosis of asynchronous discrete-event systems: a net unfolding approach. *IEEE Trans. Automat. Contr.*, 48(5):714–727, 2003.
- [BJ16] Daniil Berezun and Neil Jones. Partial evaluation and normalisation by traversals, 2016.
- [Bou09] Pierre Boudes. Thick subtrees, games and experiments. In *TLCA*, pages 65–79, 2009.
- [Cas17] Simon Castellan. *Concurrent structures in game semantics. (Structures concurrentes en sémantique des jeux)*. PhD thesis, University of Lyon, France, 2017.
- [CCPW18] Simon Castellan, Pierre Clairambault, Hugo Paquet, and Glynn Winskel. The concurrent game semantics of probabilistic PCF. In Anuj Dawar and Erich Grädel, editors, *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2018, Oxford, UK, July 09-12, 2018*, pages 215–224. ACM, 2018.
- [CCRW17] Simon Castellan, Pierre Clairambault, Silvain Rideau, and Glynn Winskel. Games and strategies as event structures. *Logical Methods in Computer Science*, 13(3), 2017.
- [CCW14] Simon Castellan, Pierre Clairambault, and Glynn Winskel. Symmetry in concurrent games. In Henzinger and Miller [HM14], page 28.
- [CCW15] Simon Castellan, Pierre Clairambault, and Glynn Winskel. The parallel intensionally fully abstract games model of PCF. In *30th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2015, Kyoto, Japan, July 6-10, 2015* [DBL15], pages 232–243.
- [CH10] Pierre Clairambault and Russ Harmer. Totality in arena games. *Ann. Pure Appl. Logic*, 161(5):673–689, 2010.
- [CHMO15] Conrad Cotton-Barratt, David Hopkins, Andrzej S. Murawski, and C.-H. Luke Ong. Fragments of ML decidable by nested data class memory automata. In Andrew M. Pitts, editor, *Foundations of Software Science and Computation Structures - 18th International Conference, FoSSaCS 2015, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2015, London, UK, April 11-18, 2015. Proceedings*, volume 9034 of *Lecture Notes in Computer Science*, pages 249–263. Springer, 2015.
- [CKV13] Ioana Cristescu, Jean Krivine, and Daniele Varacca. A compositional semantics for the reversible p -calculus. In *28th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2013, New Orleans, LA, USA, June 25-28, 2013*, pages 388–397. IEEE Computer Society, 2013.
- [CKV15] Ioana Domnina Cristescu, Jean Krivine, and Daniele Varacca. Rigid families for CCS and the π -calculus. In Martin Leucker, Camilo Rueda, and Frank D. Valencia, editors, *Theoretical Aspects of Computing - ICTAC 2015 - 12th International Colloquium Cali, Colombia, October 29-31, 2015, Proceedings*, volume 9399 of *Lecture Notes in Computer Science*, pages 223–240. Springer, 2015.
- [Cla15] Pierre Clairambault. Bounding linear head reduction and visible interaction through skeletons. *Logical Methods in Computer Science*, 11(2), 2015.
- [CP18] Pierre Clairambault and Hugo Paquet. Fully abstract models of the probabilistic lambda-calculus. In Dan R. Ghica and Achim Jung, editors, *27th EACSL Annual Conference on Computer Science Logic, CSL 2018, September 4-7, 2018, Birmingham, UK*, volume 119 of *LIPICs*, pages 16:1–16:17. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2018.
- [Cur98] Pierre-Louis Curien. Abstract Böhm trees. *Mathematical Structures in Computer Science*, 8(6):559–591, 1998.
- [DBL97] *Proceedings, 12th Annual IEEE Symposium on Logic in Computer Science, Warsaw, Poland, June 29 - July 2, 1997*. IEEE Computer Society, 1997.
- [DBL15] *30th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2015, Kyoto, Japan, July 6-10, 2015*. IEEE Computer Society, 2015.
- [EHS15] Clovis Eberhart, Tom Hirschowitz, and Thomas Seiller. An intensionally fully-abstract sheaf model for π . In Lawrence S. Moss and Pawel Sobocinski, editors, *6th Conference on Algebra and Coalgebra in Computer Science, CALCO 2015, June 24-26, 2015, Nijmegen, The Netherlands*, volume 35 of *LIPICs*, pages 86–100. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2015.
- [Ghi07] Dan R. Ghica. Geometry of synthesis: a structured approach to VLSI design. In Martin Hofmann and Matthias Felleisen, editors, *Proceedings of the 34th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2007, Nice, France, January 17-19, 2007*, pages 363–375. ACM, 2007.
- [GM03] Dan R. Ghica and Guy McCusker. The regular-language semantics of second-order idealized a_{gol} . *Theor. Comput. Sci.*, 309(1-3):469–502, 2003.


- [GM08] Dan R. Ghica and Andrzej S. Murawski. Angelic semantics of fine-grained concurrency. *Ann. Pure Appl. Logic*, 151(2-3):89–114, 2008.
- [God96] Patrice Godefroid. *Partial-Order Methods for the Verification of Concurrent Systems - An Approach to the State-Explosion Problem*, volume 1032 of *Lecture Notes in Computer Science*. Springer, 1996.
- [GT12] Dan R. Ghica and Nikos Tzevelekos. A system-level game semantics. *Electr. Notes Theor. Comput. Sci.*, 286:191–211, 2012.
- [Har99] Russ Harmer. *Games and full abstraction for nondeterministic languages*. PhD thesis, University of London, 1999.
- [HHM07] Russell Harmer, Martin Hyland, and Paul-André Mellies. Categorical combinatorics for innocent strategies. In *22nd IEEE Symposium on Logic in Computer Science (LICS 2007), 10-12 July 2007, Wroclaw, Poland, Proceedings*, pages 379–388. IEEE Computer Society, 2007.
- [HM14] Thomas A. Henzinger and Dale Miller, editors. *Joint Meeting of the Twenty-Third EACSL Annual Conference on Computer Science Logic (CSL) and the Twenty-Ninth Annual ACM/IEEE Symposium on Logic in Computer Science (LICS), CSL-LICS '14, Vienna, Austria, July 14 - 18, 2014*. ACM, 2014.
- [HMH14] Naohiko Hoshino, Koko Muroya, and Ichiro Hasuo. Memoryful geometry of interaction: from coalgebraic components to algebraic effects. In Henzinger and Miller [HM14], pages 52:1–52:10.
- [HO00] J. M. E. Hyland and C.-H. Luke Ong. On full abstraction for PCF: I, II, and III. *Inf. Comput.*, 163(2):285–408, 2000.
- [HP12] Tom Hirschowitz and Damien Pous. Innocent strategies as presheaves and interactive equivalences for CCS. *Sci. Ann. Comp. Sci.*, 22(1):147–199, 2012.
- [HW06] Jonathan Hayman and Glynn Winskel. Independence and concurrent separation logic. In *21th IEEE Symposium on Logic in Computer Science (LICS 2006), 12-15 August 2006, Seattle, WA, USA, Proceedings*, pages 147–156. IEEE Computer Society, 2006.
- [HW08] Jonathan Hayman and Glynn Winskel. The unfolding of general petri nets. In Ramesh Hariharan, Madhavan Mukund, and V. Vinay, editors, *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2008, December 9-11, 2008, Bangalore, India*, volume 2 of *LIPICs*, pages 223–234. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2008.
- [JNW93] André Joyal, Mogens Nielsen, and Glynn Winskel. Bisimulation and open maps. In *Proceedings of the Eighth Annual Symposium on Logic in Computer Science (LICS '93), Montreal, Canada, June 19-23, 1993*, pages 418–427. IEEE Computer Society, 1993.
- [Lai97] James Laird. Full abstraction for functional languages with control. In *LICS* [DBL97], pages 58–67.
- [LFVY15] Ugo Dal Lago, Claudia Faggian, Benoît Valiron, and Akira Yoshimizu. Parallelism and synchronization in an infinitary context (long version). *CoRR*, abs/1505.03635, 2015.
- [LL07] Søren B. Lassen and Paul Blain Levy. Typed normal form bisimulation. In Jacques Duparc and Thomas A. Henzinger, editors, *Computer Science Logic, 21st International Workshop, CSL 2007, 16th Annual Conference of the EACSL, Lausanne, Switzerland, September 11-15, 2007, Proceedings*, volume 4646 of *Lecture Notes in Computer Science*, pages 283–297. Springer, 2007.
- [LS88] Joachim Lambek and Philip J Scott. *Introduction to higher-order categorical logic*, volume 7. Cambridge University Press, 1988.
- [LS14] Paul Blain Levy and Sam Staton. Transition systems over games. In Henzinger and Miller [HM14], pages 64:1–64:10.
- [Mel03] Paul-André Mellies. Asynchronous games 1: A group-theoretic formulation of uniformity. *Manuscript, Available online*, 2003.
- [Mel05] Paul-André Mellies. Asynchronous games 4: A fully complete model of propositional linear logic. In *LICS*, pages 386–395. IEEE Computer Society, 2005.
- [MM07] Paul-André Mellies and Samuel Mimram. Asynchronous games: Innocence without alternation. In Luís Caires and Vasco Thudichum Vasconcelos, editors, *CONCUR*, volume 4703 of *LNCS*, pages 395–411. Springer, 2007.
- [MW08] Andrzej S. Murawski and Igor Walukiewicz. Third-order idealized algol with iteration is decidable. *Theor. Comput. Sci.*, 390(2-3):214–229, 2008.

