



HAL
open science

Disentangling Parallelism and Interference in Game Semantics

Simon Castellan, Pierre Clairambault

► **To cite this version:**

Simon Castellan, Pierre Clairambault. Disentangling Parallelism and Interference in Game Semantics. 2021. hal-03182043v1

HAL Id: hal-03182043

<https://hal.science/hal-03182043v1>

Preprint submitted on 26 Mar 2021 (v1), last revised 11 Jul 2024 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

DISENTANGLING PARALLELISM AND INTERFERENCE IN GAME SEMANTICS

SIMON CASTELLAN AND PIERRE CLAIRAMBAULT

Inria, Univ Rennes, CNRS, IRISA
e-mail address: Simon.Castellan@inria.fr

Univ Lyon, EnsL, UCBL, CNRS, LIP, F-69342, LYON Cedex 07, France
e-mail address: Pierre.Clairambault@ens-lyon.fr

ABSTRACT. Game semantics is a denotational semantics presenting compositionally the computational behaviour of various kinds of effectful programs. One of its celebrated achievements is to have obtained full abstraction results for programming languages with a variety of computational effects, in a *single* framework. This is known as the *semantic cube* or *Abramsky’s cube*, which for sequential deterministic programs establishes a correspondence between certain conditions on strategies (“innocence”, “well-bracketing”, “visibility”) and the absence of matching computational effects.

Outside of the sequential deterministic realm, there are still a wealth of game semantics-based full abstraction results; but they no longer fit in a unified canvas. In particular, Ghica and Murawski’s fully abstract model for shared state concurrency (\mathbf{IA}_{\parallel}) does *not* have a matching notion of pure parallel program – we say that parallelism and interference (*i.e.* state plus semaphores) are *entangled*. In this paper we construct a causal version of Ghica and Murawski’s model, also fully abstract for \mathbf{IA}_{\parallel} . We provide compositional conditions *parallel innocence* and *sequentiality*, respectively banning interference and parallelism, and leading to four full abstraction results. To our knowledge, this is the first extension of Abramsky’s *semantic cube* programme beyond the sequential deterministic world.

INTRODUCTION

How to *prove* that a program P is correct, or equivalent to P' ? This simple question, prerequisite for formally validating software, lies at the heart of decades of work in semantics. Its study prompted a wealth of developments, each with its methodology and scope. *Operational semantics* axiomatizes execution directly on syntax, while *denotational semantics* gives meaning to programs by embedding them in a syntax-independent mathematical space.

Operational semantics is powerful and extensible, perfectly fit for formalization in a proof assistant – it is, for instance, behind the celebrated CompCert project [Ler09]. On the other hand, its deployment often follows from ad-hoc choices, and it is not robust to variations in the language. It is tied to syntax and struggles with compositionality¹. *Denotational*

Key words and phrases: Denotational semantics, game semantics, concurrent games.

¹Operational semantics *can* be made compositional, but behind lie denotational structures: for instance, the operational semantics behind the recent *Compositional CompCert* [SBCA15] “bears much in common” (quoting the paper) with Ghica and Tzevelekos’ operational reconstruction of game semantics [GT12].

semantics is syntax-independent, and often more principled. It is a great tool to reason about program equivalence (two programs being equivalent if they denote the same object), to prove general properties of languages (*e.g.* termination), and it comes with compositional reasoning principles. The wider mathematical space in which programs are embedded sometimes suggests new useful constructs (it is the birth story of Linear Logic [Gir87]). In exchange, it is more mathematically demanding and often quite brittle: distinct fragments of the same language may require radically different representations. Traditional denotational semantics (*e.g.* Scott domains) model programs as functions, through their input/output behaviour. Effects (*e.g.* state, non-determinism, *etc*) can be captured via monads which do not readily combine. Though combining effects has been a driving question in denotational semantics these past decades, it is hardly a streamlined process. For instance, though there is significant recent research activity around domain settings supporting probabilities and higher-order [SYW⁺16, VKS19], it is unclear how they combine with non-determinism [Gou17], let alone concurrency; nor how all these models relate together.

Game semantics [HO00, AJM00], though also denotational, takes a different approach: instead of a function it represents a program as a *strategy*, a collection of (representations of) its interactions against execution environments. Once executions are first-class citizens (called *plays*) one can characterise those achievable with specific effects. This led to a wealth of fully abstract models, rewarded in 2017 by the Alonzo Church Award (from the ACM SIGLOG, the EATCS, the EACSL, and the Kurt Gödel Society). To cite the announcement:

“Game semantics has changed the landscape of programming language semantics by giving a unified view of the denotational universes of many different languages. This is a remarkable achievement that was not previously thought to be within reach.”

But are games models truly “unified”? For *deterministic sequential* programs, absolutely: various degrees of control and state are indeed captured as additional conditions on one single canvas [AM99a] – this is the *semantic cube* or *Abramsky cube*. But beyond the sequential deterministic world, the picture is not so clear. The classic fully abstract models for *finite non-determinism* [HM99], for *probabilistic choice* [DH00] or for *parallelism* [GM08] all rely on the presence of state. Until recently, there were no fully abstract model for any of these *without state* – or in the language of game semantics, there were no notions of non-deterministic, probabilistic or parallel *innocence*. Following the phrasing of the title, our understanding of these effects was *entangled* with state.

However, this picture is currently shifting. Recently, two notions of non-deterministic innocence were proposed independently [CCW14, TO15] – the two settings also handling probabilistic innocence [TO14, CCPW18]. Technically, these settings differ significantly. But conceptually, both enrich strategies with explicit *branching* information. Though the novelty may seem minor, this is in fact a major schism with respect to traditional game semantics, in that this branching information is typically not observable. So instead of a strategy being merely a formal description of how a program is observed by a certain type of contexts, the model starts to carry more intensional, *causal* information, typically inaccessible to the environment but which nonetheless finds its use in capturing compositionally the computational behaviour expressible by certain programming features. This suggests that to disentangle parallelism and state, we must adequately represent the *branching structure* of parallel computation, the (non-observable) causal patterns of pure parallel programs.

Enter *concurrent games*. Concurrent games are a family of game semantics models questioning in various ways the premise that the basic building block should be totally, chronologically ordered *plays*. Pioneered by Mellès and others [AM99b, Mel04, MM07, FP09], they have lately been under intense development, prompted by new definitions due to Rideau and Winskel [RW11]. The name comes from their relationship with the so-called *true concurrency* approach to concurrency theory, following which one represents causal dependence and independence of events explicitly rather than resorting to interleavings. Besides making concurrent games a natural target to model concurrent languages and process calculi [CC16, CY19], it provides us with the required causal description of programs.

Contributions. We *disentangle* parallelism and state – or rather parallelism and *interference*, which we intend to also encompass semaphores. More precisely, we provide a fully abstract model of Idealized Parallel Algol ($\text{IA}_{//}$), the paradigmatic language used in the game semantics literature to study shared memory concurrency on top of a higher-order language. Our model is a causal version of that of Ghica and Murawski [GM08], which additionally supports compositional conditions of *parallel innocence* and *sequentiality* respectively eliminating interference and parallelism. Accordingly the paper presents *four* full abstraction results, following all combinations of parallelism and interference on top of the pure language PCF. Thus this is a *semantic square* [AM99a], the first such result pushing Abramsky’s programme beyond the sequential deterministic world.

Of the four full abstraction results glued together, three are classics: Hyland and Ong’s full abstraction for PCF [HO00], Abramsky and McCusker’s full abstraction for Idealized Algol (IA) [AM96], and Ghica and Murawski’s full abstraction for $\text{IA}_{//}$ [GM08]. The fourth result is a variation of the full abstraction for PCF with respect to parallel evaluation initially presented in conference format in [CCW15] – in particular, the notion of parallel innocence comes from there² and was developed as part as the first author’s PhD thesis [Cas17].

These four results [HO00, AM96, GM08, CCW15] vary significantly in their technical underpinnings. For the purposes of this paper, this left us with the task, more challenging than anticipated, of providing the *glue*. Accordingly, a significant part of the paper revisits the results of [HO00] and [AJM00] in a language closer to concurrent games, mixing ideas from HO [HO00], AJM [AJM00] and *asynchronous* [Mel05] games. In doing so we hope that this paper, gathering in a single framework several important developments of the field, could also serve as a modern entry point to game semantics. Accordingly we wrote it with the newcomer in mind, not assuming prior knowledge on game semantics. The development is self-contained, with however a number of details postponed to the appendix. In the development, we also take the time to show how our model relates to other game semantics frameworks, hopefully conveying some panoramic perspective on the field. More generally, we made an important effort in staying as pedagogical as possible. This of course, has a cost in that the paper is intimidatingly lengthy; and we hope the readers will excuse us for that.

Outline. In Section 1, we start by describing $\text{IA}_{//}$ and its fragments. In Section 2, we introduce our version of alternating games, its interpretation of PCF, and link with more traditional game semantics. In Section 3, we show how (the absence of) control and state may be captured via conditions of strategies – we present *Abramsky’s cube* and some of its

²The paper [CCW15] had two main contributions: a new games model called *thin concurrent games*, and *parallel innocence*. The detailed construction of the former appears in [CCW19], but not parallel innocence. The present paper provides, among other things, detailed proofs for the second contribution of [CCW15].

consequences. In Section 4, we present our causal fully abstract model for IA_{\parallel} , based on *thin concurrent games*. In Section 5 we develop one of the key contributions of this paper, *parallel innocence*: we leverage the causal description of programs offered by thin concurrent games to characterize the causal shapes definable with pure parallel higher-order programs. In Section 6, we study the *sequential* fragment of our causal games model, and by linking it with the sequential model of Sections 2 and 3 we show full abstraction results for IA and PCF . Finally, in Section 7 we prove our last full abstraction result, for PCF_{\parallel} .

1. IA_{\parallel} AND ITS FRAGMENTS

Idealized Parallel Algol (IA_{\parallel}) is a higher-order, simply-typed, call-by-name concurrent language with shared memory and semaphores. We also introduce fragments:

PCF_{\parallel} is the fragment without interference,
 IA is the fragment without parallelism, and
 PCF has neither interference nor parallelism.

1.1. **Types.** The **types** of IA_{\parallel} are the following, highlighting types relative to interference.

$$A, B ::= \mathbb{U} \mid \mathbb{B} \mid \mathbb{N} \mid A \rightarrow B \quad \text{PCF} \\
\mid \mathbf{ref} \mid \mathbf{sem} \quad \quad \quad +\text{interference}$$

Above, \mathbb{U} is a *unit* type with only one value, and \mathbb{B} and \mathbb{N} are types for *booleans* and *natural numbers*. In the presence of interference, \mathbf{ref} is a type for *references* storing natural numbers, while \mathbf{sem} is for *semaphores*. We refer to \mathbb{U}, \mathbb{B} and \mathbb{N} as *ground types*, and use \mathbb{X}, \mathbb{Y} to range over those. Let us now give the term constructions and typing rules.

1.2. **Terms and Typing.** We define the terms of the language directly via typing rules.

Contexts are lists $x_1 : A_1, \dots, x_n : A_n$. **Typing judgments** have the form $\Gamma \vdash M : A$ with Γ a context and A a type. In addition to Figure 1, we consider present an explicit exchange rule allowing us to permute the order of variable declarations in contexts. The eliminator rules for basic datatypes are restricted to eliminate only to ground types – general eliminators are defined as syntactic sugar: *e.g.* a conditional to \mathbf{ref} may be obtained as

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

The **bad variable** and **bad semaphore** constructs \mathbf{mkvar} and \mathbf{mksem} are a common occurrence in the game semantical literature. While a “*good*” reference is tied to a memory location, many game models also comprise so-called “*bad variables*” inhabiting \mathbf{ref} but not behaving as actual variables. Full abstraction results in the concerned games models [AM96, GM08] require a corresponding syntactic construct \mathbf{mkvar} allowing one to *form* bad variables by appending arbitrary read and write methods. The same holds for semaphores.

1.3. **Further syntactic sugar.** First of all, for any type A there is a **divergence** $\vdash \perp_A : A$, any looping program. Given $\Gamma \vdash M, N : \mathbb{U}$, an equality test $\Gamma \vdash M =_{\mathbb{U}} N : \mathbb{B}$ may be defined as $M; N; \mathbf{tt}$. Likewise, for $\Gamma \vdash M, N : \mathbb{B}$ we define $\Gamma \vdash M =_{\mathbb{B}} N : \mathbb{B}$ as $\mathbf{if} M N (\mathbf{if} N \mathbf{ff} \mathbf{tt})$, and $\Gamma \vdash M =_{\mathbb{N}} N : \mathbb{B}$ similarly, with the obvious recursive program.

PCF		
$\overline{\Gamma \vdash \mathbf{skip} : \mathbb{U}}$	$\overline{\Gamma \vdash \mathbf{tt} : \mathbb{B}}$	$\overline{\Gamma \vdash \mathbf{ff} : \mathbb{B}}$
$\overline{\Gamma, x : A \vdash M : B}$	$\overline{\Gamma \vdash M : A \rightarrow B}$	$\overline{\Gamma \vdash N : A}$
$\overline{\Gamma \vdash \lambda x^A. M : A \rightarrow B}$	$\overline{\Gamma \vdash MN : B}$	$\overline{\Gamma \vdash M : \mathbb{U} \quad \Gamma \vdash N : \mathbb{X}}$
$\overline{\Gamma \vdash M : \mathbb{B} \quad \Gamma \vdash N_1 : \mathbb{X} \quad \Gamma \vdash N_2 : \mathbb{X}}$	$\overline{\Gamma \vdash M : \mathbb{N}}$	$\overline{\Gamma \vdash M : \mathbb{N}}$
$\overline{\Gamma \vdash \mathbf{if} M N_1 N_2 : \mathbb{X}}$	$\overline{\Gamma \vdash \mathbf{succ} M : \mathbb{N}}$	$\overline{\Gamma \vdash \mathbf{pred} M : \mathbb{N}}$
$\overline{\Gamma \vdash M : \mathbb{N}}$	$\overline{\Gamma, x : \mathbb{X} \vdash M : \mathbb{Y} \quad \Gamma \vdash N : \mathbb{X}}$	$\overline{\Gamma \vdash M : A \rightarrow A}$
$\overline{\Gamma \vdash \mathbf{iszero} M : \mathbb{B}}$	$\overline{\Gamma \vdash \mathbf{let} x = N \mathbf{in} M : \mathbb{Y}}$	$\overline{\Gamma \vdash \mathcal{Y} M : A}$
+interference		
$\overline{\Gamma, x : \mathbf{ref} \vdash M : \mathbb{X}}$	$\overline{\Gamma \vdash M : \mathbf{ref} \quad \Gamma \vdash N : \mathbb{N}}$	$\overline{\Gamma \vdash M : \mathbf{ref}}$
$\overline{\Gamma \vdash \mathbf{newref} x := n \mathbf{in} M : \mathbb{X}}$	$\overline{\Gamma \vdash M := N : \mathbb{U}}$	$\overline{\Gamma \vdash !M : \mathbb{N}}$
$\overline{\Gamma, x : \mathbf{sem} \vdash M : \mathbb{X}}$	$\overline{\Gamma \vdash M : \mathbf{sem}}$	$\overline{\Gamma \vdash N : \mathbf{sem}}$
$\overline{\Gamma \vdash \mathbf{newsem} x := n \mathbf{in} M : \mathbb{X}}$	$\overline{\Gamma \vdash \mathbf{grab} M : \mathbb{U}}$	$\overline{\Gamma \vdash \mathbf{release} N : \mathbb{U}}$
$\overline{\Gamma \vdash M : \mathbb{N} \rightarrow \mathbb{U} \quad \Gamma \vdash N : \mathbb{N}}$	$\overline{\Gamma \vdash M : \mathbb{U} \quad \Gamma \vdash N : \mathbb{U}}$	
$\overline{\Gamma \vdash \mathbf{mkvar} MN : \mathbf{ref}}$	$\overline{\Gamma \vdash \mathbf{mksem} MN : \mathbf{sem}}$	
+parallelism		
$\overline{\Gamma, x_1 : \mathbb{X}, x_2 : \mathbb{X} \vdash M : \mathbb{Y} \quad \Gamma \vdash N_1 : \mathbb{X} \quad \Gamma \vdash N_2 : \mathbb{X}}$	$\overline{\Gamma \vdash \mathbf{let} \left(\begin{array}{l} x_1 = N_1 \\ x_2 = N_2 \end{array} \right) \mathbf{in} M : \mathbb{Y}}$	

 Figure 1: Typing rules for $\mathbf{IA}_{//}$

We refer to constants of ground type as **values**; we use v to range over those, and n, b or c to range over values of respective types \mathbb{N}, \mathbb{B} or \mathbb{U} . We introduce a n -ary case construct branching on all values of ground types. By abuse of notation, we write $V \subseteq_f \mathbb{X}$ for any finite subset of the values of ground type \mathbb{X} . Writing $V = \{v_1, \dots, v_n\}$, we set

$$\begin{array}{ll}
 \mathbf{case} M \mathbf{of} & \mathbf{let} x = M \mathbf{in} \\
 v_1 \mapsto N_1 & \mathbf{if} x =_{\mathbb{X}} v_1 \mathbf{then} N_1 \\
 v_2 \mapsto N_2 & \mathbf{else if} x =_{\mathbb{X}} v_2 \mathbf{then} N_2 \\
 \dots & \dots \\
 v_n \mapsto N_n & \mathbf{else if} x =_{\mathbb{X}} v_n \mathbf{then} N_n \\
 & \mathbf{else} \perp
 \end{array}
 \stackrel{\text{def}}{=}$$

of type \mathbb{Y} in context Γ if $\Gamma \vdash M : \mathbb{X}$ and $\Gamma \vdash N_i : \mathbb{Y}$ for all $1 \leq i \leq n$.

The **let** construct is crucial in this paper: as we shall see later on, strategies may evaluate a variable *once*, and provide a different continuation for each possible value. This behaviour cannot be replicated strictly without **let**, see Section 3.3.1 for a more detailed discussion.

1.4. Operational semantics. We give a small-step operational semantics, following [GM08]. We fix a countable set \mathcal{L} of **memory locations**. A **store** is a partial map $s : \mathcal{L} \rightarrow \mathbb{N}$ with finite domain where \mathbb{N} stands, overloading notations, for natural numbers. **Configurations** of the operational semantics are tuples $\langle M, s \rangle$ where s is a store with $\text{dom}(s) = \{\ell_1, \dots, \ell_n\}$ and $\Sigma \vdash M : A$ with $\Sigma = \ell_1 : \mathbf{ref}, \dots, \ell_i : \mathbf{ref}, \ell_{i+1} : \mathbf{sem}, \dots, \ell_n : \mathbf{sem}$.

Reduction rules have the form $\langle M, s \rangle \rightsquigarrow \langle M', s' \rangle$ where $\text{dom}(s) = \text{dom}(s')$; we write \rightsquigarrow^* for the reflexive transitive closure. If $\vdash M : \mathbb{X}$, we write $M \Downarrow$ if $\langle M, \emptyset \rangle \rightsquigarrow^* \langle v, \emptyset \rangle$ for some value v . We give in Figure 2 the reduction rules – there and from now on in the paper we use the notation \uplus to denote the usual set-theoretic union, when it is known disjoint. For rules which do not interact with the state, we omit the state component – it is simply left unchanged by stateless basic reductions, and propagated upwards by stateless context rules.

1.5. Fragment languages. Besides PCF, we consider three main languages of interest:

$$\begin{aligned} \text{PCF}_{//} &= \text{PCF} + \text{parallelism} \\ \text{IA} &= \text{PCF} + \text{interference} \\ \text{IA}_{//} &= \text{PCF} + \text{interference} + \text{parallelism} \end{aligned}$$

IA is a variant of *Idealized Algol with active expressions* [AM96], differing only in that it has semaphores. This is not a significant difference, as semaphores are definable from state in a sequential language. Likewise, $\text{IA}_{//}$ is close to the language of [GM08]: it differs only in that the parallelism operation is more general. For $\Gamma \vdash M : \mathbb{U}$ and $\Gamma \vdash N : \mathbb{U}$ we may define their parallel composition $\Gamma \vdash M \parallel N : \mathbb{U}$ (as in [GM08]) by

$$M \parallel N = \mathbf{let} \begin{pmatrix} x = M \\ y = N \end{pmatrix} \mathbf{in skip}.$$

Conversely, for *e.g.* $\Gamma \vdash N_1 : \mathbb{N}$, $\Gamma \vdash N_2 : \mathbb{N}$ and $\Gamma, x_1 : \mathbb{N}, x_2 : \mathbb{N} \vdash M : A$, the present parallel let construction is definable via state and parallel composition of commands:

$$\mathbf{let} \begin{pmatrix} x_1 = N_1 \\ x_2 = N_2 \end{pmatrix} \mathbf{in} M = \mathbf{newref} \ v_1 := 0 \ \mathbf{in} \\ \mathbf{newref} \ v_2 := 0 \ \mathbf{in} (v_1 := N_1 \parallel v_2 := N_2); M[!v_1/x_1, !v_2/x_2]$$

1.6. Observational Equivalence and Full Abstraction. Here, \mathcal{L} may refer to any of the fragments above. A \mathcal{L} -**context** for the judgment $\Gamma \vdash A$ is a term $C[\]$ of \mathcal{L} with a hole, s.t. for any $\Gamma \vdash M : A$ in \mathcal{L} we have $\vdash C[M] : \mathbb{U}$ obtained by replacing the hole with M . Two terms $\Gamma \vdash M, N : A$ of \mathcal{L} are \mathcal{L} -**observationally equivalent** iff

$$M \sim_{\mathcal{L}} N \quad \Leftrightarrow \quad \text{for all } C[\] \text{ a } \mathcal{L}\text{-context for } \Gamma \vdash A, \quad (C[M] \Downarrow \Leftrightarrow C[N] \Downarrow)$$

We omit \mathcal{L} when it is clear from the context. Observational equivalence is usually regarded as the canonical equivalence on programs: \mathcal{L} -observationally equivalent programs are interchangeable as long as the evaluation context is in \mathcal{L} . Accordingly, denotational semantics often aims to capture observational equivalence. An interpretation of programs $\llbracket - \rrbracket$ into some mathematical universe is called **fully abstract** whenever

$$M \sim N \quad \Leftrightarrow \quad \llbracket M \rrbracket = \llbracket N \rrbracket$$

for all $\Gamma \vdash M, N : A$. *Full abstraction* is a gold standard in denotational semantics, as it captures the best possible match between a language and its semantics, ensuring that the denotational semantics is *complete* for proving equivalence between programs.

<p style="text-align: center;">Basic red. for PCF</p> $ \begin{aligned} (\lambda x^A. M) N &\rightsquigarrow M[N/x] \\ \text{skip}; N &\rightsquigarrow N \\ \text{if } b N_{\text{tt}} N_{\text{ff}} &\rightsquigarrow N_b \\ \text{succ } n &\rightsquigarrow n + 1 \\ \text{pred } 0 &\rightsquigarrow 0 \\ \text{pred } (n + 1) &\rightsquigarrow n \\ \text{iszero } 0 &\rightsquigarrow \text{tt} \\ \text{iszero } (n + 1) &\rightsquigarrow \text{ff} \\ \mathcal{Y} M &\rightsquigarrow M(\mathcal{Y} M) \\ \text{let } x = v \text{ in } M &\rightsquigarrow M[v/x] \end{aligned} $	<p style="text-align: center;">Basic reductions for interference</p> $ \begin{aligned} \text{newref } x \text{ in } v &\rightsquigarrow v \\ \text{newsem } x \text{ in } v &\rightsquigarrow v \\ (\text{mkvar } M N) := n &\rightsquigarrow M n \\ !(\text{mkvar } M N) &\rightsquigarrow N \\ \text{grab}(\text{mksem } M N) &\rightsquigarrow M \\ \text{release}(\text{mksem } M N) &\rightsquigarrow N \end{aligned} $ <p style="text-align: center;">Interfering reductions</p> $ \begin{aligned} \langle !\ell, s \uplus \{\ell \mapsto n\} \rangle &\rightsquigarrow \langle n, s \uplus \{\ell \mapsto n\} \rangle \\ \langle \ell := n, s \uplus \{\ell \mapsto _ \} \rangle &\rightsquigarrow \langle \text{skip}, s \uplus \{\ell \mapsto n\} \rangle \\ \langle \text{grab}(\ell), s \uplus \{\ell \mapsto 0\} \rangle &\rightsquigarrow \langle \text{skip}, s \uplus \{\ell \mapsto 1\} \rangle \\ \langle \text{release}(\ell), s \uplus \{\ell \mapsto n\} \rangle &\rightsquigarrow \langle \text{skip}, s \uplus \{\ell \mapsto 0\} \rangle \quad (n > 0) \end{aligned} $
<p style="text-align: center;">Basic reduction for parallelism</p> $ \text{let } \begin{pmatrix} x_1 = v_1 \\ x_2 = v_2 \end{pmatrix} \text{ in } M \rightsquigarrow M[v_1/x_1, v_2/x_2] $	
<p style="text-align: center;">Stateless context rules</p> $ \begin{array}{c} \frac{M \rightsquigarrow M'}{M N \rightsquigarrow M' N} \quad \frac{M \rightsquigarrow M'}{\text{if } M N_1 N_2 \rightsquigarrow \text{if } M' N_1 N_2} \quad \frac{M \rightsquigarrow M'}{\text{succ } M \rightsquigarrow \text{succ } M'} \quad \frac{M \rightsquigarrow M'}{!M \rightsquigarrow !M'} \\ \\ \frac{M \rightsquigarrow M'}{\text{iszero } M \rightsquigarrow \text{iszero } M'} \quad \frac{N \rightsquigarrow N'}{M := N \rightsquigarrow M := N'} \quad \frac{M \rightsquigarrow M'}{\text{grab}(M) \rightsquigarrow \text{grab}(M')} \\ \\ \frac{M \rightsquigarrow M'}{\text{release}(M) \rightsquigarrow \text{release}(M')} \quad \frac{M \rightsquigarrow M'}{M := v \rightsquigarrow M' := v} \\ \\ \frac{N \rightsquigarrow N'}{\text{let } x = N \text{ in } M \rightsquigarrow \text{let } x = N' \text{ in } M} \\ \\ \frac{N_1 \rightsquigarrow N'_1}{\text{let } \begin{pmatrix} x_1 = N_1 \\ x_2 = N_2 \end{pmatrix} \text{ in } M \rightsquigarrow \text{let } \begin{pmatrix} x_1 = N'_1 \\ x_2 = N_2 \end{pmatrix} \text{ in } M} \\ \\ \frac{N_2 \rightsquigarrow N'_2}{\text{let } \begin{pmatrix} x_1 = N_1 \\ x_2 = N_2 \end{pmatrix} \text{ in } M \rightsquigarrow \text{let } \begin{pmatrix} x_1 = N_1 \\ x_2 = N'_2 \end{pmatrix} \text{ in } M} \end{array} $	
<p style="text-align: center;">Stateful context rules</p> $ \begin{array}{c} \frac{\langle M[\ell/x], s \uplus \{\ell \mapsto n\} \rangle \rightsquigarrow \langle M'[\ell/x], s' \uplus \{\ell \mapsto n'\} \rangle}{\langle \text{newref } x := n \text{ in } M, s \rangle \rightsquigarrow \langle \text{newref } x := n' \text{ in } M', s' \rangle} \quad (\ell \in \mathcal{L} \text{ fresh}) \\ \\ \frac{\langle M[\ell/x], s \uplus \{\ell \mapsto n\} \rangle \rightsquigarrow \langle M'[\ell/x], s' \uplus \{\ell \mapsto n'\} \rangle}{\langle \text{newsem } x := n \text{ in } M, s \rangle \rightsquigarrow \langle \text{newsem } x := n' \text{ in } M', s' \rangle} \quad (\ell \in \mathcal{L} \text{ fresh}) \end{array} $	

 Figure 2: Operational semantics of IA_{\parallel}




Figure 3: An affine arena



Figure 4: $\mathbb{U} \multimap (\mathbb{U} \otimes \mathbb{U})$ 

Figure 5: An alternating play

2. GAME SEMANTICS FOR PCF





2.1. Games and Strategies. We present first a game semantics of PCF. Though it is sequential, our presentation is non-standard, somewhat mixing features of AJM [AJM00], HO [HO00] and asynchronous games [Mel05] – this is to facilitate the interplay between all the games models involved. We skip a number of details, found in Appendix B.

Game semantics presents higher-order computation as an exchange of tokens between two players, called “Player” and “Opponent”. *Player* stands for the program under evaluation – events/moves attributed to Player are observable computational events resulting from its execution: calls to variables, program phrases converging to a value. *Opponent* stands for the execution environment. Their interaction follows rules depending on the type of the program under scrutiny. In setting up a game semantics the first step is to extract from the type a structure, called a *game* or an *arena*, which presents all the observable computational events available when interacting on this type, along with their respective causal dependencies.

2.1.1. Affine arenas. We first introduce our representation of types as games in the *affine* case, *i.e.* if any computational event can appear at most once – this is merely to first help the reader build up intuition before handling replication.

Consider $(\mathbb{U} \rightarrow \mathbb{U}) \rightarrow \mathbb{B}$, where affinity implies that each argument may be called at most once. Once a call-by-name execution on that type is initiated, the available observable events are the following: (1) the term may directly converge to \mathbf{tt} or \mathbf{ff} , without evaluating its argument; (2) it may call its argument (*i.e.* it evaluates to $\lambda f^{\mathbb{U} \rightarrow \mathbb{U}}. M$ with M having f in head position). In the case (2) the control goes back to the environment, which plays for f : it may prompt f to return the unique value \mathbf{skip} , or to itself call its argument. Finally, if f calls its argument, the corresponding sub-term may reduce to a value.

Overall, these events along with their causal dependencies give rise to the diagram in Figure 3. It is read from top to bottom, with the dashed lines representing the dependency relation. Nodes are called *moves* or *events*, and are labeled with a polarity, $-$ for events due to the environment, and $+$ for events due to the program. Finally, the wiggly line between \mathbf{tt}^+ and \mathbf{ff}^+ indicates *conflict*: it represents the fact that only one of these two values may be observed in one execution, whereas all the other pairs of events could conceivably appear together. The reader may convince themselves that indeed, the diagram does represent the observable events in a call-by-name evaluation of $(\mathbb{U} \rightarrow \mathbb{U}) \rightarrow \mathbb{B}$ as outlined in the previous paragraph. We insist that those are the computational events that are *observable* in the interface with the environment: the program may perform internal computation; a program


 Figure 6: \mathbb{U}

 Figure 7: \mathbb{B}

 Figure 8: \mathbb{N}

 Figure 9: $!\mathbb{U}$

in an extension of PCF with state could possibly store and read values from a local variable, *etc.* But those are not *observable* by a context, thus are not represented in the arena.

To formalize the arena as a mathematical structure, we use *event structures*³:

Definition 2.1. An **event structure** (*es*) is a triple $E = (|E|, \leq_E, \#_E)$, where $|E|$ is a (countable) set of **events**, \leq_E is a partial order called **causal dependency** and $\#_E$ is an irreflexive symmetric binary relation on $|E|$ called **conflict**, satisfying:

$$\begin{aligned} \text{finite causes: } & \forall e \in |E|, [e]_E = \{e' \in |E| \mid e' \leq_E e\} \text{ is finite} \\ \text{conflict inheritance: } & \forall e_1 \#_E e_2, \forall e_2 \leq_E e'_2, e_1 \#_E e'_2. \end{aligned}$$

An **event structure with polarities** (*esp*) is an event structure A together with a function $\text{pol}_A : |A| \rightarrow \{-, +\}$ assigning to each event a *polarity*.

Figure 3 displays an *esp*. The wiggly line indicates conflict, but we will not put wiggly lines between all conflicting pairs of events, as long as missing conflicts may be deduced by *conflict inheritance*. A conflict that cannot be deduced by inheriting an earlier conflict is called a **minimal conflict**. As with Figure 3, we will represent types as *esps*. In fact, the event structures arising via the interpretation of types have a very restricted form. In the definition below, we use the notation $e \rightarrow_E e'$ in an event structure E to mean **immediate causality**, *i.e.* $e <_E e'$ with no other event strictly in between.

Definition 2.2. An **arena** is an *esp* $(A, \leq_A, \#_A, \text{pol}_A)$ satisfying:

$$\begin{aligned} \text{alternating: } & \text{if } a_1 \rightarrow_A a_2, \text{pol}_A(a_1) \neq \text{pol}_A(a_2), \\ \text{forestial: } & \text{if } a_1 \leq_A a \text{ and } a_2 \leq_A a, \text{ then } a_1 \leq_A a_2 \text{ or } a_2 \leq_A a_1, \\ \text{race-free: } & \text{if } a_1, a_2 \in |A| \text{ are in minimal conflict, then } \text{pol}_A(a_1) = \text{pol}_A(a_2). \end{aligned}$$

Besides, a **--arena** additionally satisfies the condition:

$$\text{negative: if } a \in \min(A), \text{ then } \text{pol}_A(a) = -,$$

where $\min(A)$ stands for the set of minimal events of A .

Types will only yield **--arenas**, but throughout the paper we will use the general case. Finally, though we motivated Definition 2.1 with arenas, event structures will have other uses. Notably, from Section 4 onwards, strategies will also be event structures.

2.1.2. Basic Constructions. We give a few basic constructions on event structures and arenas which will allow us to construct in a systematic way, from any type of PCF, a **--arena**.

We give **--arenas** for the ground types of PCF, in Figures 6, 7 and 8, using the same notations \mathbb{U}, \mathbb{B} and \mathbb{N} for the arenas as for the types. For \mathbb{N} , even though the picture only shows conflict between neighbours, all positive events are meant to be in pairwise conflict.

³More precisely, those are *prime event structures with binary conflict*.

We write $\mathbf{1}$ for the **empty es**, with no event. If A is an *esp*, we write A^\perp for its **dual**, the *esp* with same events, causality and conflict, but the opposite polarities, *i.e.* $\text{pol}_{A^\perp}(a) = -\text{pol}_A(a)$ for all $a \in |A|$. The *simple parallel composition* is defined as follows.

Definition 2.3. If E_1, E_2 are two es, their **simple parallel composition** $E_1 \parallel E_2$ has

$$\begin{aligned} \text{events:} & & |E_1 \parallel E_2| &= \{1\} \times |E_1| \uplus \{2\} \times |E_2| \\ \text{causality:} & (i, e) \leq_{E_1 \parallel E_2} (j, e') & \Leftrightarrow & i = j \ \& \ e \leq_{E_i} e' \\ \text{conflict:} & (i, e) \#_{E_1 \parallel E_2} (j, e') & \Leftrightarrow & i = j \ \& \ e \#_{E_i} e'. \end{aligned}$$

Moreover, if E_1 and E_2 have polarities (*i.e.* are *esp*), then $E_1 \parallel E_2$ also has polarities, defined as $\text{pol}_{E_1 \parallel E_2}(1, e) = \text{pol}_{E_1}(e)$ and $\text{pol}_{E_1 \parallel E_2}(2, e) = \text{pol}_{E_2}(e)$.

By extension, we often write $X \parallel Y$ for the tagged disjoint union $(\{1\} \times X) \uplus (\{2\} \times Y)$ of two sets X and Y . In the simple parallel composition of arenas A and B , the two are side by side with no interaction. The arena $A \parallel B$ adequately represents a tensor type $A \otimes B$ where the two resources A and B may be accessed in any order – although PCF does not have such a type, this construction will play an important role in the sequel. We also introduce the **product** $A_1 \& A_2$ of A_1 and A_2 --arenas, defined as for $A_1 \parallel A_2$ with conflict

$$(i, e) \#_{A_1 \& A_2} (j, e') \Leftrightarrow i \neq j \vee (i = j \wedge e \#_{A_i} e'),$$

i.e. A_1 and A_2 are in conflict. The constructions \parallel and $\&$ have obvious n -ary generalizations. We also introduce another construction on arenas, the affine arrow $A \multimap B$:

Definition 2.4. Let A, B be arenas with B **pointed**, *i.e.* with *exactly one* minimal $b_0 \in |B|$.

The **affine arrow** $A \multimap B$ has the components of $A^\perp \parallel B$ except for causality, set as:

$$\leq_{A \multimap B} = \leq_{A^\perp \parallel B} \cup \{(2, b_0), (1, a) \mid a \in |A|\}.$$

This completes an interpretation of PCF types as pointed --arenas capturing the causal dependency between computational events in an affine evaluation. For instance, on $A \multimap B$ computation starts in B , but as soon as the initial move of B has been played computation in A may start, with polarity reversed. At this point, the reader may verify that indeed, the arena $(\mathbb{U} \multimap \mathbb{U}) \multimap \mathbb{B}$ obtained by applying these constructions is indeed the one in Figure 3.

2.1.3. *General arrow.* Definition 2.4 suffices for the types of PCF (which yield pointed arenas). But we aim to show that strategies have the structure of a *Seely category*, a traditional categorical model for Intuitionistic Linear Logic – and that structure includes tensors, which do not preserve pointedness. To generalize $A \multimap B$ for B non-pointed, it is natural to set one copy of A for each initial move of B . More concretely, $A \multimap B$ has events and polarities

$$|A \multimap B| = (\|_{b \in \min(B)} A)^\perp \parallel B,$$

where $\min(B)$ is the set of *minimal events* of B . The order has $(2, b) \leq (2, b')$ iff $b \leq_B b'$, $(1, (b, a)) \leq (1, (b', a'))$ iff $b = b'$ and $a \leq_A a'$, $(2, b) \leq (1, (b', a))$ iff $b = b'$, and $(1, (b, a)) \leq (2, b')$ never, exactly matching the arrow arena of HO games [HO00]. But having two copies of A is in tension with affineness, so we use conflict to tame this copying. The construction is illustrated in Figure 4, displaying the arena $\mathbb{U} \multimap (\mathbb{U} \otimes \mathbb{U})$. There are two copies of the \mathbb{U} on the left, but still, linearity is guaranteed by the addition of conflict.

To define conflict, writing $\chi_{A, B} : |A \multimap B| \rightarrow |A^\perp \parallel B|$ for the obvious map, we use:

Lemma 2.5. *Consider A and B two $--$ arenas.*

Then, there is a unique $\#_{A \multimap B}$ making $A \multimap B$ a $--$ arena such that for all down-closed finite $x \subseteq |A \multimap B|$, $x \in \mathcal{C}(A \multimap B)$ iff $\chi_{A,B} x \in \mathcal{C}(A^\perp \parallel B)$ with $\chi_{A,B}$ injective on x .

Proof. See Appendix B.1.1. □

2.1.4. *Playing on Arenas.* Now we formulate a notion of execution, relying on the fact that event structures support a natural notion of *state* or *position*, called *configuration*.

Definition 2.6. A (finite) **configuration** of es E is a finite set $x \subseteq |E|$ which is

$$\begin{aligned} \text{down-closed: } & \forall e \in x, \forall e' \in |E|, e' \leq_E e \implies e' \in x \\ \text{consistent: } & \forall e, e' \in x, \neg(e \#_E e'). \end{aligned}$$

We write $\mathcal{C}(E)$ for the set of finite configurations on E .

For $x, y \in \mathcal{C}(E)$, we write $x \multimap y$ if there is $e \in |E|$ such that $e \notin x$ and $y = x \cup \{e\}$; \multimap is the **covering relation**. If $x \multimap x \cup \{e\}$, we say that x **enables** e or **extends by** e , written $x \vdash_E e$. Configurations of an arena represent *valid execution states*. We may now leverage this to define *plays*, which provide a mathematical notion of *execution*.

Definition 2.7. An **alternating play** on arena A is a sequence $s = s_1 \dots s_n$ which is:

$$\begin{aligned} \text{valid: } & \forall 1 \leq i \leq n, \{s_1, \dots, s_i\} \in \mathcal{C}(A), \\ \text{non-repetitive: } & \forall 1 \leq i, j \leq n, s_i = s_j \implies i = j, \\ \text{alternating: } & \forall 1 \leq i \leq n-1, \text{pol}_A(s_i) \neq \text{pol}_A(s_{i+1}), \\ \text{negative: } & \text{if } n \geq 1, \text{ then } \text{pol}_A(s_1) = -. \end{aligned}$$

We write $\Downarrow\Uparrow\text{-Plays}(A)$ for the set of alternating plays on A .