- [Plo77] Gordon D. Plotkin. LCF considered as a programming language. *Theor. Comput. Sci.*, 5(3):223–255, 1977.
- [RW11] Silvain Rideau and Glynn Winskel. Concurrent strategies. In *LICS*, pages 409–418. IEEE Computer Society, 2011.
- [Sch14] Ulrich Schöpp. On the relation of interaction semantics to continuations and defunctionalization. *Logical Methods in Computer Science*, 10(4), 2014.
- [TO15] Takeshi Tsukada and C.-H. Luke Ong. Nondeterminism in game semantics via sheaves. In *30th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2015, Kyoto, Japan, July 6-10, 2015* [DBL15], pages 220–231.
- [VJA18] Matthijs Vákár, Radha Jagadeesan, and Samson Abramsky. Game semantics for dependent types. *Inf. Comput.*, 261(Part):401–431, 2018.
- [Win86] Glynn Winskel. Event structures. In Wilfried Brauer, Wolfgang Reisig, and Grzegorz Rozenberg, editors, *Advances in Petri Nets*, volume 255 of *LNCS*, pages 325–392. Springer, 1986.
- [Win07] Glynn Winskel. Event structures with symmetry. *Electr. Notes Theor. Comput. Sci.*, 172:611–652, 2007.
- [Win12] Glynn Winskel. Deterministic concurrent strategies. *Formal Asp. Comput.*, 24(4-6):647–660, 2012.
- [Win13] Glynn Winskel. Strategies as profunctors. In Frank Pfenning, editor, *FoSSaCS*, volume 7794 of *LNCS*, pages 418–433. Springer, 2013.

APPENDIX A. EXAMPLES AND COUNTER-EXAMPLES

A.1. Necessity for uniformity witnesses. In this first appendix, we give a few examples illustrating why uniformity is achieved by having strategies carry uniformity witnesses, rather than simply as a lifting property with respect to the symmetry in the game.

A.1.1. Necessity of witness. Our first example illustrates that simply requiring strategies on expanded arenas to satisfy a lifting property with respect to the reindexing isos in the game (as in AJM games) is unsound: it is too strict, and rejects some valid uniform strategies. In particular, it becomes required when one wishes to express uniformity for strategies with a non-deterministic branching behaviour. For instance, consider the following strategy.




where subscripts are there just to distinguish occurrences of moves in the discussion.

The strategy throws a non-deterministic coin. If it gets heads it calls its argument, then diverges. On the other hand if it gets tails it calls for its argument, then returns. This behaviour is easily definable in IPA, however a naive attempt at defining uniformity rejects it. Indeed, the two configurations $\{\mathbf{run}_3^{-,0}, \mathbf{run}_1^{+,0}, \mathbf{done}_1^{-,0}\}$ and $\{\mathbf{run}_3^{-,0}, \mathbf{run}_2^{+,0}, \mathbf{done}_2^{-,0}\}$ have the same image in the game, so in particular they coincide up to reindexing iso. However, the latter can be extended with $\{\mathbf{done}_3^{+,0}\}$ whereas the former cannot. This is because while the two events $\mathbf{run}_1^{+,0}$ and $\mathbf{run}_2^{+,0}$ have the same image in the game (and hence are symmetric), they correspond to entirely distinct events in the dynamic behaviour of the program and should not be considered symmetric by the strategy. Hence we need to equip strategies with *uniformity witnesses* that record which events are symmetric.

A.1.2. Non-uniqueness of witness. If a uniformity witness is needed, one may still hope to prove that it can be recovered in a unique or canonical way, so that being uniform can still be considered as a property rather than a piece of structure. Our next example shows that it is not the case: uniform strategies can very well differ only via their uniformity witnesses (this is best read after Section 3).

Consider the tcg \mathcal{A} with events $\{\ominus_0, \ominus_1, \oplus_0, \oplus_1\}$, all compatible, with trivial causality and symmetry consisting in all polarity-preserving bijections. Now consider the strategy (without symmetry) $\sigma : S \rightarrow A$ described as follows:



By strong receptivity, any isomorphism family \tilde{S} on S making σ a \sim -strategy must contain the two permutations id and π of the set $\{\ominus_0, \ominus_1\}$. We now consider how a valid

isomorphism family \tilde{S} may allow these two permutations to extend (in \tilde{S}) to permutations on the full set $\{\ominus_0, \ominus_1, \oplus_0, \oplus_1\}$. By thinness, the identity bijection can only extend by the identity. However, the thin axiom does not restrict the possible extensions of π . This means that π can extend *e.g.* by (\oplus_0, \oplus_0) or by (\oplus_0, \oplus_1) . In fact, there are exactly two isomorphism families \tilde{S}_1 and \tilde{S}_2 on S making σ a \sim -strategy, with maximal extension of π respectively given by π_1 (left), and π_2 (right):

$$\begin{array}{ccc} \ominus_0 & \mapsto & \ominus_1 \\ \ominus_1 & \mapsto & \ominus_0 \\ \oplus_0 & \mapsto & \oplus_1 \\ \oplus_1 & \mapsto & \oplus_0 \\ \pi_1 \in \tilde{S}_1 & & \end{array} \quad \begin{array}{ccc} \ominus_0 & \mapsto & \ominus_1 \\ \ominus_1 & \mapsto & \ominus_0 \\ \oplus_0 & \mapsto & \oplus_0 \\ \oplus_1 & \mapsto & \oplus_1 \\ \pi_2 \in \tilde{S}_2 & & \end{array}$$

This yields two \sim -strategies $\sigma_1 : \mathcal{S}_1 \rightarrow \mathcal{A}$ and $\sigma_2 : \mathcal{S}_2 \rightarrow \mathcal{A}$, which differ only by how they react to Opponent permuting their copy indices: either by also permuting (σ_1) or by doing nothing (σ_2). The two resulting \sim -strategies are not weakly isomorphic or weakly equivalent because there are no maps between \mathcal{S}_1 and \mathcal{S}_2 .


A.1.3. Operational interpretation. Here we attempt to give an operational reading on the difference between σ_1 and σ_2 above, by recasting the phenomenon in IPA (this is best read after Section 5). In the interpretation of IPA, these subtle differences between uniformity witnesses can convey indirect information reminiscent from the difference between causality due to the program order, and that due to the dependencies in memory.

Consider the two terms M_1, M_2 of type $(\mathbf{com} \rightarrow \mathbf{com}) \rightarrow \mathbf{com} \rightarrow \mathbf{com}$:

$$M_1 = \lambda f x. \mathbf{newref} \ r \ \mathbf{in} \ f(\mathbf{incr} \ r; \mathbf{wait} \ (!r = 2). x; \perp)$$

$$M_2 = \lambda f x. \mathbf{newref} \ r \ \mathbf{in} \ (f(\mathbf{incr} \ r; \perp) \parallel (\mathbf{wait} \ (!r = 2). (x \parallel x)))$$

where $\mathbf{incr} \ r$ is a short-hand for $r := !r + 1$ and $\mathbf{wait} \ b. M$ for $Y(\lambda l. \mathbf{if} \ b \ M \ l)$. For both terms, the strategy includes the following pattern, where the variable f calls its argument twice:



We omit copy indices to reduce clutter. As soon as f calls its argument *twice*, the two programs call x twice. The reader will recognize in the four lowest events the same pattern as in the previous example – and the isomorphism families for $\llbracket M_1 \rrbracket$ and $\llbracket M_2 \rrbracket$ restricted to this pattern indeed behave respectively like σ_1 and σ_2 .