The notation $\Downarrow\Uparrow\text{-Plays}(A)$ means to suggest that an alternating play has two possible states: O if s has even length and the last move (if any) is by Player, and P otherwise: each new move transitions between them. We denote the empty play with ε , and the prefix ordering with \sqsubseteq . In the sequel we sometimes apply $\Downarrow\Uparrow\text{-Plays}(-)$ to esps other than arenas.

Plays record individual executions, by giving a chronological account of events observed throughout computation. For instance, Figure 5 displays a play on the arena $(\mathbb{U} \multimap \mathbb{U}) \multimap \mathbb{B}$ of Figure 3. It is also read from top to bottom. Each move corresponds to a node in Figure 3 – as each move in the arena corresponds to a given type component, the identity of each move in Figure 5 is signified by its position under the matching type component.

2.1.5. *Strategies.* Given a term of type A we may, given the adequate technical machinery, ask whether a given play describes a valid execution for that term. The play of Figure 5, for instance, describes a valid execution for $\lambda f^{\mathbb{U} \rightarrow \mathbb{U}}. f \text{ skip}; \mathbf{tt} : (\mathbb{U} \rightarrow \mathbb{U}) \rightarrow \mathbb{B}$: after Opponent starts computation, reduction immediately gets stuck with a variable f in head position. This is an observable event, corresponding to Player calling its argument with \mathbf{q}^+ . Then, Opponent proceeds to call his argument with \mathbf{q}^- , triggering the evaluation of the subterm **skip**. This (trivially) converges to a value, which is observable and corresponds to \checkmark^+ . The control goes back to f (Opponent), which evaluates to **skip** as well via observable \checkmark^- . This triggers the evaluation of **tt**, leading to the observable \mathbf{tt}^+ that terminates computation.

Figure 5 represents one possible execution of $\lambda f^{\mathbb{U} \rightarrow \mathbb{U}}. f \text{ skip}; \mathbf{tt} : (\mathbb{U} \rightarrow \mathbb{U}) \rightarrow \mathbb{B}$. In general a term is represented by a *strategy*, which aggregates all possible executions.




Figure 10: $(\mathbb{U} \rightarrow \mathbb{U}) \rightarrow \mathbb{B}$ with replication Figure 11: A configuration of $!(\mathbb{U} \rightarrow \mathbb{U}) \rightarrow \mathbb{B}$

Definition 2.8. A **alternating strategy** $\sigma : A$ on $--$ -arena A is $\sigma \subseteq \downarrow\uparrow\text{-Plays}(A)$ which is:

$$\begin{aligned}
 \text{non-empty:} & \quad \varepsilon \in \sigma \\
 \text{prefix-closed:} & \quad \forall s \sqsubseteq s' \in \sigma, s \in \sigma \\
 \text{deterministic:} & \quad \forall sa_1^+, sa_2^+ \in \sigma, a_1 = a_2 \\
 \text{receptive:} & \quad \forall s \in \sigma, sa^- \in \downarrow\uparrow\text{-Plays}(A) \implies sa \in \sigma
 \end{aligned}$$

An **alternating prestrategy** $\sigma : A$ satisfies all these conditions except for *receptive*.

In this definition we have started using a convention followed throughout this paper: when introducing an event, we sometimes annotate it with a superscript to indicate its polarity. For instance, “ $\forall sa_1^+ \in \sigma, \dots$ ” is a shorthand for “ $\forall sa_1 \in \sigma$ such that $\text{pol}(a_1) = +, \dots$ ”.

We will see later on how to compute the strategy for a term. It is a strength of game semantics that this may be done either compositionally by induction on the syntax following the methodology of denotational semantics, or operationally via an abstract machine [GT12].

2.2. Replication and symmetry. In this paper we introduce early on the machinery for replication. It requires a small jump in abstraction, but fixes the arenas once and for all.

2.2.1. Arenas with symmetry. Figure 3 displays the arena corresponding to affine executions⁴ on type $(\mathbb{U} \rightarrow \mathbb{U}) \rightarrow \mathbb{B}$. To go beyond affineness, we *expand* the arena to allow multiple calls to arguments – for $(\mathbb{U} \rightarrow \mathbb{U}) \rightarrow \mathbb{B}$, we obtain an infinite arena as drawn in Figure 10.

In the picture, it seems like *e.g.* all moves \mathbf{q}^+ are interchangeable. This is true in spirit but every move must be a distinct event of the arena. Concretely, the expanded arena is computed following the methodology of linear logic: the type $(\mathbb{U} \rightarrow \mathbb{U}) \rightarrow \mathbb{B}$ is represented by $!(\mathbb{U} \multimap \mathbb{U}) \multimap \mathbb{B}$ rather than $(\mathbb{U} \multimap \mathbb{U}) \multimap \mathbb{B}$. Here, $!$ is an *exponential modality* in the sense of linear logic. The full definition of $!A$ will appear in Definition 2.10, but its events are $|!A| = \mathbb{N} \times |A|$, pairs (n, a) where n is called a *copy index*. So in reality, a precise picture of the arena for $(\mathbb{U} \rightarrow \mathbb{U}) \rightarrow \mathbb{B}$ with replication would be a version of Figure 10 where some events are tagged by *copy index* – see Figure 11 for an example of a configuration of $!(\mathbb{U} \multimap \mathbb{U}) \multimap \mathbb{B}$ with explicit indices pictured as grey subscripts.

Expanding the arena so opens up the way to replication without compromising the *non-repetitive* condition: a strategy may replay the “same” move but with different copy indices. But then, it is necessary to identify strategies behaving in the same way save for

⁴Affineness is enforced by *non-repetitive* in Definition 2.7. Rather than expand arenas, it is tempting to simply lift it. For this to be sound, it becomes then necessary to include additional structure in plays: the *justification pointers*. This is the choice made in *HO games* [HO00]. This will be detailed in Section 2.4.

the choice of copy indices. To that end, following the approach initiated in [CCW14] we enrich arenas with a notion of *symmetry*, capturing reindexings between configurations.

Definition 2.9. An **isomorphism family** on event structure E is a set $\mathcal{S}(E)$ of bijections between configurations of E , satisfying the additional conditions:

- groupoid:* $\mathcal{S}(E)$ contains identity bijections; is closed under composition and inverse.
- restriction:* for all $\theta : x \simeq y \in \mathcal{S}(E)$ and $x \supseteq x' \in \mathcal{C}(E)$, there is a (necessarily) unique $\theta \supseteq \theta' \in \mathcal{S}(E)$ such that $\theta' : x' \simeq y'$.
- extension:* for all $\theta : x \simeq y \in \mathcal{S}(E)$, $x \subseteq x' \in \mathcal{C}(E)$, there is a (not necessarily unique) $\theta \subseteq \theta' \in \mathcal{S}(E)$ such that $\theta' : x' \simeq y'$.

Then $(E, \mathcal{S}(E))$ is an **event structure with symmetry (ess)**. If A has polarities preserved by $\mathcal{S}(A)$, A is an **event structure with symmetry and polarities (essp)**.

If A is an ess, we refer to the elements of $\mathcal{S}(A)$ as **symmetries**. We write $\theta : x \cong_A y$ to mean that $\theta : x \simeq y$ is a bijection such that $\theta \in \mathcal{S}(A)$, and write $x = \text{dom}(\theta)$ and $y = \text{cod}(\theta)$. It is an easy exercise to prove that symmetries are automatically order-isomorphisms [Win07], where configurations inherit a partially ordered structure from the causal dependency of A . We regard isomorphism families as *proof-relevant* equivalence relations: they convey the information of which configurations are interchangeable, witnessed by an explicit bijection.

From now on, **arenas** have an isomorphism family. It comprises only identity symmetries on basic arenas $\mathbb{U}, \mathbb{B}, \mathbb{N}$ and $\mathbf{1}$. The previous constructions on arenas extend transparently: A^\perp has the same symmetries as A . The symmetries on $A \parallel B$ are those of the form

$$\begin{aligned} \theta_A \parallel \theta_B & : x_A \parallel x_B \cong y_A \parallel y_B \\ & (1, a) \mapsto (1, \theta_A(a)) \\ & (2, b) \mapsto (2, \theta_B(b)) \end{aligned}$$

for $\theta_A : x_A \cong_A y_A$ and $\theta_B : x_B \cong_B y_B$. Those on $A \& B$ are the symmetries on $A \parallel B$ that are bijections between configurations of $A \& B$, *i.e.* one of θ_A and θ_B must be empty. Note that these constructions \parallel and $\&$ apply to arbitrary event structures with symmetry.

For $A \multimap B$, if $x, y \in \mathcal{C}(A \multimap B)$ and $\theta : x \simeq y$ is any bijection, defining first $\chi_{A,B} \theta$ as

$$\chi_{A,B} x \stackrel{\chi_{A,B}^{-1}}{\simeq} x \stackrel{\theta}{\simeq} y \stackrel{\chi_{A,B}}{\simeq} \chi_{A,B} y,$$

we set $\theta : x \cong_{A \multimap B} y$ when θ is an order-isomorphism satisfying $\chi_{A,B} \theta : \chi_{A,B} x \cong_{A^\perp \parallel B} \chi_{A,B} y$.

The main arena construction introducing new symmetric events is the **exponential**:

Definition 2.10. Let A be a $--$ -arena. The $--$ -arena $!A$ has components

$$\begin{aligned} \text{events:} & \quad |!A| = \mathbb{N} \times |A| \\ \text{causality:} & \quad (i, a) \leq_{!A} (j, a') \Leftrightarrow i = j \ \& \ a \leq_A a' \\ \text{conflict:} & \quad (i, a) \#_{!A} (j, a') \Leftrightarrow i = j \ \& \ a \#_A a' \\ \text{polarities:} & \quad \text{pol}_{!A}(i, a) = \text{pol}_A(a) \end{aligned}$$

along with isomorphism family comprising as symmetries those bijections of the form

$$\begin{aligned} \theta & : \parallel_{n \in \mathbb{N}} x_n \cong \parallel_{n \in \mathbb{N}} y_n \\ & (n, a) \mapsto (\pi(n), \theta_n a) \end{aligned}$$

for some permutation $\pi \in \varsigma(\mathbb{N})$ and some family $(\theta_n)_{n \in \mathbb{N}}$ with $\theta_n : x_n \cong_A y_{\pi(n)}$ for all $n \in \mathbb{N}$.

This definition applies in general to any *ess*. Figure 9 shows the plain *esp* of $!U$ with copy indices indicated as grey subscripts – its symmetries are all order-isomorphisms between configurations. While $!(-)$ does not match a type construction of PCF, we shall follow Girard [Gir87] and define the arrow type of arenas *with* replication as $A \rightarrow B = !A \multimap B$.

2.2.2. Symmetry on plays and strategies. Symmetry will allow us to identify strategies, but it should also affect how strategies play. In the presence of explicit copy indices, a fundamental property is *uniformity*. Intuitively, a strategy is uniform if its behaviour does not depend (up to symmetry) on the specific copy indices used by its environment.

The first step towards uniformity is to transport symmetry to plays.

Definition 2.11. Let A be an arena and $s, t \in \downarrow\uparrow\text{-Plays}(A)$. We say that s and t are **symmetric**, written $s \cong_A t$, if s and t have the same length, and we have

$$\theta_{s,t}^j = \{(s_i, t_i) \mid 1 \leq i \leq j\} : \{s_1, \dots, s_j\} \cong_A \{t_1, \dots, t_j\}$$

a symmetry in $\mathcal{S}(A)$ for all $1 \leq j \leq n$; writing $s = s_1 \dots s_n$ and $t = t_1 \dots t_n$.

Those readers familiar with AJM games may find comfort in the following fact.

Fact 2.12. For an arena A , the tuple $\langle |A|, \text{pol}_A, \downarrow\uparrow\text{-Plays}(A), \cong_A \rangle$ is an AJM game [AJM00].

This ignores the Question/Answer labeling in AJM games, which we shall handle later on. The proof is a straightforward exercise. For the experts, we mention that this association of arenas to AJM games does not respect the arena constructions because constructions on AJM games enforce local alternation, while $\downarrow\uparrow\text{-Plays}(-)$ does not. As in HO games [HO00], in our presentation local alternation will only follow from the P-visibility condition.

From the connection with AJM games it seems natural to import the AJM uniformity:

Definition 2.13. For A an arena and $\sigma, \tau : A$ alternating prestrategies, we write $\sigma \approx \tau$ iff:

$$\begin{aligned} \rightarrow\text{-simulation: } & \forall sa^+ \in \sigma, t \in \tau, s \cong_A t \implies \exists b^+, tb^+ \in \tau \ \& \ sa^+ \cong_A tb^+ \\ \leftarrow\text{-simulation: } & \forall s \in \sigma, tb^+ \in \tau, s \cong_A t \implies \exists a^+, sa^+ \in \sigma \ \& \ sa^+ \cong_A tb^+ \\ \rightarrow\text{-receptive: } & \forall sa^- \in \sigma, t \in \tau, sa^- \cong_A tb^- \implies tb^- \in \tau \\ \leftarrow\text{-receptive: } & \forall s \in \sigma, tb^- \in \tau, sa^- \cong_A tb^- \implies sa^- \in \sigma \end{aligned}$$

This defines a per \approx on prestrategies⁵ on A . A prestrategy $\sigma : A$ is **uniform** iff $\sigma \approx \sigma$.

Uniformity is crucial. For the interpretation to respect β -equivalence we must identify strategies that play the “same moves”, but with different copy indices. For instance, we must consider equal the two strategies $\tau_0, \tau_1 : U \rightarrow U$ with unique maximal play:

$$\begin{array}{ccc} \tau_0 & : & !U \multimap U \\ & & \mathbf{q}^- \\ & & \mathbf{q}_0^+ \\ & & \checkmark_0^- \\ & & \checkmark^+ \end{array} \qquad \begin{array}{ccc} \tau_1 & : & !U \multimap U \\ & & \mathbf{q}^- \\ & & \mathbf{q}_1^+ \\ & & \checkmark_1^- \\ & & \checkmark^+ \end{array}$$

But this quotient is risky. Let us apply both τ_0 and τ_1 to $\sigma : !U$ with only maximal play $\mathbf{q}_0^- \checkmark_0^+$. Though we have yet to define composition, the application of τ_0 to σ must converge, while that of τ_1 to σ must diverge. So τ_0 and τ_1 cannot be safely identified as

⁵For *strategies*, \rightarrow , \leftarrow -receptive are subsumed by *receptive*. But these are necessary for uniformity to apply to *prestrategies* which might not be receptive – this generalization will be used in the technical development.

they are distinguishable. In fact here, the culprit is σ : it is not *uniform*. Since $\mathbf{q}_0^- \checkmark_0^+ \in \sigma$, uniformity of σ would imply that $\mathbf{q}_1^- \checkmark_1^+ \in \sigma$ as well, breaking the counter-example.

From now on, all **alternating (pre)strategies** are assumed uniform.

2.3. Interpretation of PCF. The interpretation follows the methodology of denotational semantics, resting on the fact that arenas and strategies form a category with adequate structure. In the main text we only outline this fairly routine construction – though this should be enough to read the paper – but the construction is detailed in Appendix B.

2.3.1. Category. The category $\Downarrow\Uparrow\text{-Strat}$ has objects $--$ arenas, and morphisms the *alternating strategies* (*strategies* for short) on $A \multimap B$. The **composition** of $\sigma : A \multimap B$ and $\tau : B \multimap C$

$$\tau \odot \sigma : A \multimap C$$

follows the usual game semantics process of *parallel interaction* followed by *hiding*.

First, the **pre-interactions** are sequences $u \in |(A \multimap B) \multimap C|^*$ satisfying *valid* of Definition 2.7. A pre-interaction u has three **restrictions**, with the following types:

$$u \upharpoonright A, B \in |A \multimap B|^*, \quad u \upharpoonright B, C \in |B \multimap C|^*, \quad u \upharpoonright A, C \in |A \multimap C|^*,$$

defined in the obvious way – see Appendix B.1.2. Given prestrategies $\sigma : A \multimap B$ and $\tau : B \multimap C$, an **interaction** $u \in \tau \circledast \sigma$ is a pre-interaction $u \in |(A \multimap B) \multimap C|^*$ satisfying:

$$u \upharpoonright A, B \in \sigma, \quad u \upharpoonright B, C \in \tau, \quad u \upharpoonright A, C \in \Downarrow\Uparrow\text{-Plays}(A \multimap C).$$

The **composition** of σ and τ comprises all $s \in \Downarrow\Uparrow\text{-Plays}(A \multimap C)$ with a **witness**:

$$\tau \odot \sigma = \{u \upharpoonright A, C \mid u \in \tau \circledast \sigma\};$$

it follows that $\tau \odot \sigma : A \multimap C$ is a prestrategy; and a strategy if σ and τ are.

Composition is associative on prestrategies, but admits identities only for strategies: the **copycat** strategies. If A is a $--$ arena and $s \in |A \multimap A|$, there are left and right *restrictions*

$$s \upharpoonright l \in |A|^*, \quad s \upharpoonright r \in |A|^*,$$

defined in the obvious way (see Appendix B.1.4). For $s \in \Downarrow\Uparrow\text{-Plays}(A \multimap A)$, s is a *copycat play* iff (1) for all even-length prefix $s' \sqsubseteq s$ we have $s' \upharpoonright l = s' \upharpoonright r$, and (2) for all $(1, (a_1, a_2)) \in |s|$, if $a_2 \in \min(A)$, then $a_1 = a_2$ – a move initial on the left must be justified by the same move on the right. Writing \mathfrak{c}_A for the set of all copycat plays, we have $\mathfrak{c}_A : A \multimap A$ a strategy as required. For any strategy $\sigma : A \multimap B$ we have $\mathfrak{c}_B \odot \sigma \odot \mathfrak{c}_A = \sigma$, making $\Downarrow\Uparrow\text{-Strat}$ a category.

Remark 2.14. Our model shares with AJM games [AJM00] the management of the equivalence \approx on strategies. All our constructions on strategies must preserve \approx . For most of them it is clear, but composition requires some care (see Appendix B.3.1). Operations on strategies therefore lift transparently to \approx -equivalence classes, and one can then consider $\Downarrow\Uparrow\text{-Strat}$ to have as morphisms \approx -equivalence classes of strategies (as is done in [AJM00]). This is fine, but it does contrast with how we (also following the practice in AJM games) often refer to specific concrete strategies as being “the interpretation of” specific terms. So we refrain from quotienting, and consider $\Downarrow\Uparrow\text{-Strat}$ as having *concrete* strategies as morphisms, and homsets $\Downarrow\Uparrow\text{-Strat}(A, B)$ additionally equipped with an equivalence relation \approx which all operations preserve. This way the interpretation of terms yields *concrete* representatives, but categorical laws only hold up to \approx . In the sequel we refer only to the plain algebraic structures (as in “symmetric monoidal closed category”, “cartesian closed category”, *etc*),

with it being understood that laws for these algebraic structures only hold up to \approx and that for any construction we consider, there is a proof obligation that it preserves \approx .

2.3.2. *Further structure.* If A and B are $--$ arenas, their *tensor* is simply $A \otimes B = A \parallel B$ their parallel composition. For $\sigma_1 : A_1 \multimap B_1$ and $\sigma_2 : A_2 \multimap B_2$, the strategy

$$\sigma_1 \otimes \sigma_2 : A_1 \otimes A_2 \multimap B_1 \otimes B_2,$$

defined via adequate restrictions (see Appendix B.2.1), plays as σ_1 on A_1, B_1 and σ_2 on A_2, B_2 – this gives a symmetric monoidal structure, with structural isomorphisms copycat strategies. Moreover, $\downarrow\uparrow\text{-Strat}$ is *cartesian*. Its *terminal object* is the empty $--$ arena $\mathbf{1}$; the **product** of A and B is the $A \& B$. This forms a *cartesian product*: there are projections

$$\pi_A : A \& B \multimap A \qquad \pi_B : A \& B \multimap B$$

acting as copycat, and for $\sigma : C \multimap A$ and $\tau : C \multimap B$, their **pairing** $\langle \sigma, \tau \rangle : C \multimap A \& B$ is defined simply as the as the set-theoretic union of σ and τ (modulo the obvious relabeling).