It follows from the definition of the interpretation that this dynamics of an index exchange causing an index exchange occurs in the purely functional stage of the interpretation, as opposed to the accounting of state. As a consequence, the dynamics of indices exchange follows the causal links that originates from the program syntax tree, but ignores those that

come from communication through the memory. Hence in $\llbracket M_1 \rrbracket$ the Opponent exchange causes a Player exchange because the occurrence of x appears within the argument of f in the syntax tree. In contrast, in $\llbracket M_2 \rrbracket$ the Opponent exchange causes no Player exchange as the dependency only flows through the memory.

A.2. Absence of pullbacks in \mathcal{E}_\sim . The category of event structures with symmetry does not have pullbacks in general. For that we first note that if a diagram has a pullback in \mathcal{E}_\sim , then, forgetting symmetry, it is also a pullback in \mathcal{E} . The reason for that is the following proposition.


Proposition A.1. *The forgetful functor $\mathcal{E}_\sim \rightarrow \mathcal{E}$ which to any event structure with symmetry $A = (A, \widetilde{A})$ associates A , has a left adjoint.*

Proof. The right adjoint associates, to any event structure A , the event structure with symmetry (A, refl_A) , where

$$\text{refl}_A = \{ \{ (a, a) \mid a \in x \} \mid x \in \mathcal{C}(A) \}$$


is the minimal symmetry on A . It is straightforward that this defines an adjunction. \square

Hence the symmetry-forgetting functor is a right adjoint, and as such preserves pullbacks. Now, in order to prove that \mathcal{E}_\sim does not have pullbacks, we are going to construct a diagram in \mathcal{E}_\sim whose pullback in \mathcal{E} has no possible isomorphism family. Indeed, consider A the following event structure:



Write \mathcal{A} for A equipped with the maximal isomorphism family: all order-isomorphisms are in the family. Write \mathcal{A}_1 for the sub-event with symmetry where \cdot_1 can only be sent to itself and to \cdot_3 ; and \cdot_2 can only be sent to itself and to \cdot_4 . Similarly, write \mathcal{A}_2 for that where \cdot_1 can only be sent to \cdot_4 and \cdot_2 to \cdot_3 .

Now we have the following diagram:



Assume this diagram has a pullback $(\mathcal{A}_3, \Pi_1, \Pi_2)$. By Proposition A.1 its underlying event structure is A and the projection maps are both identities on objects. The isomorphism $\{(\cdot_a, \cdot_b)\} : \{\cdot_a\} \cong_{\widetilde{\mathcal{A}_3}} \{\cdot_b\}$ must be in $\widetilde{\mathcal{A}_3}$ as it is in both $\widetilde{\mathcal{A}_1}$ and $\widetilde{\mathcal{A}_2}$. However, its left hand side $\{\cdot_a\}$ can be extended with \cdot_1 , so by the extension property we must have


$$\{(\cdot_a, \cdot_b), (\cdot_1, \cdot_i)\} : \{\cdot_a, \cdot_1\} \cong_{\widetilde{\mathcal{A}_3}} \{\cdot_b, \cdot_i\}$$

with $i \in \{3, 4\}$. But by construction such an iso cannot be in both $\widetilde{\mathcal{A}_1}$ and $\widetilde{\mathcal{A}_2}$, absurd.

A.3. Weak equivalence is not a congruence. This key observation is one of the key facts guiding the design of our games.

Let \mathcal{A} , \mathcal{B} and \mathcal{C} be “games”, *i.e.* essps in the context of this discussion. Consider $\sigma_1 : \mathcal{S}_1 \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$, $\sigma_2 : \mathcal{S}_2 \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$, and $\tau : \mathcal{T} \rightarrow \mathcal{B}^\perp \parallel \mathcal{C}$ be “strategies”, *i.e.* maps of essps in the context of this discussion. Assume further than σ_1 and σ_2 are weakly equivalent, *i.e.* that there are $f : \mathcal{S}_1 \rightarrow \mathcal{S}_2$ and $g : \mathcal{S}_2 \rightarrow \mathcal{S}_1$ such that $g \circ f \sim \text{id}_{\mathcal{S}_1}$, $f \circ g \sim \text{id}_{\mathcal{S}_2}$, and the two obvious triangles commute up to symmetry. It first came as a surprise to us that then, $\tau \circ \sigma_1$ and $\tau \circ \sigma_2$ might *not* be weakly equivalent. Indeed, the extension property of isomorphism families ensures that two symmetric configurations have “bisimilar futures”. So it is natural to expect $\tau \circ \sigma_1$ and $\tau \circ \sigma_2$ to behave similarly, and indeed they do, but in a way less strict than that expressed by weak equivalence.

To be more precise, first consider the essp \mathcal{P} (the “pentagram”):



Its isomorphism family is the maximal one, *i.e.* all bijections between configurations are in the family. In \mathcal{P} , two events will eventually be played. It does not matter which ones, since they are all symmetric – the only thing that matters is the multiplicity.

We consider \mathcal{P} as a strategy on a game \mathcal{B} with the same events as \mathcal{P} ($\{\oplus_1, \oplus_2, \oplus_3, \oplus_4, \oplus_5\}$), the maximal isomorphism family, and no conflict. We write α_1 for the obvious labeling

$$\begin{aligned} \alpha_1 : \mathcal{P} &\rightarrow \mathcal{B} \\ \oplus_i &\mapsto \oplus_i \end{aligned}$$

which indeed informs a strategy on \mathcal{B} .

We will also be interested in another strategy:


$$\alpha_2 : \mathcal{S} \rightarrow \mathcal{B}$$

where \mathcal{S} has events $\{\oplus_1, \oplus_2\}$ and again maximal isomorphism family. The map α_2 sends \oplus_i in \mathcal{S} to \oplus_i in \mathcal{B} . The strategies α_1 and α_2 behave similarly, since both will eventually play two events; and we do not care which ones since all possible choices are symmetric in the game. Despite that, α_1 and α_2 are *not* weakly equivalent. In fact, *there is no map from α_1 to α_2* : such a map would require us to build a map of event structures from P to S , but the reader can check that this would induce a 2-coloring of P , which is not bipartite.

We will now obtain α_1 and α_2 respectively as compositions $\tau \circ \sigma_1$ and $\tau \circ \sigma_2$, for weakly equivalent σ_1 and σ_2 . We introduce the game

$$\mathcal{A} = \oplus_a \text{ ~~~~~ } \oplus_b$$

with again the maximal isomorphism family. The strategy $\tau : \mathcal{T} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$ selects α_1 or α_2 depending on Opponent's choice in \mathcal{A} . Its events are represented below.



Its isomorphism family is, again, the maximal one: all order-isomorphisms between configurations are valid symmetries. One can check that this satisfies indeed the axioms for an isomorphism family; crucially the extension axiom uses the fact that \mathcal{P} and \mathcal{S} are bisimilar (and that the symmetry on \mathcal{B} is the maximal one).

Finally, consider σ_1, σ_2 on \mathcal{A} , with σ_1 playing only Θ_a and σ_2 playing only Θ_b . They are clearly weakly equivalent, since $\{(\Theta_a, \Theta_b)\}$ is in \hat{A} . But by construction we have $\tau \circ \sigma_1 \cong \alpha_1$ and $\tau \circ \sigma_2 \cong \alpha_2$, which as we observed are not weakly equivalent.

Note that the games \mathcal{A} and \mathcal{B} are both tcgs; but crucially τ is not *thin* (Definition 3.13). Indeed, for instance, the symmetry $\{(\Theta_a, \Theta_b)\}$ extends to both $\{(\Theta_a, \Theta_b), (\Theta_1, \Theta_1)\}$ and $\{(\Theta_a, \Theta_b), (\Theta_1, \Theta_2)\}$, which is forbidden by Definition 3.13. For *thin* strategies, positive extensions of the symmetry must be canonically chosen, making it impossible that composite strategies as above are bisimilar but not weakly equivalent.