Finally, for any $--$ arenas Γ, A and B , there is an iso $\Gamma \otimes A \multimap B \cong \Gamma \multimap (A \multimap B)$, *i.e.* a bijection on events preserving and respecting all structure. This yields a bijection

$$\Lambda_{\Gamma, A, B} : \downarrow\uparrow\text{-Strat}(\Gamma \otimes A, B) \simeq \downarrow\uparrow\text{-Strat}(\Gamma, A \multimap B)$$

between the corresponding sets of strategies. Exploiting this, we define the **evaluation**

$$\text{ev}_{A, B} = \Lambda_{A \multimap B, A, B}^{-1}(\mathbb{C}_{A \multimap B}) : (A \multimap B) \otimes A \multimap B,$$

and the universal property for monoidal closure is then a direct verification. We conclude:

Proposition 2.15. *The category $\downarrow\uparrow\text{-Strat}$ is cartesian and symmetric monoidal closed.*

2.3.3. *Exponential.* On $\downarrow\uparrow\text{-Strat}$, $!$ gives an *exponential* in the sense of Linear Logic [Gir87]: a functor, with natural transformations $\text{der}_A : !A \multimap A$ and $\text{dig}_A : !A \multimap !!A$ making $(!, \text{der}, \text{dig})$ a comonad. Moreover, there are natural isomorphisms $\text{mon}_{A, B}^2 : !A \otimes !B \multimap !(A \& B)$ and $\text{mon}_{A, B}^0 : \mathbf{1} \multimap !\mathbf{1}$, satisfying the coherence laws of a *Seely category* [Mel09]. So the Kleisli category $\downarrow\uparrow\text{-Strat}_!$ is cartesian closed, and hence a model of the simply-typed λ -calculus. The construction is routine, and follows the lines of AJM games [AJM00] – see Appendix B.3.

In the sequel, given a Seely category \mathcal{C} and a morphism $f \in \mathcal{C}(!A, B)$, we shall write $f^\dagger \in \mathcal{C}(!A, !B)$ for its **promotion**, defined as $!f \circ \text{dig}_A$ – in particular, recall that Kleisli composition of $f \in \mathcal{C}_!(A, B)$ and $g \in \mathcal{C}_!(B, C)$ may be defined as $g \circ_! f = g \circ f^\dagger \in \mathcal{C}_!(A, C)$.

2.3.4. *Recursion.* Strategies on arena A may be partially ordered by inclusion; this forms a *pointed dcpo*. All operations on strategies are continuous with respect to \subseteq .

Writing $\downarrow\uparrow\text{-Strat}_!(A, B)$ for the dcpo of strategies on $A \rightarrow B = !A \multimap B$, the operation

$$\begin{aligned} F & : \downarrow\uparrow\text{-Strat}_!(\mathbf{1}, (A \rightarrow A) \rightarrow A) \rightarrow \downarrow\uparrow\text{-Strat}_!(\mathbf{1}, (A \rightarrow A) \rightarrow A) \\ & \quad \sigma \mapsto \lambda f^{A \rightarrow A}. f(\sigma f) \end{aligned}$$

written in λ -calculus syntax following the cartesian closed structure of $\downarrow\uparrow\text{-Strat}_!$, is continuous. Its least fixed point $\mathcal{Y}_A \in \downarrow\uparrow\text{-Strat}_!(\mathbf{1}, (A \rightarrow A) \rightarrow A)$ is transported to $\downarrow\uparrow\text{-Strat}_!(\Gamma, (A \rightarrow A) \rightarrow A)$ by composition with the terminal projection. For each $\sigma \in \downarrow\uparrow\text{-Strat}_!(\Gamma, A \rightarrow A)$,

$$\mathcal{Y}_A \sigma \approx \sigma(\mathcal{Y}_A \sigma)$$




Figure 12: Basic strategies for PCF




Figure 13: Interpretation of basic combinators

 Figure 14: Strategy for **pred**

so a fixed point operator up to \approx , as needed to interpret recursion.

It is a curiosity already in AJM games [AJM00] that the recursive equation for the fixpoint combinator must be solved in the domain of *concrete* strategies, rather than \approx -equivalence classes. To the best of our knowledge it is not known if the partial order induced by inclusion on \approx -equivalence classes of strategies has the adequate completeness properties to solve this, *i.e.* if the quotient of $\downarrow\uparrow\text{-Strat}$ and $\downarrow\uparrow\text{-Strat}_!$ by \approx are dcpo-enriched categories.

2.3.5. Interpretation. *Types* of PCF are interpreted as $--$ -arenas: we set $\llbracket \mathbb{U} \rrbracket = \mathbb{U}$, $\llbracket \mathbb{B} \rrbracket = \mathbb{B}$, $\llbracket \mathbb{N} \rrbracket = \mathbb{N}$ and $\llbracket A \rightarrow B \rrbracket = !\llbracket A \rrbracket \multimap \llbracket B \rrbracket$ yielding for any type A an arena $\llbracket A \rrbracket$. A *context* $\Gamma = x_1 : A_1, \dots, x_n : A_n$ is interpreted as $\llbracket \Gamma \rrbracket = \&_{1 \leq i \leq n} \llbracket A_i \rrbracket$. A term $\Gamma \vdash M : A$ yields

$$\llbracket M \rrbracket \in \downarrow\uparrow\text{-Strat}_!(\llbracket \Gamma \rrbracket, \llbracket A \rrbracket).$$

We skip the details of the interpretation of the λ -calculus combinators, which follows the standard interpretation of the simply-typed λ -calculus in a cartesian closed category [LS88].

We specify strategies for PCF combinators. For constants, $\llbracket \text{skip} \rrbracket : \mathbb{U}$, $\llbracket \text{tt} \rrbracket : \mathbb{B}$, $\llbracket \text{ff} \rrbracket : \mathbb{B}$ and $\llbracket n \rrbracket : \mathbb{N}$ are the corresponding obvious strategies replying immediately the corresponding value. For the others the interpretation is in Figure 13, annotating strategy operations with $!$ to emphasize that they are in the Kleisli category $\downarrow\uparrow\text{-Strat}_!$. The strategies used are in Figures 12, 14 and 15. Save for **let**, the diagram displays exhaustively their maximal plays,



Figure 15: Typical play of **let**

Figure 16: Pointer annotations on plays

defining them completely. For **let**, the strategy implements a *memoization* mechanism: it evaluates on \mathbb{X} obtaining a value v , then fed to the function argument each call, without re-evaluating it. The play shown for **let** is not maximal as Opponent could play some $q_{1,i_{n+1}}^-$ after. We will see in Section 3.2 that it *is* fully informative: there is only one *innocent* strategy that includes these plays. The interpretation is completed with $\llbracket \mathcal{Y}_A M \rrbracket = \mathcal{Y}_A \odot \llbracket M \rrbracket$.

This interpretation satisfies the main property expected of a denotational semantics:

Proposition 2.16 (Adequacy). *For any $\vdash M : \mathbb{U}$, $M \Downarrow$ if and only if $\llbracket M \rrbracket \Downarrow$.*

Note there are only two strategies on \mathbb{U} : the minimal $\{\varepsilon, q^-\}$ matching any *diverging* program, and the *converging* $\{\varepsilon, q^-, q^-\checkmark^+\}$. For $\sigma : \mathbb{U}$, we write $\sigma \Downarrow$ if σ converges and $\sigma \Uparrow$ if σ diverges. We omit the proof which is standard using logical relations, see *e.g.* [HO00].

This immediately entails the proof soundness for observational equivalence:

Corollary 2.17. *Let $\Gamma \vdash M, N : A$ be any terms of PCF. If $\llbracket M \rrbracket \approx \llbracket N \rrbracket$, then $M \sim N$.*

Proof. Assume $\llbracket M \rrbracket \approx \llbracket N \rrbracket$, and consider a context $C[-]$ such that $C[M] \Downarrow$. By Proposition 2.16, $\llbracket C[M] \rrbracket \Downarrow$. But $\llbracket C[M] \rrbracket = \llbracket C[-] \rrbracket \odot \llbracket M \rrbracket \approx \llbracket C[-] \rrbracket \odot \llbracket N \rrbracket = \llbracket C[N] \rrbracket$, so $C[N] \Downarrow$ by Proposition 2.16. The other direction also holds, hence $M \sim N$. \square

Computational adequacy is the standard to express that a model accurately describes computation in the language. In fact in game semantics the connection with operational semantics is much stronger, as highlighted earlier. We will elaborate on that in Section 3.

2.4. HO games. Before exploring this computational content, we highlight the connection with HO games [HO00], based on representing plays up to symmetry as *plays with pointers*.

2.4.1. Plays with pointers. First, a convention. For A a $--$ -arena and $s \in \Downarrow\uparrow\text{-Plays}(A)$, then $|s| \in \mathcal{C}(A)$ has two order structures: it is totally ordered chronologically as prescribed by s , and has a partial order imported from \leq_A . When representing plays, we often annotate them with the immediate causal dependency generating \leq_A . For instance, Figure 16 shows it for $s \in \Downarrow\uparrow\text{-Plays}(\llbracket (\mathbb{U} \rightarrow \mathbb{U}) \rightarrow \mathbb{B} \rrbracket)$ with $|s|$ displayed in Figure 11. The dashed lines represent immediate causal dependency in \leq_A , omitted when it coincides with juxtaposition. We call

these dashed lines *pointers*, going upwards from one event to its predecessor in A . As arenas are forestial, any move has at most one pointer and only minimal events have none.

It is worth, just this once, being extremely pedantic about the representation used in Figure 16 and others. Recall that $\llbracket (\mathbb{U} \rightarrow \mathbb{U}) \rightarrow \mathbb{B} \rrbracket = !(\mathbb{U} \multimap \mathbb{U}) \multimap \mathbb{B}$. Accordingly,

$$\llbracket (\mathbb{U} \rightarrow \mathbb{U}) \rightarrow \mathbb{B} \rrbracket = \mathbb{N} \times (\mathbb{N} \times |\mathbb{U}| + |\mathbb{U}|) + |\mathbb{B}|$$

with $+$ the tagged disjoint union $A + B = \{1\} \times A \cup \{2\} \times B$, previously also written \parallel . So an event of $\llbracket (\mathbb{U} \rightarrow \mathbb{U}) \rightarrow \mathbb{B} \rrbracket$ carries a *move* from \mathbb{U} or \mathbb{B} , *tags* originating from the disjoint unions and indicating one type component, and natural numbers, the *copy indices*. In Figure 16 the information of the *moves* is conveyed by the label, *i.e.* $\mathbf{q}^-, \checkmark^+$, *etc.* The *tag* is conveyed by the position of the move under the corresponding type component. Finally, the *copy indices* are given as a sequence in grey, with the leftmost integer corresponding to the outermost $!$. For instance, the move $\mathbf{q}_{0,1}^-$ really stands for $(1, (0, (1, (1, \mathbf{q}^-)))$.

It is often convenient to display pointers, but they are not part of the structure of plays. If they are imported into plays, then copy indices become essentially disposable (up to \cong). To make this formal, we start by defining a notion of *plays with pointers* on a $--$ -arena.

Definition 2.18. An **alternating play with pointers** on A is $s_1 \dots s_n \in |A|^*$ which is:

$$\textit{alternating: } \forall 1 \leq i \leq n-1, \text{pol}_A(s_i) \neq \text{pol}_A(s_{i+1}),$$

together with, for all $1 \leq j \leq n$ s.t. s_j is non-minimal in A , the data of a **pointer** to some earlier s_i s.t. $s_i \rightarrow_A s_j$. We write $\mathcal{P}\text{-}\downarrow\uparrow\text{-Plays}(A)$ for the set of plays with pointers on A .

The *non-repetitive* condition of Definition 2.7 would make pointers redundant as each move has a unique predecessor, and the existence of pointers would boil down to the fact that plays reach only down-closed sets of events. It is a useful exercise to show that non-repetitive plays with pointers are in bijection with alternating plays, on arenas without conflict.

Reciprocally, since repetitions are now allowed, we may use them to represent executions with replication even without the expansion process of Section 2.2.

2.4.2. *Meager and concrete arenas.* Definition 2.18 applies to arenas in the sense of Section 2.2.1, but it ignores part of their structure: it takes no account of *conflict*, and *symmetry*. Indeed, plays with pointers originate from HO games, where arenas are much simpler:

Definition 2.19. A **meager arena** is a partial order with polarities $(A, \leq_A, \text{pol}_A)$ s.t.:

$$\begin{aligned} \textit{alternating: } & \text{if } a_1 \rightarrow_A a_2, \text{pol}_A(a_1) \neq \text{pol}_A(a_2), \\ \textit{forestial: } & \text{if } a_1 \leq_A a \text{ and } a_2 \leq_A a, \text{ then } a_1 \leq_A a_2 \text{ or } a_2 \leq_A a_1, \end{aligned}$$

without conflict or symmetry. A **meager $--$ -arena** additionally satisfies:

$$\textit{negative: } \text{if } a \in \min(A), \text{ then } \text{pol}_A(a) = -.$$

Clearly, Definition 2.18 applies to meager arenas. Each PCF type A may be interpreted as a meager arena $[A]$, setting $[\mathbb{U}] = \mathbb{U}$, $[\mathbb{B}] = \mathbb{B}$, $[\mathbb{N}] = \mathbb{N}$ and $[A \rightarrow B] = [A] \multimap [B]$; *i.e.* as for $\llbracket - \rrbracket$ but without the $!$ – this is exactly the interpretation in [HO00]. The arena $\llbracket A \rrbracket$ is then an *expansion* of $[A]$ – the notion of *concrete arena* makes this explicit:

Definition 2.20. A **concrete arena** is (A, A^0, lbl) with A an arena, A^0 a *meager arena*,

$$\text{lbl} : |A| \rightarrow |A^0|,$$

a **label** function, together satisfying the following additional requirements:

- locally pointed*: for all $x \in \mathcal{C}(A)$, x has at most one minimal event of each polarity,
- rigid*: lbl preserves minimality, and preserves immediate causality \rightarrow ,
- transparent*: for any $x, y \in \mathcal{C}(A)$ and bijection $\theta : x \simeq y$,
then $\theta \in \mathcal{S}(A)$ iff θ is an order-iso preserving lbl .

We shall update this in Section 7.3.1, when further structure becomes required. *Locally pointed* is phrased so as to allow non-negative arenas of the form $A^\perp \parallel B$. In most cases, for negative arenas, configurations $x \in \mathcal{C}(A)$ will have at most one minimal event.

Every basic arena \mathbb{X} may be regarded as the concrete arena $(\mathbb{X}, \mathbb{X}, \text{lbl}_{\mathbb{X}})$ with $\text{lbl}_{\mathbb{X}}$ the identity function. Concrete arenas support the arena constructions $\&$ and \rightarrow with $(A \& B)^0 = A^0 \otimes B^0$, and $(A \rightarrow B)^0 = A^0 \multimap B^0$. By induction, for every type A this gives us $(\llbracket A \rrbracket, \llbracket A \rrbracket, \text{lbl}_A)$, a pointed concrete $--$ -arena with lbl_A simply forgetting all copy indices.

Remark 2.21. *Transparent* makes explicit the nature of symmetries on arenas arising from types: as they leave all components unchanged except *copy indices*, they are exactly all reindexings. This does not always hold outside the types considered here. In particular, concrete arenas do not support \otimes : of course condition *locally pointed* fails, but more fundamentally, valid symmetries in $!(A \otimes B)$ must send $(i, (1, a))$ and $(i, (2, b))$ to the same copy index j , a non-local constraint, not reflected by condition *transparent*. This is why we do not consider all arenas to be *concrete*: they fail to cover the full Seelye category structure.

In the sequel, we only assume arenas to be concrete when it is explicitly mentioned.

2.4.3. *Pointers and symmetry.* Plays with pointers represent plays up to symmetry:

Proposition 2.22. *Consider A a concrete arena. Then, there is a function*

$$\mathcal{P} : \downarrow\uparrow\text{-Plays}(A)/\cong \rightarrow \mathcal{P}\text{-}\downarrow\uparrow\text{-Plays}(A^0),$$


injective and preserving length and prefix.

Proof. For $s \in \downarrow\uparrow\text{-Plays}(A)$, we first construct $s^\rightarrow \in \mathcal{P}\text{-}\downarrow\uparrow\text{-Plays}(A)$ by importing \rightarrow_A . Then, $\mathcal{P}(s)$ is obtained by applying lbl_A pointwise. That pointers on $\mathcal{P}(s)$ are well-formed (*i.e.* that if s_j points to s_i , then $s_i \rightarrow_{A^0} s_j$) follows from lbl preserving minimality and the immediate causal order. That \mathcal{P} is invariant under \cong boils down to *transparent*. By construction, \mathcal{P} preserves length and prefix. For injectivity, take $s, s' \in \downarrow\uparrow\text{-Plays}(A)$ such that $\mathcal{P}(s) = \mathcal{P}(s')$. Since \mathcal{P} is length-preserving, s and s' have the same length n . Consider

$$\theta = \{(s_i, s'_i) \mid 1 \leq i \leq n\} : |s| \simeq |s'|$$

the induced bijection. Since $\mathcal{P}(s) = \mathcal{P}(s')$, in particular s and s' have the same pointers, so θ is an order-isomorphism, and moreover since $\mathcal{P}(s) = \mathcal{P}(s')$ again we also have $\text{lbl}_A(s_i) = \text{lbl}_A(s'_i)$ for all $1 \leq i \leq n$. Hence, θ is a symmetry, so by *transparent*, $s \cong s'$ as required. \square

However, \mathcal{P} is not surjective. Writing $A = \mathbb{U} \rightarrow \mathbb{B}$, the play $s \in \mathcal{P}\text{-}\downarrow\uparrow\text{-Plays}(\llbracket A \rrbracket)$ set as



is not the image of any play in $\downarrow\uparrow\text{-Plays}(\llbracket A \rrbracket)$, for two reasons: (1) not every move is duplicated in $\llbracket A \rrbracket$, e.g. there is only one copy of \checkmark^- for every copy of \mathbf{q}^+ – this linearity discipline is enforced by the *non-repetitive* condition, which is absent in $\mathcal{P}\text{-}\downarrow\uparrow\text{-Plays}(\llbracket A \rrbracket)$; and (2) likewise, $\llbracket A \rrbracket$ and $\mathcal{P}\text{-}\downarrow\uparrow\text{-Plays}(\llbracket A \rrbracket)$ do not account for *conflict* between \mathbf{tt}^+ and \mathbf{ff}^+ in $\llbracket A \rrbracket$.

2.4.4. HO strategies. This extends to *strategies*. For concrete arena A and $\sigma : A$, then $\mathcal{P}(\sigma) = \{\mathcal{P}(s) \mid s \in \sigma\}$ is a strategy on A^0 in the Hyland-Ong sense, i.e. a prefix-closed, deterministic set of plays with pointers closed under Opponent extensions. We have:

Proposition 2.23. *Consider A a concrete arena, and prestrategies $\sigma, \tau : \llbracket A \rrbracket$. Then, $\sigma \approx \tau$ iff $\mathcal{P}(\sigma) = \mathcal{P}(\tau)$.*

Proof. If. Consider $\sigma, \tau : A$ s.t. $\mathcal{P}(\sigma) = \mathcal{P}(\tau)$. For $\sigma \approx \tau$ we first check \rightarrow -simulation. Consider $sa^+ \in \sigma, t \in \tau$ s.t. $s \cong_A t$. But $\mathcal{P}(sa) \in \mathcal{P}(\tau)$, so there is $t'b' \in \tau$ s.t. $\mathcal{P}(t'b') = \mathcal{P}(sa)$. Hence by Proposition 2.22, $sa \cong_A t'b'$. So $t, t' \in \tau$ and $t \cong_A t'$, with $t'b' \in \tau$. By uniformity of τ , $tb \in \tau$ for some b with $t'b' \cong_A tb$, so $tb \cong_A sa$ as well. The condition \leftarrow -simulation is symmetric. For \rightarrow -receptive, assume $sa^- \in \sigma, t \in \tau$ and $sa^- \cong_A tb^-$. Since $\mathcal{P}(sa^-) \in \tau$, there is $t'b' \in \tau$ s.t. $\mathcal{P}(sa) = \mathcal{P}(t'b')$, i.e. $sa \cong_A tb$. But then $t'b' \in \tau$ and $t'b' \cong_A tb$, so by uniformity of τ we have $tb \in \tau$. Finally, \leftarrow -receptive is symmetric.