A.4. Failure of extension for copycat on general games. In the main text, we give (Definition 3.34) a candidate $\mathbb{C}_{\hat{\mathcal{A}}}$ for the isomorphism family on copycat for any essp \mathcal{A} . The valid symmetries on $\mathbb{C}_{\mathcal{A}}$ are simply those order-isomorphisms between configurations of $\mathbb{C}_{\mathcal{A}}$ which map to valid symmetries on $\mathcal{A}^\perp \parallel \mathcal{A}$.

We prove in Proposition 3.3.2 that this satisfies the extension property of isomorphism families if the game is a tcg. This boiled down to Lemma B.3, which shows that tcgs are *race-preserving* : races in the isomorphism family always originate to races in the game. As this phenomenon played an important role in the design of the theory, we find it useful to include here an example demonstrating the fact that without this race-preservation property (so if the games are plain event structures with polarities and symmetry, rather than tcgs), the extension property fails in general for the isomorphism family on copycat.

Consider an essp $A = \{a^-, b^+\}$. We form $!_2 A$ with events/causality/polarities/conflict those of $A \parallel A$ (we write $a^{-,i}$ for (i, a^-)), and isomorphism families the set of bijections between configurations included in the two maximal ones:

$$\begin{aligned} & \{(a^{-,1}, a^{-,1}), (b^{+,1}, b^{+,1}), (a^{-,2}, a^{-,2}), (b^{+,2}, b^{+,2})\} \\ & \{(a^{-,1}, a^{-,2}), (b^{+,1}, b^{+,2}), (a^{-,2}, a^{-,1}), (b^{+,2}, b^{+,1})\} \end{aligned}$$

So maximal symmetries either globally preserve the copy indices, or globally swap them. It is not possible for a symmetry to, *e.g.* send $a^{-,1}$ to $a^{-,1}$ and $b^{+,1}$ to $b^{+,2}$. It is, in fact, a binary version of the $!_{AJM}$ operation of Definition 3.4, applied to \mathcal{A} .

From Definition 3.34, $\mathbb{C}_{!_2 A}$ is equipped with a candidate isomorphism family $\mathbb{C}_{!_2 \hat{A}}$. We now show that this however fails the (*Extension*) axiom of isomorphism families. From the definition, the diagram below represents a valid symmetry in $\mathbb{C}_{!_2 \hat{A}}$.

$$\begin{array}{ccccccc}
 !_2 A^\perp & \parallel & !_2 A & \cong_{\mathbb{C}_{!_2 A}} & !_2 A^\perp & \parallel & !_2 A \\
 b^{-,1} & \cdots & a^{-,1} & \cdots & b^{-,1} & \cdots & a^{-,2}
 \end{array}$$

The issue will come from the fact that the symmetry follows irreconcilable courses in the left and the right components of $\mathbb{C}_{!_2}$: in the left component it preserves copy indices, whereas in the right component it swaps them. So it the left hand side of this symmetry extends as depicted below

$$\begin{array}{ccccccc}
 !_2 A^\perp & \parallel & !_2 A & \cong_{\mathbb{C}_{!_2 A}} & !_2 A^\perp & \parallel & !_2 A \\
 b^{-,1} & \cdots & a^{-,1} & \cdots & b^{-,1} & \cdots & a^{-,2} \\
 & \searrow & & & & & \\
 & & & & & & b^{+,1}
 \end{array}$$

the only matching extension on the right hand side is with $b^{+,1}$ as well ($b^{+,2}$ is not possible as it would require to play $b^{-,2}$ first), but $(b^{+,1}, b^{+,1})$ is not a valid extension of the symmetry above, as for that it would need to swap the copy index instead of preserving it.

APPENDIX B. POSTPONED PROOFS

B.1. Pullbacks of dual \sim -receptive maps. Let $\sigma : \mathcal{S} \rightarrow \mathcal{A}$ and $\tau : \mathcal{T} \rightarrow \mathcal{A}^\perp$ be \sim -receptive maps of ess. The set $\tilde{\mathcal{S}} \wedge \tilde{\mathcal{T}}$ is an isomorphism family on $S \wedge T$ and the ess $(S \wedge T, \tilde{\mathcal{S}} \wedge \tilde{\mathcal{T}})$ is a pullback in \mathcal{E}_\sim of σ and τ , written $\mathcal{S} \wedge \mathcal{T}$.

Proof. The (Groupoid) and (Restriction) axioms are direct consequences of the corresponding conditions for $\tilde{\mathcal{S}}$ and $\tilde{\mathcal{T}}$.

(Extension). Let $\theta : w \cong_{\tilde{\mathcal{S}} \wedge \tilde{\mathcal{T}}} z$. Assume w can be extended by an event $e \in S \wedge T$ to w' . Write $s = \Pi_1 e$ and $t = \Pi_2 e$, and assume e.g. σs is positive in \mathcal{A} . We then have:

$$\begin{array}{ccccc}
 \Pi_1 w' & \simeq & (\sigma \wedge \tau) w' & \simeq & \Pi_2 w' \\
 \downarrow s & & & & \uparrow t \\
 \Pi_1 w & \xrightarrow{\sigma} (\sigma \wedge \tau) w \xrightarrow{\tau} & \Pi_2 w & & \\
 \theta_S \Downarrow_{\tilde{\mathcal{S}}} & & \theta_T \Downarrow_{\tilde{\mathcal{T}}} & & \\
 \Pi_1 z & \xrightarrow{\sigma} (\sigma \wedge \tau) z \xrightarrow{\tau} & \Pi_2 z & &
 \end{array}$$

We first use the extension property on θ_S as $\Pi_1 w \xrightarrow{s} c$: θ_S extends by (s, s') . Since $\sigma \theta_S = \tau \theta_T$, this means that $\tau \theta_T$ extends by $(\sigma s, \sigma s')$ which is negative in \mathcal{A}^\perp . By \sim -receptivity of τ , θ_T extends by (t, t') with $\tau t = \sigma s$ and $\tau t' = \sigma s'$. The picture is now:

$$\begin{array}{ccccc}
 \Pi_1 w' & \simeq & (\sigma \wedge \tau) w' & \simeq & \Pi_2 w' \\
 \downarrow s & & & & \uparrow t \\
 \Pi_1 w & \xrightarrow{\sigma} (\sigma \wedge \tau) w \xrightarrow{\tau} & \Pi_2 w & & \\
 \Downarrow_{\tilde{\mathcal{S}}} & \theta_S \Downarrow_{\tilde{\mathcal{S}}} & \theta_T \Downarrow_{\tilde{\mathcal{T}}} & \Downarrow_{\tilde{\mathcal{T}}} & \\
 \Pi_1 z & \xrightarrow{\sigma} (\sigma \wedge \tau) z \xrightarrow{\tau} & \Pi_2 z & & \\
 \uparrow s' & & \downarrow t' & & \\
 z_1 & \simeq & \sigma z_1 = \tau z_2 & \simeq & z_2
 \end{array}$$

The obtained $\varphi : z_1 \simeq z_2$ is secured by construction, so as observed in Definition 2.11 its graph is ordered by \leq_φ compatible with \leq_S and \leq_T . Therefore restricting φ to the causal history of (s', t') yields $e' = [(s', t')]_\varphi$ a prime secured bijection, *i.e.* an event $e' \in S \wedge T$ such that $z \xrightarrow{e'} z'$. Finally, $\theta \cup \{(e, e')\} \in \tilde{S} \wedge \tilde{T}$ because $\theta_S \cup \{(s, s')\} \in \tilde{S}$ and $\theta_T \cup \{(t, t')\} \in \tilde{T}$.

If σs is negative, the dual reasoning uses extension on \tilde{T} and then \sim -receptivity of σ .

It is a pullback. Clearly the maps $\Pi_1 : S \wedge T \rightarrow S$ and $\Pi_2 : S \wedge T \rightarrow T$ preserve symmetry: they map θ to θ_S and θ_T respectively. We only need to check the universal property. Assume we have two morphisms of ess $\varphi : \mathcal{X} \rightarrow \mathcal{S}$ and $\psi : \mathcal{X} \rightarrow \mathcal{T}$ such that the square commutes:

$$\begin{array}{ccc}
 & \mathcal{X} & \\
 \varphi \swarrow & & \searrow \psi \\
 \mathcal{S} \wedge \mathcal{T} & & \mathcal{T} \\
 \sigma \swarrow & & \searrow \tau \\
 \tilde{\mathcal{S}} & & \mathcal{A}
 \end{array}$$

Because $S \wedge T$ is a pullback in \mathcal{E} there is a map of event structures $\langle \varphi, \psi \rangle : X \rightarrow S \wedge T$ making the two triangles commute, which is unique in \mathcal{E} . This uniqueness lifts to \mathcal{E}_\sim as the forgetful functor $\mathcal{E}_\sim \rightarrow \mathcal{E}$ is faithful. To conclude we need only to prove that $\langle \varphi, \psi \rangle$ preserves symmetry and is thus a morphism in \mathcal{E}_\sim . Let $\theta : x \cong_{\tilde{X}} y$. It is transported to a bijection $\langle \varphi, \psi \rangle \theta : \langle \varphi, \psi \rangle x \simeq \langle \varphi, \psi \rangle y$ such that $(\langle \varphi, \psi \rangle \theta)_S = \varphi \theta$ and $(\langle \varphi, \psi \rangle \theta)_T = \psi \theta$, thus $\langle \varphi, \psi \rangle \theta \in \tilde{S} \wedge \tilde{T}$ by definition. \square

B.2. Composition of \sim -receptivity and componentwise courtesy. To prove that, we first introduce the following more local characterisation for \sim -receptivity.

Lemma B.1. *Let \mathcal{A} be a tcg and $\sigma : \mathcal{S} \rightarrow \mathcal{A}$ be a map of ess. Then, σ is \sim -receptive iff for all $x \in \mathcal{C}(\mathcal{S})$ and $x \xrightarrow{s_1^-}$, for all $\text{id}_{\sigma x} \cup \{(\sigma s_1, a_2)\} \in \tilde{\mathcal{A}}$, there exists a unique s_2 such that $\sigma s_2 = a_2$, and we have $\text{id}_x \cup \{(s_1, s_2)\} \in \tilde{\mathcal{S}}$.*

Proof. only if. Particular case of the definition of \sim -receptivity.

if. Assume $\theta : x_1 \cong_{\tilde{\mathcal{S}}} x_2$, $x_1 \xrightarrow{s_1^-}$ and $\sigma x_2 \xrightarrow{a_2^-}$ such that $\sigma \theta \cup \{(\sigma s_1, a_2)\} \in \tilde{\mathcal{A}}$. By (Extension), there is s'_1 such that $\theta \xrightarrow{(s_1, s'_1)}$. Since σ is a map of ess, we must have $\sigma \theta \cup \{(\sigma s_1, \sigma s'_1)\} \in \tilde{\mathcal{A}}$ as well. By (Groupoid), it follows that $\text{id}_{\sigma x_2} \cup \{(\sigma s'_1, a_2)\} \in \tilde{\mathcal{A}}$. By hypothesis, we get a unique s_2 such that $\sigma s_2 = a_2$, satisfying $\text{id}_{x_2} \cup \{(s'_1, s_2)\} \in \tilde{\mathcal{S}}$. And finally, by (Groupoid) again, $\theta \cup \{(s_1, s_2)\} \in \tilde{\mathcal{S}}$. \square

Using that, we prove:

Let $\sigma : \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$ and $\tau : \mathcal{T} \rightarrow \mathcal{B}^\perp \parallel \mathcal{C}$ be pre- \sim -strategies, such that σ is (A, B) -courteous and τ is (B, C) -courteous. Then, $\tau \odot \sigma$ is \sim -receptive and (A, C) -courteous.

Proof. As a preliminary to the proof, we note that thanks to (A, B) -courtesy of σ and (B, C) -courtesy of τ , the immediate dependencies of negative $p \in T \odot S$ in $T \otimes S$ have to be visible as well, and must map to the same component – this is the key argument of the proof. Indeed assume $p' \rightarrow p$ in $T \otimes S$ with visible p mapping to a negative event in $A^\perp \parallel C$,

for instance in C . Then by general properties of the pullback in \mathcal{E} , since p maps to C , its immediate causal dependencies $p' \rightarrow p$ in $T \otimes S$ must be such that $\Pi_2 p' \rightarrow \Pi_2 p$ in $A \parallel T$ (consequence of *e.g.* Lemma 2.7 of [CCRW17] with Proposition 2.16), but since p maps to C those must actually both be in T , and since τ is (B, C) -courteous p' must map to C as well, therefore it is visible.

From that, it is clear that $\tau \circ \sigma$ is (A, C) -courteous. We now show that it is \sim -receptive. We prove it via Lemma B.1. Take $z \in \mathcal{C}(T \circ S)$, assume z extends via some negative p , say in C . The configuration z has a witness $[z] \in \mathcal{C}(T \otimes S)$, however in general this witness might not extend with p , as it may need to perform some invisible events prior to that. In our case though, the preliminary above shows that this is not possible: the immediate dependencies in $T \otimes S$ of p are visible as well, and hence in $z \subseteq [z]$. Now, if we also have that $(\tau \circ \sigma)z$ extends with c^- with $\text{id}_{(\tau \circ \sigma)z} \cup \{((\tau \circ \sigma)p, c)\} \in \tilde{A} \parallel \tilde{B} \parallel \tilde{C}$, then

$$\Pi_2[z] \xrightarrow[\mathcal{C}]{\Pi_2 p} \text{id}_{\Pi_2[z]} \cup \{((A \parallel \tau)(\Pi_2 p), c)\} \in \tilde{A} \parallel \tilde{B} \parallel \tilde{C}$$

so using \sim -receptivity of $A \parallel \tau$, we can uniquely lift c to $A \parallel T$, hence to $T \otimes S$ and $T \circ S$, and that lifting is by construction compatible with $\tilde{T} \otimes \tilde{S}$ and $\tilde{T} \circ \tilde{S}$. \square

B.3. Proofs for the copycat \sim -strategy. This section contains proofs relative to the construction of the isomorphism family for the copycat strategy. We start with a simple characterisation of the valid symmetries announced in the main text.

The set $\mathbb{C}_{\tilde{A}}$ is equivalently defined as comprising the bijections

$$\theta_1 \parallel \theta_2 : x_1 \parallel x_2 \simeq_{\tilde{A}^\perp \parallel \tilde{A}} y_1 \parallel y_2$$

satisfying the further condition that for all $a \in x_1 \cap x_2$, we have $\theta_1(a) = \theta_2(a)$.

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

If θ is an order-iso, then take $a \in x_1 \cap x_2$. Assume without loss of generality that $\text{pol}_A(a) = +$, so that $(1, a) \rightarrow (2, a)$ in \mathbb{C}_A . But then since θ is an order-iso, it preserves immediate causal dependency, therefore $(1, \theta_1 a) \rightarrow (2, \theta_2 a)$. But since these two events are in different components of $A^\perp \parallel A$, this necessarily means that $\theta_1 a = \theta_2 a$ as required (using *e.g.* the characterisation of immediate causality of copycat in Lemma 3.3 of [CCRW17]).

Reciprocally, assume that for all $a \in x_1 \cap x_2$, $\theta_1 a = \theta_2 a$. Using again Lemma 3.3 of [CCRW17], it is immediate that θ preserves immediate causal links. The same reasoning applies to θ^{-1} (it is easy to show that the hypothesis is stable under inverse), so it reflects immediate causal links as well; and is an order-iso. \square

We now set to prove Proposition 3.3.2. Most verifications are direct; the main issue being to show that $\mathbb{C}_{\tilde{A}}$ always satisfies the axioms of isomorphism families. The first two axioms are immediate consequences of the definition of $\mathbb{C}_{\tilde{A}}$ in Definition 3.34:

Lemma B.2. *For any teg \mathcal{A} , the family $\mathbb{C}_{\tilde{A}}$ satisfies the axioms (Groupoid) and (Restriction) of isomorphism families.*

The main difficulty is to show the (*Extension*) axiom. And for good reasons: indeed, this axiom fails if \mathcal{A} is an essp with no further constraints (as illustrated in Appendix A.4). That it holds when \mathcal{A} is a teg boils down to the property below.

Lemma B.3. *Let \mathcal{A} be a tcg. Then \tilde{A} is race-preserving, in the sense that for any $\theta : x \cong_{\tilde{A}} y$, for any $\theta \sqsubseteq^+ \theta_1 : x_1 \cong_{\tilde{A}} y_1$ and $\theta \sqsubseteq^- \theta_2 : x_2 \cong_{\tilde{A}} y_2$, if x_1 and x_2 are compatible ($x_1 \cup x_2 \in \mathcal{C}(A)$), then so are θ_1 and θ_2 : $\theta_1 \cup \theta_2 \in \tilde{A}$ as well.*

Proof. We first prove that \tilde{A}_+ and \tilde{A}_- are race-preserving. Let $\theta : x \cong_{\tilde{A}_+} y$ with a positive extension $\theta_1 : x_1 \cong_{\tilde{A}_+} y_1$ and a negative extension $\theta_2 : x_2 \cong_{\tilde{A}_+} y_2$, with $x_1 \cup x_2 \in \mathcal{C}(A)$.