Only if. Consider $\sigma, \sigma' : A$ s.t. $\sigma \approx \sigma'$, and take $\mathcal{P}(s) \in \mathcal{P}(\sigma)$ for some $s \in \sigma$. By induction on s , we build some $s' \in \sigma'$ s.t. $s \cong_A s'$: for positive extensions this follows from $\sigma \approx \sigma'$; for negative extensions from the *extension* condition on isomorphism families and the \rightarrow -receptive condition on uniformity. But then by Proposition 2.22 we have $\mathcal{P}(s) = \mathcal{P}(s')$, so $\mathcal{P}(\sigma) \subseteq \mathcal{P}(\sigma')$. The argument is symmetric, so $\mathcal{P}(\sigma) = \mathcal{P}(\sigma')$ as desired. \square

Plays with pointers permit a presentation of strategies up to \approx , avoiding copy indices. They provide the foundation for HO games [HO00], where the interpretation of types is essentially $[-]$ (without conflict), and plays carry pointers. We include the classical example showing that though one may choose copy indices or pointers, one cannot avoid both.

Example 2.24. The *Kierstead terms* $\vdash K_x, K_y : ((\mathbb{B} \rightarrow \mathbb{B}) \rightarrow \mathbb{B}) \rightarrow \mathbb{B}$ are defined as

$$K_x = \lambda F^{(\mathbb{B} \rightarrow \mathbb{B}) \rightarrow \mathbb{B}}. F(\lambda x. F(\lambda y. x)), \quad K_y = \lambda F^{(\mathbb{B} \rightarrow \mathbb{B}) \rightarrow \mathbb{B}}. F(\lambda x. F(\lambda y. y)).$$

Their respective interpretations in $\downarrow\uparrow\text{-Strat}_1$ have distinctive plays:



Here pointers are redundant, and computed from the identity of moves. In particular, in both plays the \mathbf{q}^+ “points to” the unique \mathbf{q}^- with compatible copy indices. Mapping these through \mathcal{P} , we get two plays with pointers that only differ through their pointers. In HO games, the Kierstead terms are only distinguished by pointers⁶. It is crucial to keep them separate: it is a surprisingly challenging exercise to find a PCF context that separates them.

⁶It is necessary to go up to third-order types to find such examples. Pointers are redundant up to second-order types, which is the starting point of *algorithmic game semantics* [GM03].




Figure 17: Illustration of the operational contents of game semantics

Plays with pointers are powerful, and indeed the game semantics literature is strongly biased towards HO games (as opposed to AJM games). This, however, has two costs. Firstly, plays with pointers are not a natural inductive structure, making their manipulation sometimes inelegant or unwieldy (so-called “pointer surgery”). Propositions have been made for clean formalizations, *e.g.* through nominal sets [GG12]. Another cost is that replication is so hard-wired into the model that it does not enjoy a clean linear decomposition. Enforcing linearity is slightly awkward and relies on additional structure [McC98].

In this work we stick with $\downarrow\uparrow$ -Strat rather than adopting plays with explicit pointers. Among other things this will ease the relationship with the forthcoming *thin concurrent games*, which we do not know how to formulate with pointers in general. Besides, in $\downarrow\uparrow$ -Strat, pointers *can* be directly obtained from the arena, and as such may be used as in HO games⁷. In fact, pointers play a central role in this paper. From now on, all representations of plays will display pointers. In contrast, we will often omit copy indices as most of the time they convey no useful information; one can regard this convention as drawing $\mathcal{P}(s)$ rather than s .

3. SEQUENTIAL COMPUTATIONAL EFFECTS IN GAME SEMANTICS

We now explore the model constructed above, introducing the traditional “semantic cube”.

The plays of a term are computed denotationally, by induction on syntax. However, given a term, an experienced game semanticist will be able to directly list its plays, without going through the intricate definition of the interpretation. This is because as discussed before, plays represent the operational behaviour of the term: rather than denotationally, they can be obtained directly from the term by operational means [DHR96, Jab15, GT12, LS14]. This is illustrated in Figure 17. Opponent moves trigger the evaluation of a subterm, which appears boxed. The following Player move then corresponds to the head (*i.e.* leftmost)


⁷Another work blurring the lines between HO and AJM is [AJ09] where AJM games are equipped with a *function* able to rebuild pointers without the need to explicitly integrate them in plays. All the data of a game in the sense of [AJ09] can be computed from an arena in our sense, but our arenas are more primitive.

variable occurrence (or constant) of the subterm being evaluated. The pointers from Player moves correspond to the stage where the variable in head position was abstracted, or to the function call being returned by the value in head position. More specifically, Figure 17 represents the interaction of the term under study with the applicative context:

$$C[] = [] (\lambda x^{\mathbb{B}}. \lambda y^{\mathbb{B}}. \mathbf{if} \ x (\mathbf{if} \ y \ \mathbf{skip} \ \mathbf{skip}) \ \mathbf{skip})$$

Figure 17 is strongly inspired by the *Pointer Abstract Machine (PAM)* [DHR96].

3.1. Well-Bracketing. Now that *executions as plays* are first-class citizens, independent of programs, we may start classifying them according to the computational capabilities that they witness. For instance, is the following play a possible execution of a term?



We argue informally why this cannot be an execution in PCF. The first action of the term is to ask its argument, so it has the form $\lambda f^{\mathbb{U} \rightarrow \mathbb{U}}. f M$; we annotate the figure with the corresponding operational state as in Figure 17. In the last line, \mathbf{tt} at toplevel indicates the overall computation has terminated to \mathbf{tt} . This is confusing, since operationally the Opponent move in the third line corresponded to triggering the evaluation of the argument of f . How can evaluating the argument of f cause the whole computation to terminate?

Nevertheless, this play is indeed a realistic execution, for the term

$$\lambda f^{\mathbb{U} \rightarrow \mathbb{U}}. \mathbf{callcc} (\lambda k^{\mathbb{B} \rightarrow \mathbb{U}}. f (k \ \mathbf{tt}); \perp) : (\mathbb{U} \rightarrow \mathbb{U}) \rightarrow \mathbb{B}$$

where **callcc** is the *call-with-current-continuation* primitive originating in *Scheme*, and which famously may be typed with Peirce’s law [Gri90]. The precise operational semantics of **callcc** will not be useful for this paper, but informally **callcc** M immediately calls M , feeding it a special function k , the “continuation”. When the continuation is called with value v , **callcc** interrupts M and returns v at toplevel, breaking the call stack discipline.

Can the play above be realised without **callcc** (or some other *control operator*, as such primitives are called)? We can show that the answer is no, by capturing plays that “respect the call stack discipline”, and refining the whole interpretation to show that this invariant is preserved. This is the goal of the notion of *well-bracketing*. First we enrich arenas:

Definition 3.1. A **Question/Answer labeling** on arena A is a function

$$\lambda_A : |A| \rightarrow \{\mathcal{Q}, \mathcal{A}\}$$

invariant under symmetry (if $\theta : x \cong_A y$, then for all $a \in x$, $\lambda_A(a) = \lambda_A(\theta(a))$) and satisfying:

- question-opening:* if $a \in |A|$ is minimal, then $\lambda_A(a) = \mathcal{Q}$,
- answer-closing:* if $\lambda_A(a) = \mathcal{A}$, then a is maximal for \leq_A ,
- answer-linear:* if $\lambda_A(a_1) = \lambda_A(a_2) = \mathcal{A}$ with $a \rightarrow_A a_1, a_2$, then $a_1 = a_2$ or $a_1 \#_A a_2$.

From now on, **arenas** have a Question/Answer labeling. *Questions* intuitively correspond to variable calls, while *Answers* correspond to returns. Basic arenas are enriched as shown in Figure 18. For other constructions the labeling is inherited transparently, with $\lambda_{!A}(i, a) = \lambda_A(a)$, $\lambda_{A_1 \otimes A_2}(i, a) = \lambda_{A_i}(a)$, $\lambda_{A \multimap B}(2, b) = \lambda_B(b)$, and $\lambda_{A \multimap B}(1, (b, a)) = \lambda_A(a)$.




Figure 18: Question/Answer labeling on basic arenas

If $s \in \downarrow\uparrow\text{-Plays}(A)$ and s_i is an answer, it cannot be minimal in A by *question-opening*. Its antecedent in A – its *justifier* – must appear in s as some s_j with $j < i$, and is a question by *answer-closing*. We say that s_i **answers** s_j . If a question in s has an answer in s we say it is **answered** in s . The last unanswered question of s , if any, is the **pending question**.

We now capture executions respecting the call stack discipline as *well-bracketed* plays.

Definition 3.2. Let $s \in \downarrow\uparrow\text{-Plays}(A)$ be an alternating play.


It is **well-bracketed** if for all prefix $ta^A \sqsubseteq s$, a answers the pending question of t .

All plays encountered in the paper until now are well-bracketed, with the exception of the example at the beginning of Section 3.1. We can then define well-bracketed strategies:

Definition 3.3. Let $\sigma : A$ be a strategy on A .


It is **well-bracketed** iff for all $sa^+ \in \sigma$, if s is well-bracketed then sa is well-bracketed.

In other words, a well-bracketed strategy is never the first to break the call stack discipline. Asking all plays to be well-bracketed is too strict, as illustrated by the play



of copycat: the last move does not answer the pending question, but because Opponent broke the normal control flow first. There is a luf subcategory $\downarrow\uparrow\text{-Strat}^{\text{wb}}$ of $\downarrow\uparrow\text{-Strat}$, having well-bracketed strategies as morphisms. The interpretation of PCF in $\downarrow\uparrow\text{-Strat}_!$ in fact yields only well-bracketed strategies, *i.e.* has target $\downarrow\uparrow\text{-Strat}_!^{\text{wb}}$. This shows that indeed, the execution at the beginning of Section 3.1 cannot be realised in PCF.

3.2. Visibility and Innocence. Likewise, is this play a possible execution of a term?



Again, this seems unfeasible in PCF. Again, on the right hand side we show, assuming a term realising this play, its corresponding operational states. At the third and fifth moves, the *same* subterm is being evaluated; yet we get two distinct answers. In an extension




Figure 19: A strategy with references




Figure 20: Non P-visible play

of PCF with a primitive $+$ for non-deterministic choice, this play would be realisable by $\lambda f^{\mathbb{B} \rightarrow \mathbb{U}}. f(\mathbf{tt} + \mathbf{ff})$. But does it make computational sense in a *deterministic* language?

Once more, the answer is yes: the play above describes a valid execution of the term

$$\lambda f^{\mathbb{B} \rightarrow \mathbb{U}}. \mathbf{newref} \ r \ \mathbf{in} \ f(\mathbf{let} \ x \ =!r \ \mathbf{in} \ r := 1; (x > 0)) : (\mathbb{B} \rightarrow \mathbb{U}) \rightarrow \mathbb{U}$$

in PCF extended with references: $\mathbf{newref} \ r \ \mathbf{in} \ M$ allocates a reference r initialized to 0. We show in Figure 19 an operational description as to how this term indeed realises this play.

Again, this cannot be realised in PCF. To show this, we give a version of *innocence* [HO00], formalizing that without state, evaluating the same subterm yields the same response. The first step is a mathematical way to state that two plays “correspond to the same subterm”, like the two prefixes of the play of Figure 19 terminating with a \mathbf{q}^- on the left.

The operation computing (a mathematical notion of) “current subterm” is the *P-view*:

Definition 3.4. Let $s \in \downarrow\uparrow\text{-Plays}(A)$. Its **P-view** is the subsequence defined by induction:

$$\begin{aligned} \ulcorner \varepsilon \urcorner &= \varepsilon \\ \ulcorner sa^+ \urcorner &= \ulcorner s \urcorner a^+ \\ \ulcorner sa_1^+ s' a_2^- \urcorner &= \ulcorner s \urcorner a_1^+ a_2^- \quad \text{if } a_1^+ \rightarrow_A a_2^- \\ \ulcorner sa_1^+ s' a_2^- \urcorner &= a_2^- \quad \text{if } a_2^- \text{ is negative minimal in } A \end{aligned}$$

We take the immediate prefix for P -ending plays and follow the pointer for O -ending plays. For instance, the prefixes of length 3 and 5 of the play on Figure 19 have the same P-view, capturing that they correspond to the same subterm. This is a powerful definition – really, the distinguishing feature of HO games [HO00] – and it often takes newcomers a while to digest. Interestingly, our forthcoming *parallel innocence* will be phrased quite differently.

But this is not yet conclusive: if $s \in \downarrow\uparrow\text{-Plays}(A)$, it might be that $\ulcorner s \urcorner \notin \downarrow\uparrow\text{-Plays}(A)$. For instance, in Figure 20 we gray out moves not selected in computing the P-view of $s \in \downarrow\uparrow\text{-Plays}(A)$ for $A = [(\mathbb{U} \rightarrow \mathbb{U}) \rightarrow \mathbb{U}]$. The subsequence of $\ulcorner s \urcorner$ in black is an alternating sequence of $|A|$, but fails *valid* of Definition 2.7. Indeed, the “justifier” of \checkmark^+ , its immediate dependency in A , is not selected – thus $\ulcorner s \urcorner$ is not down-closed. Accordingly, we say:

Definition 3.5. A play $e \in \downarrow\uparrow\text{-Plays}(A)$ is **P-visible** if for all prefix $\forall t \sqsubseteq e$, $\ulcorner t \urcorner \in \downarrow\uparrow\text{-Plays}(A)$.

Likewise, a strategy $\sigma : A$ is **P-visible** iff all its plays are P-visible.

So, “computing P-views never drops pointers”, or “Player always points in the P-view”. On P-visible $s \in \downarrow\uparrow\text{-Plays}(A)$, the P-view always yields a well-formed (P-visible) play.

We now define *innocent strategies* as those that behave the same in any situation where the same subterm is being evaluated, *i.e.* whose behaviour only depends on the *P-view*:

Definition 3.6. A P -visible alternating strategy $\sigma : A$ is **innocent** if it satisfies:

innocence: for all $sa^+ \in \sigma$, for all $t \in \sigma$, if $\ulcorner s \urcorner = \ulcorner t \urcorner$ then $ta^+ \in \sigma$.

That $ta^+ \in \downarrow\uparrow\text{-Plays}(A)$ is well-formed relies on $\ulcorner s \urcorner = \ulcorner t \urcorner$, so that the causal dependencies of a in A appear in t ⁸. All structural morphisms of $\downarrow\uparrow\text{-Strat}$ are innocent. Innocent strategies compose – though this is infamously tricky to prove, prompting a significant line of work investigating the structures arising from the composition of innocent strategies [Cur98, HHM07, CD15]. We do not review here the proof of stability under composition.

The interpretation of PCF yields only innocent strategies, *i.e.* targets the cartesian closed lluf subcategory $\downarrow\uparrow\text{-Strat}_!^{\text{inn}}$ of innocent strategies. Hence, the play at the beginning of Figure 3.2 cannot be realised in PCF. We also get a cartesian closed lluf subcategory $\downarrow\uparrow\text{-Strat}_!^{\text{wb,inn}}$ with well-bracketing. Finally, the weaker P -visibility is also preserved under the categorical operations, forming lluf sub-cartesian closed categories $\downarrow\uparrow\text{-Strat}_!^{\text{vis}}$ and $\downarrow\uparrow\text{-Strat}_!^{\text{wb,vis}}$.

3.3. Full Abstraction for PCF. We have now eliminated all non PCF-definable behaviour. We review the corresponding *definability* and *intensional full abstraction* arguments.

3.3.1. Definability. Call a **P-view** on arena A any $s \in \downarrow\uparrow\text{-Plays}(A)$ invariant under P-view, *i.e.* $\ulcorner s \urcorner = s$ – those are exactly the $s \in \downarrow\uparrow\text{-Plays}(A)$ such that for all $ts_i^+ s_{i+1}^- \sqsubseteq s$, we have $s_i \rightarrow_A s_{i+1}$, in other words *Opponent always points to the previous move*. We motivated P-views as a way to address specific “subterms” of a strategy – it might therefore not be a surprise that those are the key to reconstruct a term from an innocent strategy. We write

$$\ulcorner \sigma \urcorner = \{ \ulcorner s \urcorner \mid s \in \sigma \}$$

for the set of P-views of σ . If σ is innocent, then it is simple that $\ulcorner \sigma \urcorner \subseteq \sigma$. Moreover, σ can then be recovered as the set of P-visible $s \in \downarrow\uparrow\text{-Plays}(A)$ such that for all $t \sqsubseteq s$, $\ulcorner t \urcorner \in \ulcorner \sigma \urcorner$.

For $\sigma : A$ innocent, $\ulcorner \sigma \urcorner$ is not a strategy as in general it fails receptivity. It is however easily verified to be a *prestrategy* – and in particular *uniform*. Moreover, we have:

Proposition 3.7. For $\sigma, \tau : A$ innocent strategies on A , we have $\ulcorner \sigma \urcorner = \ulcorner \tau \urcorner$ iff $\sigma = \tau$.

Likewise, $\ulcorner \sigma \urcorner \approx \ulcorner \tau \urcorner$ if and only if $\sigma \approx \tau$.

Proof. We only detail the second statement. Firstly, if $\sigma \approx \tau$, it is direct that $\ulcorner \sigma \urcorner \approx \ulcorner \tau \urcorner$ as $\ulcorner \sigma \urcorner \subseteq \sigma$ and $\ulcorner \tau \urcorner \subseteq \tau$ and the bisimulation game of Definition 2.13 preserves P-views.

If $\ulcorner \sigma \urcorner \approx \ulcorner \tau \urcorner$, take $sa^+ \in \sigma, t \in \tau$ s.t. $s \cong_A t$. In particular $\ulcorner s \urcorner \cong_A \ulcorner t \urcorner$ and $\ulcorner s \urcorner a^+ \in \ulcorner \sigma \urcorner$. By \rightarrow -extension, there is b^+ s.t. $\ulcorner t \urcorner b^+ \in \ulcorner \tau \urcorner$, so $tb^+ \in \tau$ by innocence. This proves \rightarrow -extension, \leftarrow -extension is symmetric and \rightarrow, \leftarrow -receptivity follow by receptivity of σ, τ . \square

So innocent strategies have two representations: a full $\sigma : A$ satisfying Definition 3.6; or, following Proposition 3.7, the set $\ulcorner \sigma \urcorner$. Anticipating on later developments, we refer to $\ulcorner \sigma \urcorner$ as the *causal* presentation of σ . In traditional innocent game semantics, the forest of P-views is called (notably by Curien [Cur06]) the *meager* representation, while the set of plays is *fat*. Here this is misleading, because plays in $\ulcorner \sigma \urcorner$ still carry explicit copy indices. In particular $\ulcorner \sigma \urcorner$ has branches matching all copyable Opponent moves, which is “fat”.

To recover the meager representation, we show:

⁸In traditional Hyland-Ong games based on plays with points, one would conclude the above definition with something like “... then $ta \in \sigma$, where a has the same pointer as in sa ”, which is rarely made very formal. Here, because pointers are derived the above definition is rigorous and self-contained.

Proposition 3.8. *Consider A a concrete arena and $\sigma, \tau : A$ innocent strategies.*

If $\mathcal{P}(\ulcorner\sigma\urcorner) = \mathcal{P}(\ulcorner\tau\urcorner)$, then $\sigma \approx \tau$.

Proof. Let $\sigma, \tau : \llbracket A \rrbracket$ be innocent strategies on A and assume that $\mathcal{P}(\ulcorner\sigma\urcorner) = \mathcal{P}(\ulcorner\tau\urcorner)$.

By Proposition 2.23, $\ulcorner\sigma\urcorner \approx \ulcorner\tau\urcorner$. Then, by Proposition 3.7, it follows that $\sigma \approx \tau$. \square

This, at last, provides the *meager* representation.