Using (Extension) of \tilde{A}_+ twice to θ_1 and θ_2 , we get to the following picture:

$$\begin{array}{ccc}
 \theta'_1 : x_1 \cup x_2 \cong_{\tilde{A}_+} y'_1 & & \theta'_2 : x_1 \cup x_2 \cong_{\tilde{A}_+} y'_2 \\
 \downarrow \cup & & \downarrow \cup \\
 \theta_1 : x_1 \cong_{\tilde{A}_+} y_1 & & \theta_2 : x_2 \cong_{\tilde{A}_+} y_2 \\
 \times \searrow & & \swarrow \cup \\
 & \theta : x \cong_{\tilde{A}_+} y &
 \end{array}$$

By the (Groupoid) axiom on \tilde{A}_+ , we have $\text{id}_y \subseteq \theta'_1 \circ \theta'^{-1}_2 : y'_2 \cong_{\tilde{A}_+} y'_1$. By (Restriction), we build $\varphi = \theta'_1 \circ \theta'^{-1}_2 \upharpoonright y_2$. By construction, we have $\text{id}_y \sqsubseteq^- \varphi \in \tilde{A}_+$, so $\varphi = \text{id}_{y_2}$ (as \mathcal{A}_+^\perp is thin). It follows that $\theta_2 \subseteq \theta'_1$, hence $\theta'_1 = \theta_1 \cup \theta_2$ as required. A dual reasoning shows that \tilde{A}_- is race-preserving as well.

Now, we deduce the result for \tilde{A} , using the decomposition of Lemma 3.18. Assume $\theta = \theta^- \circ \theta^+$ has extensions $\theta \sqsubseteq^+ \theta_1$ and $\theta \sqsubseteq^- \theta_2$, with decompositions $\theta_1^- \circ \theta_1^+$ and $\theta_2^- \circ \theta_2^+$. By monotonicity of the decomposition, we have $\theta^+ \sqsubseteq^+ \theta_1^+$, $\theta^+ \sqsubseteq^- \theta_2^+$, $\theta^- \sqsubseteq^+ \theta_1^-$ and $\theta^- \sqsubseteq^- \theta_2^-$. By race-preservation of \tilde{A}_+ it follows first that $\theta_1^+ \cup \theta_2^+ \in \tilde{A}_+$, and then by race-preservation of \tilde{A}_- it follows that $\theta_1^- \cup \theta_2^- \in \tilde{A}_-$. Thus $(\theta_1^- \cup \theta_2^-) \circ (\theta_1^+ \cup \theta_2^+) = (\theta_1^- \circ \theta_1^+) \cup (\theta_2^- \circ \theta_2^+) = \theta_1 \cup \theta_2 \in \tilde{A}$. \square

That \tilde{A} is race-preserving is actually a sufficient condition for the (*Extension*) axiom to hold on $\mathbb{C}_{\tilde{A}}$. With this we can finally complete the proof.

Let \mathcal{A} be a tcg. Then, writing $\mathbb{C}_{\mathcal{A}} = (\mathbb{C}_A, \mathbb{C}_{\tilde{A}})$, the map

$$\mathbb{c}_{\mathcal{A}} : \mathbb{C}_{\mathcal{A}} \rightarrow \mathcal{A}^\perp \parallel \mathcal{A}$$

is a \sim -strategy.

Proof. By Lemma B.2 it remains to prove (*Extension*). Let $\theta_1 \parallel \theta_2 : x \parallel y \cong_{\mathbb{C}_{\tilde{A}}} x' \parallel y'$.

Assume e.g. $x \parallel y \xrightarrow{(2,a)} \text{c}$. There are two cases:

- If $\text{pol}_A(a) = -$, then by (Extension) for $\tilde{A}^\perp \parallel \tilde{A}$ we have $\theta_1 \parallel \theta_2 \subseteq \theta_1 \parallel \theta'_2 \in \tilde{A}^\perp \parallel \tilde{A}$ whose domain is $x \parallel y \cup \{a\}$. Its codomain is $x' \parallel y' \cup \{a'\}$. Since $\text{pol}_A(a) = -$, we cannot have $a' \in x'$ – indeed $x' \supseteq^+ x' \cap y' \sqsubseteq^- y'$, so we would have $a' \in y'$ as well, absurd. So we have $x' \cap (y' \cup \{a'\}) = x' \cap y' \sqsubseteq^+ x'$, and $x' \cap (y' \cup \{a'\}) = x' \cap y' \sqsubseteq^- y' \sqsubseteq^- y' \cup \{a'\}$, which establishes that $x' \parallel (y' \cup \{a'\}) \in \mathcal{C}(\mathbb{C}_A)$.

Likewise we have $\theta_1 \cap \theta'_2 = \theta_1 \cap \theta_2$, hence we still have $\theta_1 \cap \theta'_2 \sqsubseteq^+ \theta_1$ but also $\theta_1 \cap \theta'_2 \sqsubseteq^- \theta_2 \sqsubseteq^- \theta'_2$, therefore $\theta_1 \parallel \theta'_2 \in \mathbb{C}_{\tilde{A}}$.

- If $\text{pol}_A(a) = +$ is positive then $a \in x$ as well. Thus, $[a] \subseteq x \cap y$. Therefore, we have $(x \cap y) \cup \{a\} \in \mathcal{C}(A)$, and $(x \cap y) \cup \{a\} \subseteq x$. Define $\theta'_1 = \theta_1 \upharpoonright (x \cap y) \cup \{a\}$. We have:

$$\theta'_1 \supseteq^+ \theta_1 \cap \theta_2 \sqsubseteq^- \theta_2$$

By construction, the domains of θ'_1 (which is $(x \cap y) \cup \{a\}$) and the domain of θ_2 (which is y) are compatible, so by Lemma B.3, $\theta'_2 = \theta'_1 \cup \theta_2 \in \tilde{A}$, and by construction its domain is

$y \cup \{a\}$. To sum up, we have:

$$\theta_1 \supseteq^+ \theta_1 \cap \theta_2' \subseteq^- \theta_2'$$

Hence $\theta_1 \parallel \theta_2' \in \mathbb{C}_{\tilde{\mathcal{A}}}$ provides the required extension.

We have established that $\mathbb{C}_{\tilde{\mathcal{A}}}$ is an isomorphism family. It is obvious that $\mathbb{C}_{\mathcal{A}} : \mathbb{C}_{\mathcal{A}} \rightarrow \mathcal{A}^\perp \parallel \mathcal{A}$ preserves symmetry. It remains to show that it is \sim -receptive, for which we apply Lemma B.1. Assume $x \parallel y \in \mathcal{C}(\mathbb{C}_{\mathcal{A}})$ can be extended by $(2, a^-)$ in $\mathbb{C}_{\mathcal{A}}$ and by $(2, b^-)$ in $\mathcal{A}^\perp \parallel \mathcal{A}$ (in which case it is immediate that it is a valid extension in $\mathbb{C}_{\mathcal{A}}$ as well), such that:

$$\text{id}_x \parallel (\text{id}_y \cup \{(a, b)\}) \in \tilde{\mathcal{A}}^\perp \parallel \tilde{\mathcal{A}}$$

We need to check that this is a valid extension in $\mathbb{C}_{\tilde{\mathcal{A}}}$ as well. By the characterisation of Proposition 3.3.2, we only have to check that $\text{id}_x c = (\text{id}_y \cup \{(a, b)\}) c$ for each $c \in x \cap (y \cup \{a\})$, but in fact we must have $c \in x \cap y$. Indeed, we cannot have $a \in x$, as by $x \supseteq^+ x \cap y \subseteq^- y$ and $\text{pol}_{\mathcal{A}}(a) = -$ that would imply $a \in y$ as well, absurd. So the verification is obvious.

Finally, that copycat is thin is an immediate consequence of \mathcal{A}_- and \mathcal{A}_+^\perp being thin along with the characterisation of symmetries in copycat of Proposition 3.3.2. \square

B.4. Positivisation of mediating maps. We start with the following lemma, which intuitively allows us to canonically “transport” a configuration along a negative symmetry.

Lemma B.4. *Let $\sigma : \mathcal{S} \rightarrow \mathcal{A}$ be a pre- \sim -strategy, $x \in \mathcal{C}(S)$, along with $\theta_- : \sigma x \cong_{\tilde{\mathcal{A}}_-} y$.*

Then, there is a unique $\varphi : x \cong_{\tilde{\mathcal{S}}} x'$ s.t. $\sigma \varphi = \theta_+ \circ \theta_- : \sigma x \cong_{\tilde{\mathcal{A}}} \sigma x'$ for some $\theta_+ : y \cong_{\tilde{\mathcal{A}}_+} \sigma x'$.

Proof. Uniqueness. If there are two such $\varphi_1 : x \cong_{\tilde{\mathcal{S}}} x'_1$ and $\varphi_2 : x \cong_{\tilde{\mathcal{S}}} x'_2$, then from the hypotheses $\psi = \varphi_2 \circ \varphi_1^{-1} : x'_1 \cong_{\tilde{\mathcal{S}}} x'_2$ such that $\sigma \psi \in \tilde{\mathcal{A}}_+$; by Lemma 3.27 it follows that ψ is an identity and $x'_1 = x'_2$, $\varphi_1 = \varphi_2$.

Existence. Direct by induction of x and θ_- . For negative extensions, it follows from the (*Extension*) property for $\tilde{\mathcal{A}}_+$ and \sim -receptivity of σ . For positive extensions, it follows from the (*Extension*) axiom for $\tilde{\mathcal{S}}$ and axiom (d) of tcgs on \mathcal{A} . \square

Let $\sigma : \mathcal{S} \rightarrow \mathcal{A}$ and $\sigma' : \mathcal{S}' \rightarrow \mathcal{A}$ be pre- \sim -strategies, and $f : \mathcal{S} \rightarrow \mathcal{S}'$ such that $\sigma' \circ f \sim_{\tilde{\mathcal{A}}} \sigma$. Then, there exists a unique $f' : \mathcal{S} \rightarrow \mathcal{S}'$ such that $f \sim_{\tilde{\mathcal{S}}} f'$, and $\sigma' \circ f' \sim_{\tilde{\mathcal{A}}_+} f$.

Proof. We show that for all $x \in \mathcal{C}(S)$, there is a unique $\varphi_x : f x \cong_{\tilde{\mathcal{S}}} x'$ such that $\sigma x \cong_{\tilde{\mathcal{A}}_+} \sigma' x'$.

Uniqueness is again a consequence of Lemma 3.27.

Existence. If $x \in \mathcal{C}(S)$, then by hypothesis we know that the triangle induces

$$\theta_x : \sigma'(f x) \cong_{\tilde{\mathcal{A}}} \sigma x$$

By Lemma 3.18, θ_x decomposes as $\sigma'(f x) \cong_{\tilde{\mathcal{A}}_-}^{\theta_x^-} y \cong_{\tilde{\mathcal{A}}_+}^{\theta_x^+} \sigma x$. By Lemma B.4, there is $\varphi_x : f x \cong_{\tilde{\mathcal{S}}} x'$ such that $\sigma \varphi = \theta_x'^+ \circ \theta_x^-$. We then have $\theta_x'^+ \theta_x^{-1} : \sigma x \cong_{\tilde{\mathcal{A}}_+} \sigma' x$ as required.

It is routine to show that the assignment from x to x' is monotonic, preserves cardinality and unions, hence it is generated by a map of event structures $f' : \mathcal{S} \rightarrow \mathcal{S}'$, which by construction preserves symmetry. By construction, f' satisfies the desired properties. The uniqueness of f' follows directly from Lemma 3.27. \square

B.5. Preservation of single-threadedness. Let $\sigma : S \rightarrow A^\perp \parallel B$ and $\tau : T \rightarrow B^\perp \parallel C$ be negative single-threaded prestrategies. Then, $\tau \odot \sigma$ is single-threaded. We first prove that the interaction $T \otimes S$ satisfies the single-threadedness conditions. More precisely, we prove by induction on φ that for any secured bijection $\varphi : x_S \parallel x_C \simeq y_A \parallel y_T$ representing (via Proposition 2.16) a configuration of $T \otimes S$, then

$$\varphi = \bigsqcup_{1 \leq i \leq n} \varphi_i$$

where each φ_i is a secured bijection with a unique minimal event. Indeed, assume $\varphi \xrightarrow{(c,d)} \varphi'$ where φ' fails this condition. Necessarily, either $c = (1, s^+)$ or $d = (2, t^+)$, *w.l.o.g.* assume the first. Then, the immediate predecessors of (c, d) in $\leq_{\varphi'}$ must be $((1, s_1^-), d_1), \dots, ((1, s_p^-), d_p)$ (using Lemma 2.7 of [CCRW17]), with $s_i \rightarrow_S s$. By hypothesis, there are $1 \leq i, j \leq p$ and distinct $1 \leq k \neq l \leq n$ such that $((1, s_i^-), d_i) \in \varphi_k$ and $((1, s_j^-), d_j) \in \varphi_l$. But φ_k (resp. φ_l) must contain an event synchronized with $\text{init}(s_i)$ (resp. $\text{init}(s_j)$). Since σ is single-threaded and $s_i, s_j \in [s]$ we have $\text{init}(s_i) = \text{init}(s_j)$, which contradicts $\varphi_l \cap \varphi_k = \emptyset$.

Now, we go on to prove single-threadedness.

(1) Prime secured bijections have no non-trivial decomposition as above, therefore they have a unique minimal event. This is true in particular for the *visible* prime secured bijections. Condition (1) of single-threadedness follows then from the fact used in the proof of Lemma 4.3 that a minimal event in the interaction of negative strategies is always visible.

(2) Finally, assume there is a minimal conflict $\varphi \sim \psi$ in $T \odot S$ between visible prime secured bijections. This means that there are non-necessarily visible prime secured bijections $\varphi' \subseteq [\varphi]_{T \otimes S}$, $\psi' \subseteq [\psi]_{T \otimes S}$, such that $\varphi' \sim \psi'$ in $T \otimes S$. Writing φ'' (resp. ψ'') for φ' (resp. ψ') without its top event, minimality of $\varphi' \sim \psi'$ means that $\varphi'' \cup \psi''$ is a valid secured bijection. Therefore, it decomposes:

$$\varphi'' \cup \psi'' = \bigsqcup_{1 \leq i \leq n} \varphi_i$$

With each φ_i a secured bijection having exactly one minimal event. If $n = 1$, we are done since as remarked the unique minimal event is necessarily visible. Otherwise, there are at least two φ_i, φ_j with distinct minimal events.

Then, using Lemma 2.15, $\varphi' \sim \psi'$ implies that their top elements have the form $((1, s_{\varphi'}), d_{\varphi'})$ and $((1, s_{\psi'}), d_{\psi'})$ with $s_{\varphi'} \sim_{S} s_{\psi'}$, or $(c_{\varphi'}, (2, t_{\varphi'}))$ and $(c_{\psi'}, (2, t_{\psi'}))$ with $t_{\varphi'} \sim_{T} t_{\psi'}$, *w.l.o.g.* say the first. By receptivity and courtesy of σ , we have $\text{pol}(s_{\varphi'}) = \text{pol}(s_{\psi'}) = +$. Since $n \geq 2$, there are $(c_1, d_1) \rightarrow_{\varphi} ((1, s_{\varphi'}^+), d_{\varphi'})$ with $(c_1, d_1) \in \varphi_i$ and $(c_2, d_2) \rightarrow_{\psi} ((1, s_{\psi'}^+), d_{\psi'})$ with $(c_2, d_2) \in \varphi_j$. By Lemma 2.7 of [CCRW17], as an immediate dependency of $((1, s_{\varphi'}^+), d_{\varphi'})$, we have $(c_1, d_1) = ((1, s_1^-), d_1)$ with $s_1 \rightarrow_S s_{\varphi'}$ (similarly, $(c_2, d_2) = ((1, s_2^-), d_2)$ with $s_2 \rightarrow_S s_{\psi'}$). But by single-threadedness of σ , $\text{init}(s_{\varphi'}) = \text{init}(s_{\psi'})$, so there should be an event synchronized with $\text{init}(s_1) = \text{init}(s_2)$ both in φ_i and φ_j , absurd.