These representations have distinct advantages: composition is only directly defined on the fat representation; but it is the meager one that bridges innocent strategies and syntax and allows definability. An innocent alternating strategy $\sigma : A$ is **finite** iff $\mathcal{P}(\ulcorner\sigma\urcorner)$ is finite. Its **size** is simply the cardinal of that set. Definability simply follows the meager form:

Theorem 3.9. *Let A be a PCF type, and $\sigma : \llbracket A \rrbracket$ be a finite well-bracketed innocent strategy.*


Then, there is a PCF term $\vdash M : A$ s.t. $\llbracket M \rrbracket \approx \sigma$.

Proof. We describe the argument – for more details, the reader is referred to [HO00].

Without loss of generality, A has the form $A_1 \rightarrow \dots \rightarrow A_n \rightarrow \mathbb{X}$ where for each $1 \leq i \leq n$,

$$A_i = A_{i,1} \rightarrow \dots \rightarrow A_{i,p_i} \rightarrow \mathbb{X}_i.$$

We reason on $\ulcorner\sigma\urcorner$, by induction on the size of σ . If σ has no reaction to the (unique) minimal \mathbf{q}^- in \mathbb{X} (i.e. $\sigma = \{\varepsilon, \mathbf{q}^-\}$), any diverging term will do. Otherwise, by determinism there is exactly one move a^+ s.t. $\mathbf{q}^- a^+ \in \sigma$. If a^+ is an answer v^+ on \mathbb{X} , then M is the matching constant. Otherwise, a^+ is the initial $\mathbf{q}_{i_0}^+$ in some A_{i_0} ⁹. The situation is drawn as



with, in grey, the possible extended P-views. For each extension there is a residual substrategy.

We extract those – first, if $\mathbf{q}_{i_0}^+$ immediately returns. For each value v in \mathbb{X}_i , we form

$$\ulcorner\sigma_v\urcorner = \{\mathbf{q}^- s \mid \mathbf{q}^- \mathbf{q}_{i_0}^+ v^- s \in \ulcorner\sigma\urcorner\},$$

a causal innocent strategy on $\llbracket A \rrbracket$ of size strictly lesser than σ . By induction hypothesis there is $\vdash M_v : A$ with $\llbracket M_v \rrbracket \approx \sigma_v$. As σ is finite, there are finite many v s.t. σ_v is non-diverging.

Alternatively, for all $1 \leq j \leq p_{i_0}$, we consider P-views $\mathbf{q}^- \mathbf{q}_{i_0}^+ \mathbf{q}_{i_0,j}^- s \in \ulcorner\sigma\urcorner$ where as a P-view, s answers neither $\mathbf{q}_{i_0}^+$, nor \mathbf{q}^- by well-bracketing. Such a P-view yields

$$\mathbf{q}_{i_0,j}^- s \in \downarrow\uparrow\text{-Plays}(\llbracket A_1 \rightarrow \dots \rightarrow A_n \rightarrow A_{i_0,j} \rrbracket)$$

a P-view where moves in s formerly depending on \mathbf{q}^- in $\llbracket A \rrbracket$ are set to depend on $\mathbf{q}_{i_0,j}^-$. Considering all such P-views generates a causal innocent strategy of size strictly lesser than σ , hence by induction hypothesis there is $\vdash M_{i_0,j} : A_1 \rightarrow \dots \rightarrow A_n \rightarrow A_{i_0,j}$ s.t. $\llbracket M_{i_0,j} \rrbracket \approx \sigma_{i_0,j}$.

Finally, with all this data we may form $\vdash M : A$ as

$$\begin{aligned} & \lambda x_1^{A_1} \dots x_n^{A_n}. \text{case } x_{i_0} (M_{i_0,1} x_1 \dots x_n) \dots (M_{i_0,p_{i_0}} x_1 \dots x_n) \text{ of} \\ & \quad v_1 \mapsto M_{v_1} \\ & \quad \dots \\ & \quad v_p \mapsto M_{v_p} \end{aligned}$$

where p is such that every σ_{v_i} with $i > p$ is diverging. We get, as needed $\llbracket M \rrbracket \approx \sigma$. \square

⁹Here the subscripts indicate the type component and not copy indices, which are left un-specified.

The final statement is a careful verification following the definition of the interpretation, see [HO00]. Here, **case** is the syntax introduced in Section 1.3, involving the **let** construct. Without that, simply iterating **if** constructs would yield a strategy that re-computes

$$x_{i_0} (M_{i_0,1} x_1 \dots x_n) \dots (M_{i_0,p_{i_0}} x_1 \dots x_n)$$

each time it matches it against a value. This is what is done in [HO00] as their version of PCF does not include a **let** construct. This yields a term that is not \approx -equivalent to σ , but is nonetheless \sim -equivalent (see Section 1.6), which suffices for full abstraction. We prefer the present more intensional definability result, and hence have included the **let** construct¹⁰.

3.3.2. Intensional full abstraction. Full abstraction of a denotational model with respect to a language was defined in Section 1.6. Of course, $\downarrow\uparrow\text{-Strat}_!$ is not fully abstract for PCF as it stands. For instance, $\llbracket \lambda x^{\mathbb{U}}. x; x \rrbracket \not\approx \llbracket \lambda x^{\mathbb{U}}. x \rrbracket$: game semantics displays explicitly individual calls to x , so we see that the term on the left hand side evaluates x twice whereas the other evaluates it once. However, we do of course have $\lambda x^{\mathbb{U}}. x; x \sim \lambda x^{\mathbb{U}}. x$; this can for instance be deduced from them having the same interpretation in Scott domains [Plo77].

The celebrated “full abstraction for PCF” results are in fact what (following [AJM00]) we call *intensional full abstraction*. Fixing an interpretation $\llbracket - \rrbracket$ of PCF into a \mathcal{C} , we set

$$f \sim g \quad \Leftrightarrow \quad \forall \alpha \in \mathcal{C}(A \rightarrow B, \llbracket \mathbb{U} \rrbracket), \quad (\alpha \circ \bar{f} = \alpha \circ \bar{g}),$$

for $f, g \in \mathcal{C}(A, B)$, with $\bar{f}, \bar{g} \in \mathcal{C}(\mathbf{1}, A \rightarrow B)$ obtained via cartesian closure, and $\mathbf{1}$ the terminal object of \mathcal{C} . This mimics the definition of observational equivalence. We say that \mathcal{C} is **intensionally fully abstract** for PCF iff the quotiented model \mathcal{C}^{\sim} is fully abstract.

Theorem 3.10. *The model $\downarrow\uparrow\text{-Strat}_!^{\text{wb,inn}}$ is intensionally fully abstract for PCF.*




Proof. Consider $\vdash M, N : A$ s.t. $M \sim N$, and assume $\llbracket M \rrbracket \not\downarrow \llbracket N \rrbracket$, i.e. there is a test $\alpha \in \downarrow\uparrow\text{-Strat}_!^{\text{wb,inn}}(\llbracket A \rrbracket, \llbracket \mathbb{U} \rrbracket)$ s.t. $\alpha \odot! \llbracket M \rrbracket \neq \alpha \odot! \llbracket N \rrbracket$ – say *w.l.o.g.* that $\alpha \odot! \llbracket M \rrbracket \downarrow$ converges while $\alpha \odot! \llbracket N \rrbracket \uparrow$. One may prove (see [HO00] for details) that the corresponding interactions expose only a finite part of α , so *w.l.o.g.* we can assume α finite. By Theorem 3.9, α is defined via a PCF term, providing a context $C[-]$ s.t. $\llbracket C[M] \rrbracket = \alpha \odot! \llbracket M \rrbracket$ and $\llbracket C[N] \rrbracket = \alpha \odot! \llbracket N \rrbracket$. But then, we must have $C[M] \downarrow$ while $C[N]$ diverges by Proposition 2.16; contradiction. \square

Intensional full abstraction is full abstraction for an *a priori* non effective quotiented model: it does not directly provide effective tools to reason about observational equivalence. Instead, it is a way of stating that we have faithfully captured the intensional behaviour of programs, in the sense that the added tests in the model are not able to distinguish more – there is no “abstraction leak”. Often, it packages adequacy and finite definability.

Full abstraction is of course the preferred notion when the quotiented model is sufficiently effective and the interpretation computable (i.e. *effectively presentable* [Plo81]). But when it requires an undecidable quotient¹¹, we believe it preferable to use a different terminology: “*intensional full abstraction*” puts the emphasis on the model pre-quotient. In game semantics, it is that model pre-quotient that had the most impact. In particular it then led to *effective* fully abstract models for stateful languages, leveraging the results and insights above.

¹⁰An alternative is to include a primitive **case** evaluating its argument exactly once. The terms then obtained via definability are easily characterised syntactically – dubbed PCF *Böhm trees* by Curien, and are studied in [Cur98]. The definability process informs a concrete order-isomorphism between finite meager innocent strategies and finite PCF Böhm trees, emphasizing that meager innocent strategies *are* syntax.

¹¹For PCF this is unavoidable as observational equivalence is undecidable already for finitary PCF [Loa01].


 Figure 21: \mathbf{ref}_w

 Figure 22: \mathbf{ref}_r

 Figure 23: \mathbf{sem}

3.4. Full Abstraction for IA. The exposition in Section 3.2 suggests that also without innocence, strategies are computationally relevant for programs with mutable state. We now focus on the game semantics of IA, namely PCF extended with interference (see Section 1).

3.4.1. Interpretation of types. With respect to PCF, IA adds the type \mathbf{ref} of *integer references*, and the type \mathbf{sem} of *semaphores*. Their usual game semantic interpretation is *behavioural*, in the sense that it represents how one may interact on those types: one may read a reference or write a new value in it; and likewise one may grab a semaphore, or release it.

To capture this, we define $--$ arenas: $\mathbf{ref}_w = \&_{n \in \mathbb{N}} \mathbb{U}$, $\mathbf{ref}_r = \mathbb{N}$ and $\mathbf{sem} = \mathbb{U} \& \mathbb{U}$, and set $\llbracket \mathbf{ref} \rrbracket = \mathbf{ref}_w \& \mathbf{ref}_r$ and $\llbracket \mathbf{sem} \rrbracket = \mathbf{sem}$. Although we reuse the arena constructions for \mathbb{U} and \mathbb{N} , for specific moves in these arenas we use the naming conventions of Figures 21, 22 and 23 – in Figures 21 and 22 all distinct moves in the same row are in pairwise conflict.

3.4.2. Interacting with Memory and Semaphores. The idea behind Abramsky and McCusker’s interpretation of state is that it is not the operations of reading, writing, grabbing or releasing a semaphore that are effectful – indeed, those are just requests via the interface provided by the \mathbf{ref} and \mathbf{sem} types and associated commands. The strategy for a program with free reference or semaphore variables will simply record their accesses leaving the memory and semaphores uninterpreted. For instance, the strategy for $x : \mathbf{ref} \vdash x := 0; !x : \mathbb{N}$ includes:

$$\begin{array}{c} \mathbf{!ref} \quad \multimap \quad \mathbb{N} \\ \mathbf{q}^- \\ \mathbf{w0}^+ \\ \checkmark^- \\ \mathbf{r}^+ \\ \mathbf{42}^- \end{array}$$

a play where the value read is not the value just written. The actual effectful computation will be handled in Section 3.4.3 with the creation of new references and semaphores.

Accordingly, we set the interpretation of memory and semaphore accesses as:

$$\begin{aligned} \llbracket M := N \rrbracket &= \mathbf{assign} \odot! \langle \llbracket N \rrbracket, \llbracket M \rrbracket \rangle \\ \llbracket !M \rrbracket &= \mathbf{deref} \odot! \llbracket M \rrbracket \\ \llbracket \mathbf{grab}(M) \rrbracket &= \mathbf{grab} \odot! \llbracket M \rrbracket \\ \llbracket \mathbf{release}(M) \rrbracket &= \mathbf{release} \odot! \llbracket M \rrbracket \end{aligned}$$

using the (innocent well-bracketed) strategies of Figures 24, 25, 26, and 27. Finally, we set

$$\llbracket \mathbf{mkvar} M N \rrbracket = \langle \langle \llbracket M \rrbracket n \mid n \in \mathbb{N} \rangle, \llbracket N \rrbracket \rangle, \quad \llbracket \mathbf{mksem} M N \rrbracket = \langle \llbracket M \rrbracket, \llbracket N \rrbracket \rangle,$$

where $\llbracket M \rrbracket n$ is $\llbracket M \rrbracket$ applied to the constant strategy n (using the cartesian closed structure of $\downarrow\uparrow\text{-Strat}_!$), and using implicitly the isomorphisms $\mathbf{ref} \cong (\&_{n \in \mathbb{N}} \mathbb{U}) \& \mathbb{N}$ and $\mathbf{sem} \cong \mathbb{U} \& \mathbb{U}$.

Finite definability of finite innocent well-bracketed strategies still holds – the proof (see [AM96]) directly extends that of Theorem 3.9, using bad variables and semaphores.

$$\begin{array}{c} !(\mathbb{N} \ \& \ \mathbf{ref}) \multimap \mathbb{U} \\ \mathbf{q}_0^+ \\ n_0^- \\ \mathbf{wn}_1^+ \\ \checkmark_1^- \\ \checkmark^+ \end{array}$$

Figure 24: assign

$$\begin{array}{c} !\mathbf{ref} \multimap \mathbb{N} \\ \mathbf{q}^- \\ \mathbf{r}_0^+ \\ n_0^- \\ n^+ \end{array}$$

Figure 25: deref

$$\begin{array}{c} !\mathbf{sem} \multimap \mathbb{U} \\ \mathbf{q}^- \\ \mathbf{g}_0^+ \\ \checkmark_0^- \\ \checkmark^+ \end{array}$$

Figure 26: grab

$$\begin{array}{c} !\mathbf{sem} \multimap \mathbb{U} \\ \mathbf{q}^- \\ \mathbf{d}_0^+ \\ \checkmark_0^- \\ \checkmark^+ \end{array}$$

Figure 27: release

Proposition 3.11. *Consider A a type of \mathbf{IA} , and $\sigma : \llbracket A \rrbracket$ finite innocent well-bracketed. Then, there is a \mathbf{IA} term $\vdash M : A$ not using **newref** or **newsem**, such that $\llbracket M \rrbracket \approx \sigma$.*

3.4.3. Creation of References and Semaphores. Finally, we introduce the actual effectful behaviour. The idea is to use non-innocent $\mathbf{cell}_n : \mathbf{!ref}$, $\mathbf{lock}_n : \mathbf{!sem}$ implementing interference. For instance, \mathbf{cell}_n is a memory cell with n currently stored. When read it returns n , and upon a write request for $k \in \mathbb{N}$, it acknowledges it and proceeds as \mathbf{cell}_k . Likewise \mathbf{lock}_0 is the strategy for a free semaphore, and \mathbf{lock}_n for $n > 0$ represents a semaphore in use. Those may be simply described as the language of prefixes of the infinite trees:

$$\begin{aligned} \mathbf{cell}_n^I &= \mathbf{r}_i^- \cdot n_i^+ \cdot \mathbf{cell}_n^{I \uplus \{i\}} \mid \mathbf{wk}_i^- \cdot \checkmark_i^+ \cdot \mathbf{cell}_k^{I \uplus \{i\}} & (i \notin I) \\ \mathbf{lock}_0^I &= \mathbf{g}_i^- \cdot \checkmark_i^+ \cdot \mathbf{lock}_1^{I \uplus \{i\}} \mid \mathbf{d}_i^- & (i \notin I) \\ \mathbf{lock}_n^I &= \mathbf{g}_i^- \mid \mathbf{d}_i^- \cdot \checkmark_i^+ \cdot \mathbf{lock}_0^{I \uplus \{i\}} & (i \notin I, n > 0) \end{aligned}$$

where symbols are moves in **!ref** and **!sem** respectively, separated via \cdot for readability. Here, $I \subseteq_f \mathbb{N}$ collects the copy indices already used, ensuring *non-repetitive*. We set \mathbf{cell}_n as (the prefix language of) $\mathbf{cell}_n^\emptyset$ and \mathbf{lock}_n as (the prefix language of) $\mathbf{lock}_n^\emptyset$; it is direct that $\mathbf{cell}_n : \mathbf{!ref}$ and $\mathbf{lock}_n : \mathbf{!sem}$. However, they are not innocent. Considering the two plays:

$$\mathbf{r}_0^- \cdot 0_0^+, \quad \mathbf{w}1_1^- \cdot \checkmark_1^+ \cdot \mathbf{r}_0^- \cdot 1_0^+ \in \mathbf{cell}_0,$$

as $\mathbf{r}_0^- = \mathbf{w}1_1^- \cdot \checkmark_1^+ \cdot \mathbf{r}_0^-$, innocence requires $\mathbf{r}_0^- \cdot 1_0^+ \in \mathbf{cell}_0$ as well, which is not the case. Of course, it is precisely the role of \mathbf{cell} and \mathbf{lock} to break innocence and transfer information across distinct copies – however, \mathbf{cell} and \mathbf{lock} remain *P-visible* in the sense of Definition 3.5.

We now complete the interpretation of \mathbf{IA} . Consider $\Gamma, x : \mathbf{ref} \vdash M : A$ with

$$\llbracket M \rrbracket \in \downarrow \uparrow \text{-Strat}_!^{\mathbf{wb}, \mathbf{vis}}(\Gamma \ \& \ \mathbf{ref}, A)$$

omitting some brackets. Using the cartesian closed structure of $\downarrow \uparrow \text{-Strat}_!^{\mathbf{wb}, \mathbf{vis}}$, we consider

$$\Lambda_\Gamma^!(\llbracket M \rrbracket) \in \downarrow \uparrow \text{-Strat}_!^{\mathbf{wb}, \mathbf{vis}}(\mathbf{!ref}, \mathbf{!}\Gamma \multimap A).$$

which we compose with the memory cell. Summing up, for references and semaphores,

$$\begin{aligned} \llbracket \mathbf{newref} \ x := n \ \mathbf{in} \ M \rrbracket &= \Lambda_\Gamma^{!-1}(\Lambda_\Gamma^!(\llbracket M \rrbracket) \odot \mathbf{cell}_n) \in \downarrow \uparrow \text{-Strat}_!^{\mathbf{wb}, \mathbf{vis}}(\Gamma, A) \\ \llbracket \mathbf{newsem} \ x := n \ \mathbf{in} \ M \rrbracket &= \Lambda_\Gamma^{!-1}(\Lambda_\Gamma^!(\llbracket M \rrbracket) \odot \mathbf{lock}_n) \in \downarrow \uparrow \text{-Strat}_!^{\mathbf{wb}, \mathbf{vis}}(\Gamma, A). \end{aligned}$$

This concludes the interpretation of \mathbf{IA} in $\downarrow \uparrow \text{-Strat}_!^{\mathbf{wb}, \mathbf{vis}}$. Adequacy proceeds as in [AM96], undisturbed by the slightly different technical setting of the present paper.

Proposition 3.12 (Adequacy). *For any $\vdash M : \mathbb{U}$ in \mathbf{IA} , $M \Downarrow$ if and only if $\llbracket M \rrbracket \Downarrow$.*

3.4.4. *Full Abstraction.* We now review the full abstraction result of [AM96]. The argument revolves around a fundamental *factorisation* theorem, stated as follows.

Theorem 3.13 (Factorisation). *Let A be a type of \mathbf{IA} , and $\sigma : \llbracket A \rrbracket$ be P -visible well-bracketed. Then, there is an innocent well-bracketed $\text{Inn}(\sigma) \in \downarrow\uparrow\text{-Strat}^{\text{wb,inn}}(\mathbf{!ref}, \llbracket A \rrbracket)$ such that*

$$\sigma \approx \text{Inn}(\sigma) \odot \text{cell}_0,$$

and $\text{Inn}(\sigma)$ is finite if σ is finite.

Proof. For any O -ending $sa^- \in \sigma$, we wish $\text{Inn}(\sigma)$ to act like σ , but as an innocent strategy it may only depend on $\ulcorner sa^- \urcorner$. However, $\text{Inn}(\sigma)$ may also access the reference, so we will maintain the invariant that the reference contains (an encoding of) the full history, or more precisely of $\mathcal{P}(s)$. Between $\ulcorner sa^- \urcorner$ and $\mathcal{P}(s)$, σ knows the full play (up to symmetry).

Upon being called with a^- , $\text{Inn}(\sigma)$ reads $\mathcal{P}(s)$ from the reference, then stores $\mathcal{P}(sab)$ in the reference (for $sa^-b^+ \in \sigma$) and then plays b . See [AM96] for more details. \square

Finiteness of $\text{Inn}(\sigma)$ follows the definition of finite *innocent* strategies from Section 3.3.1: having finitely many (\cong -equivalence classes of) P -ending P -views. However, finiteness of non-innocent strategies has not yet been defined. We define it now: a strategy in $\downarrow\uparrow\text{-Strat}(A)$ is **finite** iff the set of (\cong -equivalence classes of) P -ending plays of σ is finite. Despite the common terminology, these two notions are distinct: an innocent strategy may be finite as an innocent strategy while being non-finite as a non-innocent strategy. This mismatch comes from the fact that these two notions both coincide with the domain-theoretic notion of *compactness*, but in the distinct domains $\downarrow\uparrow\text{-Strat}^{\text{wb,inn}}(A, B)$ and $\downarrow\uparrow\text{-Strat}^{\text{wb}}(A, B)$ (ordered by inclusion) for $--$ arenas A, B . By Proposition 2.22, these statements involving \cong -equivalence classes may be instead phrased with plays with pointers.

From Theorem 3.13 and Proposition 3.11 it is direct that finite definability holds for \mathbf{IA} . We can deduce immediately intensional full abstraction for \mathbf{IA} , proved as Theorem 3.13.

Theorem 3.14. *The model $\downarrow\uparrow\text{-Strat}_I^{\text{wb,vis}}$ is intensionally fully abstract for \mathbf{IA} .*

This is exactly as Theorem 3.10. However, in stark contrast with Theorem 3.10, for \mathbf{IA} the fully abstract quotient category is effectively presentable. In fact, for $\sigma, \tau : \llbracket A \rrbracket$,

$$\sigma \sim \tau \iff \mathcal{P}(\text{comp}(\sigma)) = \mathcal{P}(\text{comp}(\tau))$$

where $\text{comp}(\sigma)$ is the set of *complete* plays of σ , capturing the completed executions where both players act like P -visible well-bracketed strategies: a play is **complete** if it is well-bracketed, P -visible, *O-visible* (the dual to P -visibility, not detailed here), and such that every question has an answer. The result follows from finite definability for \mathbf{IA} [AM96].

This effective fully abstract model of \mathbf{IA} is one of the most striking results of game semantics. Observational equivalence in \mathbf{IA} remains undecidable with bounded integers, at fourth order without recursion [Mur03] and second-order with recursion [Ong02] (of course, observational equivalence is obviously undecidable in the full language as it is Turing-complete). However, the model yielded sound and complete algorithms for observational equivalence on restricted fragments [GM03], starting the field of *algorithmic game semantics*.

3.5. The Semantic Cube. Abramsky’s “semantic cube”, often called the “Abramsky cube”, starts with the observation that game semantics allows the interpretation of both control (*i.e.* **callcc**) and state in the same model, *i.e.* the same category.

3.5.1. *Control.* We have not given the interpretation of **callcc**, nor the corresponding full abstraction result [Lai97]. In fact, in the present technical setting, we cannot interpret **callcc**. This is due, in part, to the added conflict in arenas for basic datatypes with respect to standard HO games [HO00] – see Figures 6, 7, and 8. This conflict imposes that each question can be answered at most once, which is incompatible with **callcc**¹². In fact:

Proposition 3.15. *For any PCF type A , any innocent $\sigma : \llbracket A \rrbracket$ is also well-bracketed.*

Proof. First, any innocent $\sigma : \llbracket A \rrbracket$ is well-bracketed iff its P-views are well-bracketed – see [Lai97] for a proof. Hence if σ is not well-bracketed, then there is a P-view

$$s_1 \dots \mathbf{q}^{-, \mathcal{Q}} \dots \mathbf{q}^{-, \mathcal{Q}} \dots a^{+, \mathcal{A}}$$

where a answers the first \mathbf{q} shown rather than the pending question, the second \mathbf{q} shown. But this second $\mathbf{q}^{-, \mathcal{Q}}$ must be the initial move of a banged sub-arena in the interpretation of A , so we can play it again. And by innocence of σ , the following must be a play of σ :

$$s_1 \dots \mathbf{q}^{-, \mathcal{Q}} \dots \mathbf{q}_i^{-, \mathcal{Q}} \dots a^{+, \mathcal{A}} \mathbf{q}_{i+1}^{-, \mathcal{Q}} \dots a^{+, \mathcal{A}}$$

where both copies of a point to the first $\mathbf{q}^{-, \mathcal{Q}}$, absurd by *non-repetitive*. \square

This entails that in fact, Theorem 3.10 holds for $\downarrow\uparrow\text{-Strat}_i^{\text{inn}}$. But no such coincidence holds beyond innocent strategies: for Theorem 3.14 well-bracketing really is needed. In this paper we have adopted an interpretation of ground types incompatible with **callcc**. There is no technical obstacle to modelling **callcc** – one can simply drop conflicts in basic arenas and duplicate return values – but we prefer our design, closer to linear logic and the relational model (see Section 7.1.6). Furthermore, control operators will play no role in the present paper beyond the exposition of the scientific context.

3.5.2. *The Semantic Cube.* We temporarily consider, for the sake of the discussion, a setting with both control and state; say Murawski’s model for interference and control [Mur07], which is essentially equivalent (modulo the representation with pointers) to ours where basic arenas have no conflict and answers are replicated. Let us call it by *Vis*. There is

$$\text{PCF} + \text{interference} + \text{control} \quad \rightarrow \quad \text{Vis}$$

an adequate interpretation, so we can model a rich combination of effects; but that is not all.

Indeed, there are four (intensional) full abstraction results:

Theorem 3.16 (Semantic Cube). *We have four intensional full abstraction results:*

$$\begin{array}{ll} \text{Vis} & \text{is fully abstract for} \quad \text{PCF} + \text{interference} + \text{control} \\ \text{Vis} + \text{innocence} & \text{is fully abstract for} \quad \text{PCF} + \text{control} \\ \text{Vis} + \text{well-bracketing} & \text{is fully abstract for} \quad \text{PCF} + \text{interference} \\ \text{Vis} + \text{innocence} + \text{well-bracketing} & \text{is fully abstract for} \quad \text{PCF} \end{array}$$

We have reviewed two cases before, namely PCF (Theorem 3.10) and PCF + interference (Theorem 3.14). The full abstraction result for PCF + control is due to Laird [Lai97], while for PCF + interference + control appears in Murawski¹³ [Mur07].

¹²In addition to conflicts, the incompatibility with **callcc** comes from the fact that our interpretation of types only involves ! on arrows, and not on basic datatypes. To authorize control we should change *e.g.* the arena \mathbb{B} to one with replicated answers, written (in the language of tensorial logic [MT10]) as $\neg! \neg(1 \oplus 1)$.

¹³Murawski uses a different primitive for control, but the difference is superficial within \mathbf{IA} .

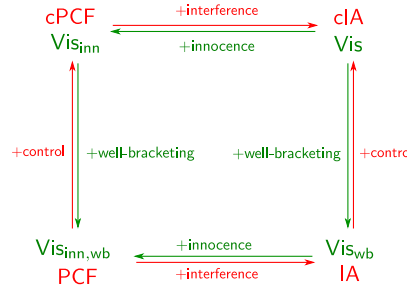


Figure 28: The Semantic Cube

This “Semantic Cube”, drawn in Figure 28, expresses that the conditions on strategies capture the behaviour generated by certain computational effects; or rather the *absence* of certain effects. The achievement is noteworthy, as it is famously difficult to *combine* semantic accounts of computational effects. But independently of purely semantic purposes, this provides us with a microscope to study behaviourally interactions between effects in programming languages. We demonstrate this with the following orthogonality property¹⁴ between interference and control which nicely illustrates the strength of game semantics:

Theorem 3.17. *Let $\vdash M : A$ a term of PCF + interference + control. Assume that*

- (1) $M \sim N_1$ where N_1 is a term of PCF + interference,
- (2) $M \sim N_2$ where N_2 is a term of PCF + control;

then $M \sim N$ where N is a term of (an infinitary extension of) PCF.

Proof. Consider $\llbracket M \rrbracket : A$. We have seen in Section 3.4.4 that for IA, strategies are indistinguishable iff they have the same *complete plays*. In the presence of control this phenomenon gets stronger: strategies are indistinguishable iff they have the *same plays* [Mur07]. Hence, $\llbracket M \rrbracket$ is an innocent well-bracketed strategy (even though M might internally use state and control). It is approximated by a sequence of finite innocent strategies which may be defined; but as the definability process is monotone this yields an infinitary PCF term. \square

It is widely believed that in a version of PCF such as ours with a **let** construct, the innocent well-bracketed games model is *intensionally universal*, meaning that each computable innocent well-bracketed strategy is definable¹⁵. With such a result, Theorem 3.17 would generalize to conclude the existence of simply a term of PCF, rather than an infinitary term.

3.6. Towards Concurrency. The reader may rightly complain that Figure 28 is not a “semantic cube”, only a “semantic square”. Though we focused on control and interference, there are fully abstract models of languages featuring general references [AHM98], exceptions [Lai01a], coroutines [Lai04], non-determinism [HM99], probabilistic choice [DH00], concurrency [Lai01b, GM08], and others. One imagines that the methodology above generalizes, and that the big “syntactic hypercube” of these effects is matched by a “semantic hypercube”.

However, there is no such “semantic hypercube”: the works cited above rely on a priori incompatible formal settings. In this paper, we present steps towards such a semantic

¹⁴We learnt of it from a talk by Paul Levy in 2014 [Lev14].

¹⁵Hyland and Ong have a *extensional* universality theorem [HO00], *i.e.* up to observational equivalence. *Intensional* universality does not appear anywhere in call-by-name, although it does in call-by-value [MT13].

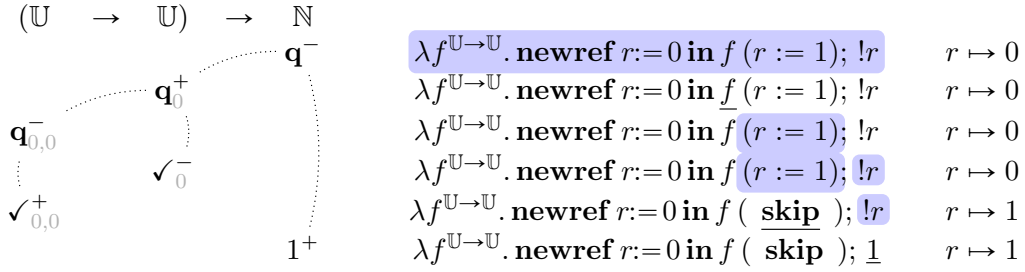


Figure 29: Operational content of a non-alternating play

hypercube. More precisely we aim to *disentangle parallelism and interference* in the same sense as the “Abramsky cube” disentangles control and interference, *i.e.* we must answer:

Question. *Build a model \rightarrow -Strat with notions of parallel innocence and sequentiality s.t.:*

$$\begin{aligned}
\rightarrow\text{-Strat} & \text{ is fully abstract for } \mathbf{IA}_{\parallel}, \\
\rightarrow\text{-Strat} + \text{parallel innocence} & \text{ is fully abstract for } \mathbf{PCF}_{\parallel}, \\
\rightarrow\text{-Strat} + \text{sequentiality} & \text{ is fully abstract for } \mathbf{IA}, \\
\rightarrow\text{-Strat} + \text{parallel innocence} + \text{sequentiality} & \text{ is fully abstract for } \mathbf{PCF},
\end{aligned}$$

all of these being intensional full abstraction results.

The model should be fully abstract for \mathbf{IA}_{\parallel} , link with $\downarrow\uparrow\text{-Strat}_{!}^{\text{wb,inn}}$ and $\downarrow\uparrow\text{-Strat}_{!}^{\text{wb}}$ (respectively fully abstract for \mathbf{PCF} and \mathbf{IA}), but also support a notion of parallel innocence yielding full abstraction for \mathbf{PCF}_{\parallel} . It is natural to start with a simple non-alternating variant of $\downarrow\uparrow\text{-Strat}$, inspired by Ghica and Murawski’s fully abstract games model for \mathbf{IA}_{\parallel} [GM08].

3.6.1. *Non-alternating plays and strategies.* We simply relax alternation in Definition 2.7.

Definition 3.18. A **non-alternating play** on $--$ -arena A is $s = s_1 \dots s_n$ which is:

$$\begin{aligned}
\text{valid:} & \quad \forall 1 \leq i \leq n, \{s_1, \dots, s_i\} \in \mathcal{C}(A), \\
\text{non-repetitive:} & \quad \forall 1 \leq i, j \leq n, s_i = s_j \implies i = j, \\
\text{negative:} & \quad n \geq 1 \implies \text{pol}(s_1) = -.
\end{aligned}$$

We write $\mathcal{O}\text{-Plays}(A)$ for the set of non-alternating plays on A .

The notation (inspired by *template games* [Mel19]), is intended to suggest that whereas alternating plays in $\downarrow\uparrow\text{-Plays}(A)$ transition between two states O and P determining which player has control, in $\mathcal{O}\text{-Plays}(A)$ there is only one state, in which either player may play. The intuition is simple: as several threads might be running in parallel, their interleaving breaks alternation. We show in Figure 29 a non-alternating play on $\llbracket (\mathbb{U} \rightarrow \mathbb{U}) \rightarrow \mathbb{N} \rrbracket$, using the same conventions as previously. Resting on the same computational intuitions as before, we show for each move a representation of the matching computational state of a term realizing that play. The figure illustrates that even \mathbf{IA} , a sequential language, allows non-alternating plays, as the environment can evaluate subterms in parallel. After the third and fourth moves, two subterms are being evaluated in parallel: $r := 1$ and $!r$, causing a *data race*.

As before, we may now define *strategies* as certain sets of plays.

Definition 3.19. A **non-alternating strategy** $\sigma : A$, is $\sigma \subseteq \mathcal{C}\text{-Plays}(A)$ which is:

- non-empty:* $\varepsilon \in \sigma$
- prefix-closed:* $\forall s \sqsubseteq s' \in \sigma, s \in \sigma$
- receptive:* $\forall s \in \sigma, sa^- \in \mathcal{C}\text{-Plays}(A) \implies sa \in \sigma$
- courteous:* $\forall sabt \in \sigma, sb \in \mathcal{C}\text{-Plays}(A), (\text{pol}(b) = - \vee \text{pol}(a) = +) \implies sbat \in \sigma$

A **non-alternating prestrategy** is only required to satisfy *non-empty* and *prefix-closed*.

Let us compare with Definition 2.8. Besides moving to non-alternating plays, we remove determinism. Of course, this is natural since the interleaving semantics of even a pure parallel language represents non-deterministically the choices of the scheduler. The new condition added is *courtesy*, it corresponds to the *saturation* condition of [GM08] (the name “courtesy” is imported from [MM07]). Courtesy expresses that the model is asynchronous. If one has $sa^+b \in \sigma$, there is an execution of σ where it plays a , then we observe b (of any polarity). But if the surrounding computing environment is asynchronous, nothing forces a to be directly observable by Opponent – a might get stuck in a buffer, in the network, *etc.* So then, courtesy imposes that σ should be *stable under asynchronous delays*: if $sa^+b \in \sigma$ and there is no dependency from a to b in the arena, then a can be postponed after b in σ .

3.6.2. *Well-bracketing.* We introduce *well-bracketing* for non-alternating strategies. In Ghica and Murawski’s games, all plays are *well-bracketed* in the sense that they satisfy two conditions, dubbed *fork* and *wait*. We adapt and introduce these conditions now.

Definition 3.20. For a $--$ -arena A , $s \in \mathcal{C}\text{-Plays}(A)$ is **well-bracketed** if for any $s' \sqsubseteq s$,

- fork:* if $s' = \dots \mathbf{q}^Q \dots m$ with $\mathbf{q} \rightarrow_A m$, \mathbf{q} must be unanswered before m is played,
- wait:* if $s' = \dots \mathbf{q}^Q \dots a^A$ with $\mathbf{q} \rightarrow_A a$, all questions justified by \mathbf{q} must be answered.

This differs from the simple *well-bracketing* of alternating plays (Definition 3.2). In a non-alternating setting it does not make sense to refer to the last unanswered question as it might originate in a different thread than the one the Player move to be played belongs to. Instead, this condition forces plays to follow the following protocol: a question, as long as it is not answered, may prompt (*i.e.* justify) other questions. It can only be answered once all the questions it justified are answered, and then it will not be able to justify anything further. For instance, the play of Figure 29 is *not* well-bracketed: the fourth move \checkmark_0^- causes a failure to *wait* because it is justified by \mathbf{q}_0^+ , although the latter has justified $\mathbf{q}_{0,0}^-$ as of yet unanswered. If we were to permute the moves $\mathbf{q}_{0,0}^-$ and \checkmark_0^- then $\mathbf{q}_{0,0}^-$ would cause a failure to *fork* as it would be justified by a question that is already answered.

Rather than imposing this condition on all plays, we impose it on strategies.

Definition 3.21. Let $\sigma : A$ be a non-alternating strategy.

It is **well-bracketed** iff for all $sa^+ \in \sigma$, if s is well-bracketed then sa is well-bracketed.

The play of Figure 29 belongs to a well-bracketed strategy: Opponent breaks well-bracketing first. This is a slight difference with Ghica and Murawski’s model: we observe all dynamic behaviour of a program of $\mathbf{IA}_{//}$, even that not reachable via a context of $\mathbf{IA}_{//}$. Of course, it is always possible to restrict to well-bracketed plays without cutting any Player behaviour (*i.e.* the strategy is cut at Opponent extensions, not Player extensions).




Figure 30: Two (pointer representation) of plays in a parallel innocent strategy

3.6.3. *Limitations.* The proximity of non-alternating strategies to both $\downarrow\uparrow$ -Strat and Ghica and Murawski’s model make them a tempting foundation for this paper. However, they have two limitations, both subtly related to their inability to record the *branching structure*.

Firstly, *parallel innocence* requires adequate causal structures, as illustrated by the following example. Is there a parallel innocent strategy that includes the two plays with pointer representation in Figure 30? Is there a program of PCF_{\parallel} that may realize both? Traditional innocence forbids that, because in a sequential program, both plays must be visiting the same piece of syntax and obtain the same result. In PCF_{\parallel} though, the program

$$\lambda f^{\mathbb{B} \rightarrow \mathbb{U}}. \text{let} \left(\begin{array}{l} x = f \mathbf{tt} \\ y = f \mathbf{ff} \end{array} \right) \text{in } x; y : (\mathbb{B} \rightarrow \mathbb{U}) \rightarrow \mathbb{U}$$

indeed realizes these two plays, corresponding to the evaluation of distinct threads. A deterministic innocent strategy is determined by (the pointer representation of) its *P-views*, so we may see the set of P-views as a witness for innocence. Analogously, what structure may witness that a non-alternating strategy is innocent? In fact, what is missing from the two P-views of Figure 30 is the *branching structure*, keeping these two P-views apart and recording how they are linked to each other. It has already been observed [CCW14, TO15] that *non-deterministic innocence* may be defined by replacing sets of P-views, as witnesses for innocence, with *trees* recording the non-deterministic branching information. Here we must do the same, but instead record the branching structure pertaining to *parallelism* as well – which plain non-alternating strategies cannot capture adequately [CC16].

Secondly, it is unclear how to endow non-alternating strategies with appropriate notions of *symmetry* and *uniformity* as in Definition 2.13. Our attempts in generalizing Definition 2.13 ended up suffering from various pathologies, typically uniformity not being preserved by hiding. The tension with hiding comes from a play $s \in \tau \odot \sigma$ being witnessed by distinct interactions between σ and τ – this suggests again the need for an explicitly branching structure, as then one recovers a *unique witness* property. For these two reasons, we move from plain non-alternating strategies to *thin concurrent games*. We will, however rely on:

Proposition 3.22. *There is a symmetric monoidal closed category with products $\mathcal{O}\text{-Strat}^{\text{wb}}$, with objects $--$ arenas, and morphisms well-bracketed non-alternating strategies on $A \multimap B$.*

Proof. The constructions play a minor role in this paper and are very similar to the alternating case, so we omit them. Some details of the construction appear in Appendix C. \square

These limitations call for a more intensional setting, representing explicitly the parallel and non-deterministic branching structures. To our knowledge, the only games setting in the literature sufficiently expressive and mature is *thin concurrent games*, one of the possible enrichments with symmetry of the concurrent strategies of Rideau and Winskel [RW11].