APPENDIX C. INDEXES

LIST OF NOTATIONS

\leq_E	9
Causal order on event structure E .	
Con_E	9
Consistent subsets of event structure E .	
$\mathcal{C}(E)$	9
Configurations of event structure E .	
$[X]_E$	9
Down-closure in E of $X \subseteq E$, written $[e]$ for a singleton $\{e\}$.	
$[e)$	9
Dependencies of e , $[e) = [e] \setminus \{e\}$.	
$x \xrightarrow{e} \text{---}$	9
Configuration x extends with event e .	
\rightarrow	9
Immediate causality.	
$e_1 \# e_2$	9
Binary conflict between events e_1 and e_2 .	
$e_1 \rightsquigarrow e_2$	9
Immediate conflict between events e_1 and e_2 .	
pol_A	10
Polarity function on esp A .	
\mathcal{E}	11
Category of event structures and maps of event structures.	
$\sigma \cong \tau$	12
(Strong) isomorphism between prestrategies.	
$A \parallel B$	12
Simple parallel composition.	
A^\perp	12
Dual.	
$\sigma : A \twoheadrightarrow B$	12
Prestrategy $\sigma : S \rightarrow A^\perp \parallel B$, with S anonymous.	
$\mathbb{c}_A : \mathbb{C}_A \rightarrow A^\perp \parallel A$	12
Copycat strategy.	
$\sigma \wedge \tau : S \wedge T \rightarrow A$	13
Interaction of $\sigma : S \rightarrow A$ and $\tau : T \rightarrow A^\perp$.	
$\varphi : x \simeq y$	13
Secured bijection between x and y .	
$\Pi_i : S_1 \wedge S_2 \rightarrow S_i$	15
Projections of an interaction pullback.	
$\mathcal{B}_{\sigma, \tau}^{\text{sec}}$	13
Set of secured bijections between (configurations of) σ and τ .	
$\tau \otimes \sigma : T \otimes S \rightarrow A \parallel B \parallel C$	16
Interaction pullback of $\sigma : S \rightarrow A^\perp \parallel B$ and $\tau : T \rightarrow B^\perp \parallel C$.	

1	16
	Empty game/arena.	
$E \downarrow V$	17
	Projection of an event structure (hiding).	
$\tau \odot \sigma : T \odot S \rightarrow A^\perp \parallel C$	17
	Composition of $\sigma : S \rightarrow A^\perp \parallel B$ and $\tau : T \rightarrow B^\perp \parallel C$.	
$\text{lbl } \alpha$	19
	Label of an indexing function α .	
$\text{ind } \alpha$	19
	Copy index of an indexing function α .	
$!A$	19
	Expanded arena.	
$\text{just}(\alpha)$	19
	Justifier (unique immediate dependency) of α in $!A$.	
$\langle -1, \dots, -n \rangle$	19
	Any injective function $\omega^n \rightarrow \omega$.	
$x \cong y$	21
	Order-isomorphism between (possibly implicitly) ordered sets.	
$\sigma \approx \tau$	21
	Weak isomorphism.	
\mathcal{A}	24
	Event structure with symmetry $\mathcal{A} = (A, \tilde{A})$.	
\tilde{A}	24
	Isomorphism family of an ess \mathcal{A} .	
$\theta \upharpoonright x$	24
	Restriction of $\theta : x' \cong y' \in \tilde{A}$ to subconfiguration $x \subseteq x'$.	
$\theta : x \cong y$	25
	Order-preserving bijection between x and y .	
$\theta : x \cong_{\tilde{A}} y$	25
	$\theta : x \cong y \in \tilde{A}$.	
$\text{dom } \theta$	25
	Domain of θ .	
$\text{codom } \theta$	25
	Codomain of θ .	
\tilde{A}_-	25
	Set of negative reindexing isos on $!A$.	
\tilde{A}_+	25
	Set of positive reindexing isos on $!A$.	
$\theta_1 \parallel \theta_2$	25
	Parallel composition $\theta_1 \parallel \theta_2 : x \parallel y \simeq x' \parallel y'$ of bijections $\theta_1 : x \simeq x'$ and $\theta_2 : y \simeq y'$.	
$f \theta$	26
	Forward image $f \theta = \{(f a_1, f a_2) \mid (a_1, a_2) \in \theta\}$ of a bijection.	
\mathcal{E}_\sim	26
	Category of event structures with symmetry and morphisms preserving symmetry.	
$f \sim_{\tilde{B}} g$	26
	Maps $f, g : \mathcal{A} \rightarrow \mathcal{B}$ in \mathcal{E}_\sim are symmetric.	

$\tilde{\mathcal{A}}_-, \tilde{\mathcal{A}}_+$	32
Negative/positive isomorphism families of a tcg \mathcal{A} .	
$x \sqsubseteq^p y, \theta \sqsubseteq^p \theta'$	32
Extension with events of polarity $p \in \{-, +\}$.	
$\mathbb{C}_{\tilde{\mathcal{A}}}$	41
Isomorphism family on copycat.	
$x \sqsubseteq_A y$	41
The Scott order $y \supseteq^- x \cap y \sqsubseteq^+ x$.	
$A \times B$	46
Product of arenas A and B , defined as $A \parallel B$.	
$\varpi_A : !(A \times B) \xrightarrow{\mathbf{Tcg}} A$	46
Left projection in Cho .	
$\varpi_B : !(A \times B) \xrightarrow{\mathbf{Tcg}} B$	46
Right projection in Cho .	
$\langle\langle \sigma, \tau \rangle\rangle$	47
Pre-pairing of \sim -strategies $\sigma : !A \xrightarrow{\mathbf{Tcg}} !B$ and $\tau : !A \xrightarrow{\mathbf{Tcg}} !C$ with disjoint codomains.	
$\langle \sigma, \tau \rangle$	48
Pairing of \sim -strategies $\sigma : !A \xrightarrow{\mathbf{Tcg}} !B$ and $\tau : !A \xrightarrow{\mathbf{Tcg}} !C$.	
$\text{init}(s)$	49
Unique minimal event of $s \in S$ for a single-threaded prestrategy $\sigma : S \rightarrow A$.	
$\sigma : A \xrightarrow{\mathbf{Cho}} B$	50
Morphism from A to B in Cho , <i>i.e.</i> single-threaded negative $\sigma : !A \xrightarrow{\mathbf{Tcg}} !B$.	
$A \Rightarrow B$	51
Arrow arena.	
\sharp_-	52
Any injective encoding of the moves of an arena B as natural numbers.	

INDEX OF NOTIONS

- ~-receptive, 27
- Arena, 10
- Arrow arena, 51
- Binary conflict, 9
- Configurations, 9
- Copy index, 19
- Copycat strategy, 12
- Courteous, 17
 - componentwise, 40
- Dual
 - of a tcg, 32
 - of an esp, 12
 - of event structures with symmetry (ess), 25
- Empty arena, 45
- Event structure, 9
 - with polarities (esp), 10
 - with symmetry (ess), 24
 - with symmetry and polarities (essp), 24
- Game, 10
- Indexing function, 19
- Isomorphism
 - of prestrategies, 12
- Isomorphism family, 24
- Justifier, 19
- Label, 19
- Lifting
 - of maps of \mathcal{EP}_\sim , 42
- Minimal conflict, 9
- Morphisms
 - of event structures, 11
 - of event structures with symmetry (ess), 26
 - of prestrategies, 11
 - preserving symmetry, 26
- Negative
 - (pre)strategy, 45
 - arena, 10
- Pairing, 48
- Prestrategy, 11
 - from A to B , 12
- Prime configurations, 9
- Prime secured bijection, 14
- Product arena, 46
- Projection of an event structure, 17
- Projections (strategies), 46
- Pullback
 - bi-, 26
 - pseudo-, 26
- Receptive, 17
- Reindexing isomorphism, 21
 - negative, 21
 - positive, 21
- Secured bijection, 13
- Simple parallel composition
 - of event structures, 12
 - of event structures with symmetry (ess), 25
 - of thin concurrent games (tcg), 32
- Single-threaded, 49
- Strategy, 17
- Symmetric morphisms of ess, 26
- Symmetry, 25
- Thin
 - concurrent games (tcg), 32
 - essp, 31
- Weak isomorphism
 - between strategies on $!A$, 21
- Weak morphism
 - between strategies on $!A$, 21