4. CAUSAL FULL ABSTRACTION FOR \mathbf{IA}_{\parallel}

Thin concurrent games were first introduced in [CCW15], but the detailed construction (along with significant improvements and simplifications) is presented in [CCW19].

In this section, we start with an introduction to thin concurrent games. We omit detailed constructions (which appear in [CCW19]), but we do attempt to give a self-contained introduction, in particular providing required reasoning principles. Then, we apply this setting to give a fully abstract model for \mathbf{IA}_{\parallel} , the causal sibling of Ghica and Murawski’s model [GM08]. In the *concurrent games* literature, strategies are often referred to as *concurrent strategies*. Here we prefer *causal strategies* to better distinguish them with non-alternating strategies, which also represent concurrent behaviour.

4.1. Arenas and causal strategies. First, we must refine the symmetry on arenas.

4.1.1. *Polarized symmetry.* Arenas for causal strategies in [CCW19] require the following:

Definition 4.1. For A an arena with isomorphism family $\mathcal{S}(A)$, a **polarized decomposition** of $\mathcal{S}(A)$ comprises isomorphism families $\mathcal{S}_-(A), \mathcal{S}_+(A)$ included in $\mathcal{S}(A)$, s.t.:

- (1) If $\theta \in \mathcal{S}_-(A) \cap \mathcal{S}_+(A)$, then θ is an identity bijection,
- (2) If $\theta \in \mathcal{S}_-(A)$ and $\theta \sqsubseteq^- \theta' \in \mathcal{S}(A)$, then $\theta' \in \mathcal{S}_-(A)$,
- (3) If $\theta \in \mathcal{S}_+(A)$ and $\theta \sqsubseteq^+ \theta' \in \mathcal{S}(A)$, then $\theta' \in \mathcal{S}_+(A)$.

where \sqsubseteq^p means that only (pairs of) events of polarity p are added.

If $\theta : x \cong_A y$ and $\theta \in \mathcal{S}_-(A)$, we write $\theta : x \cong_A^- y$ and call θ a **negative symmetry**. Likewise, $\theta : x \cong_A^+ y$ means that $\theta : x \cong_A y$ with $\theta \in \mathcal{S}_+(A)$, called a **positive symmetry**.

Arenas with a polarized decomposition are *thin concurrent games*¹⁶ in the sense of [CCW19]. Intuitively, negative symmetries (resp. positive) reindex Opponent (resp. Player) moves – though Definition 4.1 does not involve “copy indices”. By Lemma 3.19 from [CCW19], any $\theta : x \cong_A z$ factors uniquely as $\theta_+ \circ \theta_-$, with θ_+ positive and θ_- negative.

4.1.2. *Constructions on arenas.* First we extend the arena constructions accordingly.

For \mathbb{U}, \mathbb{B} and \mathbb{N} , all isomorphism families are reduced to identity bijections between configurations. For the dual of an arena A , then for all symmetry $\theta : x \cong_A y$,

$$\theta : x \cong_{A^\perp}^+ y \iff \theta : x \cong_A^- y, \quad \theta : x \cong_{A^\perp}^- y \iff \theta : x \cong_A^+ y.$$

For parallel composition and product, the sub-symmetries are inherited as for the full symmetry in Section 2.2.1 applying to the positive and negative isomorphism families separately. For the arrow, for $x, y \in \mathcal{C}(A \multimap B)$ and $\theta : x \cong y$ an order-isomorphism, we set $\theta \in \mathcal{S}_+(A \multimap B)$ iff $\chi_{A,B} \theta \in \mathcal{S}_+(A^\perp \parallel B)$; and $\theta \in \mathcal{S}_-(A \multimap B)$ iff $\chi_{A,B} \theta \in \mathcal{S}_-(A^\perp \parallel B)$.

The most interesting construction is the exponential. Recall that a symmetry in $!A$ is

$$\theta : \begin{array}{l} \parallel_{n \in \mathbb{N}} x_n \cong \parallel_{n \in \mathbb{N}} y_n \\ (n, a) \mapsto (\pi(n), \theta_n(a)) \end{array}$$

for some permutation $\pi \in \zeta(\mathbb{N})$ and for some family $(\theta_n)_{n \in \mathbb{N}}$ with $\theta_n : x_n \cong_A y_{\pi(n)}$ for all $n \in \mathbb{N}$. First we set $\theta \in \mathcal{S}_-(!A)$ iff for all $n \in \mathbb{N}$, $\theta_n : x_n \cong_A^- y_{\pi(n)}$. Finally, we set $\theta \in \mathcal{S}_+(!A)$ iff for all $n \in \mathbb{N}$ such that x_n is non-empty, we have $\pi(n) = n$, and $\theta_n : x_n \cong_A^+ y_n$.

¹⁶Thin concurrent games in [CCW19] are more general, *e.g.* they might not be alternating or forestial.

Why does $!A$ treat differently the positive and negative isomorphism families? The permutation $\pi(-)$ corresponds to reindexing the minimal events of $!A$. Because the exponential construction is intended to apply to negative games, $\pi(-)$ reindexes negative moves. But symmetries in $\mathcal{S}_+(!A)$ can only reindex positive moves, so $\pi(-)$ must be the identity. In contrast, symmetries in $\mathcal{S}_-(!A)$ can perform any reindexing on the minimal events. This intuition is further solidified by the extension of *concrete arenas* to these polarized sub-isomorphism families. This will not be required until much later, so it only appears in Section 7.3.1 – but it might still be helpful for the reader to consult it now.

4.1.3. *Causal strategies.* In contrast with traditional game models, a causal strategy is one global object: an event structure. It presents all execution threads together, with explicit information on how these executions relate via parallel and non-deterministic branching.

Definition 4.2. A **causal prestrategy** $\sigma : A$ comprises an ess $(|\sigma|, \leq_\sigma, \#_\sigma, \mathcal{S}(\sigma))$ with

$$\partial : |\sigma| \rightarrow |A|$$

a function called the **display map**, subject to the following conditions:

- rule-abiding:* for all $x \in \mathcal{C}(\sigma)$, $\partial(x) \in \mathcal{C}(A)$,
- locally injective:* for all $s_1, s_2 \in x \in \mathcal{C}(\sigma)$, if $\partial(s_1) = \partial(s_2)$ then $s_1 = s_2$,
- symmetry-preserving:* for all $\theta \in \mathcal{S}(\sigma)$, $\partial(\theta) = \{(\partial(s_1), \partial(s_2)) \mid (s_1, s_2) \in \theta\} \in \mathcal{S}(A)$,
- \sim -receptive:* for all $\theta : x \cong_\sigma y$, and extensions $x \vdash_\sigma s_1^-$, $\partial(\theta) \vdash_{\mathcal{S}(A)} (\partial(s_1^-), a_2^-)$, there is a unique $s_2^- \in |\sigma|$ s.t. $\theta \vdash_{\mathcal{S}(\sigma)} (s_1^-, s_2^-)$ and $\partial(s_2^-) = a_2^-$,
- thin:* for all $\theta : x \cong_\sigma y$, and extension $x \vdash_\sigma s_1^+$, there is a *unique* extension $y \vdash_\sigma s_2^+$ such that $\theta \vdash_{\mathcal{S}(\sigma)} (s_1^+, s_2^+)$.

Additionally, we say that σ is a **causal strategy** if it satisfies the further two conditions:

- negative:* for all $s \in |\sigma|$, if s is minimal then s is negative,
- courteous:* for all $s_1 \rightarrow_\sigma s_2$, if $\text{pol}(s_1) = +$ or $\text{pol}(s_2) = -$ then $\partial(s_1) \rightarrow_A \partial(s_2)$,
- receptive:* for all $x \in \mathcal{C}(\sigma)$, for all $\partial(x) \vdash_A a^-$, there is a unique $x \vdash_\sigma s^- \in \mathcal{C}(\sigma)$ such that $\partial(s) = a$,

As a convention, causal strategies are ranged over by symbols in bold font, as in *e.g.* σ, τ . We disambiguate some notations used in the definition. First, σ implicitly comes with polarities, imported from A as $\text{pol}_\sigma(s) = \text{pol}_A(\partial(s))$. We also used the enabling relation on isomorphism families, defined by $\theta \vdash_{\mathcal{S}(A)} (a_1, a_2)$ iff $(a_1, a_2) \notin \theta$ and $\theta \uplus \{(a_1, a_2)\} \in \mathcal{S}(A)$.

Conditions *rule-abiding*, *locally injective* and *symmetry-preserving* together amount to ∂ being a **map of event structures with symmetry** [Win07]. Conditions *courteous* and *receptive* play the same role as in Definition 3.19. The condition *\sim -receptive* forces strategies to treat uniformly any pairs of Opponent events symmetric in the game. Finally, *thin* forces strategies to pick *one* canonical representative up to symmetry for positive moves. For further explanations and discussions on those conditions, the reader is directed to [CCW19].

Causal strategies and non-alternating strategies differ fundamentally in *how* the concurrent behaviour is represented. While non-alternating strategies present observable execution traces, causal strategies present the causal constraints underlying the observed behaviour. In Figure 31, we present a causal strategy, corresponding to a linear version of Figure 29. In this diagram and others further on, we draw $\sigma : A$ by picturing the event structure σ with events displayed as their image through ∂_σ . Whenever possible, we keep the convention to draw moves under the corresponding type component. The causal dependency \leq_σ is




Figure 31: A causal strategy





Figure 32: Two augmentations of Figure 31

pictured via its immediate dependency relation \rightarrow_σ . As with arenas, we show immediate conflicts (*i.e.* not inherited) as wiggly lines. As in plays, we denote by dotted lines the immediate causal dependency in the arena A : for $s \in |\sigma|$, either $\partial_\sigma(s)$ is minimal in A , or there is a unique $s' \in |\sigma|$ such that $\partial_\sigma(s') \rightarrow_A \partial_\sigma(s)$; in which case the diagram has a dotted line between the events representing s and s' . Borrowing earlier terminology, we refer to s' as the **justifier** of s . Finally, the symmetry $\mathcal{S}(\sigma)$ is not shown at all (of course, for Figure 31 it would remain trivial). Indeed it is hard to represent, but also we regard it as not being part of the pertinent operational structure: its mere existence witnesses uniformity.

Figure 31 presents, in *one* diagram, the full behaviour of the program of Figure 29 under linear execution contexts. Drawing a strategy fully in this way is sometimes challenging. It is often convenient to refer to $-$ and draw $-$ consistent fragments of a causal strategy: an **augmentation** of σ is any (x, \leq_x) where $x \in \mathcal{C}(\sigma)$ and \leq_x is the partial order inherited from \leq_σ . For instance the two maximal augmentations of Figure 31 yield the diagrams of Figure 32. In this case, the two configurations correspond to the two resolutions of the data race described with Figure 29: if the write wins, the read yields 1 and depends on q^- (which triggered the write). If the read wins, we read 0 and the write acknowledgment depends on \checkmark^- (which triggered the read). However, this representation of a strategy via its augmentations is partial: it forgets the explicit non-deterministic branching.

4.1.4. *Non-alternating unfolding.* Causal strategies generate non-alternating strategies:

Proposition 4.3. Consider A a $--$ -arena, and $\sigma : A$ a causal strategy.

The **non-alternating unfolding** of σ is (with ∂_σ applied to plays move-by-move):

$$\mathcal{U}\text{-Unf}(\sigma) = \partial_\sigma(\mathcal{U}\text{-Plays}(\sigma)).$$

Thus defined, $\mathcal{U}\text{-Unf}(\sigma) : A$ is a non-alternating strategy on A .

Proof. *Non-empty* and *prefix-closed* are obvious. For *receptive*, take $s \in \mathcal{U}\text{-Unf}(\sigma)$ s.t. $sa^- \in \mathcal{U}\text{-Plays}(A)$. By definition, there is $t \in \mathcal{U}\text{-Plays}(\sigma)$ s.t. $s = \partial_\sigma(t)$. Thus t , written

$$t = t_1 \dots t_n \in \mathcal{U}\text{-Plays}(\sigma),$$

is such that for all $1 \leq i \leq n$, $\{t_1, \dots, t_i\} \in \mathcal{C}(\sigma)$, and is also *non-repetitive* and *negative*. Let us write $x = \{t_1, \dots, t_n\}$. Then, $\partial_\sigma(x) = \{s_1, \dots, s_n\} \in \mathcal{C}(A)$, and since $sa^- \in \mathcal{U}\text{-Plays}(A)$ we have $\partial_\sigma(x) \vdash_A a^-$. By *receptivity* of σ , there is (a unique) $m \in |\sigma|$ such that $x \vdash_\sigma m$ and $\partial_\sigma(m) = a$. It is then direct that $tm \in \mathcal{U}\text{-Plays}(\sigma)$, witnessing $sa \in \mathcal{U}\text{-Unf}(\sigma)$ as required.

For *courteous*, consider $s_1abs_2 \in \mathcal{U}\text{-Unf}(\sigma)$ s.t. $\text{pol}(a) = +$ or $\text{pol}(b) = -$, and $s_1b \in \mathcal{U}\text{-Plays}(A)$. By definition, there is $t_1mnt_2 \in \mathcal{U}\text{-Plays}(\sigma)$ s.t. $s_1abs_2 = \partial_\sigma(t_1mnt_2)$. We claim

that $x = |t_1 n| \in \mathcal{C}(\sigma)$. We know that $|t_1 mn| \in \mathcal{C}(\sigma)$; so x is consistent. If it is not down-closed, necessarily $m \rightarrow_{\sigma} n$. So, by *courtesy* of σ , $a \rightarrow_A b$; contradicting $s_1 b \in \mathcal{C}\text{-Plays}(A)$. Thus $|t_1 n| \in \mathcal{C}(\sigma)$ from which we deduce $t_1 n m t_2 \in \mathcal{C}\text{-Plays}(\sigma)$, so $s_1 b a s_2 \in \mathcal{C}\text{-Unf}(\sigma)$. \square

For instance, the non-alternating play in Figure 29 is in the unfolding of the causal strategy in Figure 31. This extraction of a non-alternating strategy is an instance of the usual relationship between interleaving and “truly concurrent” models for concurrency. In this paper this relationship will in particular allow us to import well-bracketing from $\mathcal{C}\text{-Strat}^{\text{wb}}$.

Definition 4.4. Consider $\sigma : A$ a causal strategy on $\text{--arena } A$.

We say that σ is **well-bracketed** if $\mathcal{C}\text{-Unf}(\sigma) : A$ is well-bracketed.

4.2. A Category of Causal Strategies. We now start building the categorical operations on causal strategies, aiming at a Seely category $\text{--}\rightarrow\text{-Strat}$. We first focus on composition.

For $\text{--arenas } A$ and B , a causal strategy **from } A \text{ to } B is a causal strategy**

$$\sigma : A^{\perp} \parallel B,$$

in the sequel we also write $A \vdash B$ for $A^{\perp} \parallel B$. This is unlike for $\downarrow\uparrow\text{-Strat}$ and $\mathcal{C}\text{-Strat}$ introduced earlier, for which morphisms from A to B were defined as strategies on $A \multimap B$. We do this to keep close to [CCW19]. When linking with $\downarrow\uparrow\text{-Strat}$ and $\mathcal{C}\text{-Strat}$ we shall deal with this mismatch, but that will not cause us too much trouble¹⁷.

Composition of causal strategies is more elaborate than for play-based strategies. We define it in several stages. Fix $\sigma : A \vdash B$ and $\tau : B \vdash C$ two causal (pre)strategies.

4.2.1. Synchronization. For configurations $x^{\sigma} \in \mathcal{C}(\sigma), x^{\tau} \in \mathcal{C}(\tau)$, as a convention we write

$$\partial_{\sigma}(x^{\sigma}) = x_A^{\sigma} \parallel x_B^{\sigma} \in \mathcal{C}(A \vdash B), \quad \partial_{\tau}(x^{\tau}) = x_B^{\tau} \parallel x_C^{\tau} \in \mathcal{C}(B \vdash C),$$

for the corresponding projections to the game. In defining composition, the first stage is to capture when such configurations $x^{\sigma} \in \mathcal{C}(\sigma)$ and $x^{\tau} \in \mathcal{C}(\tau)$ may successfully *synchronise*.

Definition 4.5. Consider two configurations $x^{\sigma} \in \mathcal{C}(\sigma)$ and $x^{\tau} \in \mathcal{C}(\tau)$. They are **causally compatible** if (1) they are **matching**: $x_B^{\sigma} = x_B^{\tau} = x_B$; and (2) if the composite bijection

$$\varphi_{x^{\sigma}, x^{\tau}} : x^{\sigma} \parallel x_C^{\tau} \xrightarrow{\partial_{\sigma} \parallel x_C^{\tau}} x_A^{\sigma} \parallel x_B \parallel x_C^{\tau} \xrightarrow{x_A^{\sigma} \parallel \partial_{\tau}^{-1}} x_A^{\sigma} \parallel x^{\tau},$$

using local injectivity of ∂_{σ} and ∂_{τ} , is **secured**, in the sense that the relation

$$(m, n) \triangleleft (m', n') \quad \Leftrightarrow \quad m <_{\sigma \parallel C} m' \quad \vee \quad n <_{A \parallel \tau} n',$$

defined on (the graph of) $\varphi_{x^{\sigma}, x^{\tau}}$ by importing causal constraints of σ and τ , is acyclic.

Two matching $x^{\sigma} \in \mathcal{C}(\sigma), x^{\tau} \in \mathcal{C}(\tau)$ agree on the state reached in B . By local injectivity of ∂_{σ} and ∂_{τ} , this induces a bijection as above, thought of as the induced synchronization between events of x^{σ} and x^{τ} that match in B . But this is not enough to capture a sensible notion of execution: some matching pairs might not be *reachable*, in the sense that σ and τ impose incompatible constraints as to the order in which the state should be reached. To illustrate this we show in Figures 33 and 34 two attempted synchronizations between configurations of the strategy of Figure 31 and the causal strategy for the identity $\lambda x^{\cup}. x$. In

¹⁷Plays on $A \multimap B$ carry more information than on $A \vdash B$, namely the justifier for initial moves in A . With causal strategies, that information may be read back from the causal structure. See Section 4.3.2.




Figure 33: Matching, secured

Figure 34: Matching, non-secured

both cases, the configurations are matching. In Figure 33, the synchronization is successful and yields causally compatible pairs of configurations. However, in Figure 34 the induced bijection is not *secured*: the two strategies impose opposite constraints as to the order in which the two \checkmark moves are to be played. Thus, this synchronization fails. For us, this will entail that the identity $\lambda x^U. x$ may only synchronize successfully with the augmentation of the program of Figure 29 appearing in Figure 33 – so that the only final result is 1.

4.2.2. *Interaction.* But we must present the interaction of σ and τ as an event structure. More specifically, it should be an event structure with symmetry along with a display map:

Definition 4.6. A **pre-interaction** on A, B, C is an ess $\mu = (|\mu|, \leq_\mu, \#_\mu, \mathcal{S}(\mu))$ with

$$\partial : |\mu| \rightarrow |A \parallel B \parallel C|$$

a display map subject to the following conditions:

- rule-abiding:* for all $x \in \mathcal{C}(\mu)$, $\partial(x) \in \mathcal{C}(A \parallel B \parallel C)$,
- locally injective:* for all $s_1, s_2 \in x \in \mathcal{C}(\mu)$, if $\partial(s_1) = \partial(s_2)$ then $s_1 = s_2$,
- symmetry-preserving:* for all $\theta \in \mathcal{S}(\mu)$, $\partial(\theta) \in \mathcal{S}(A \parallel B \parallel C)$,

i.e. $\partial : \mu \rightarrow A \parallel B \parallel C$ is a map of event structures with symmetry.

An **isomorphism** between pre-interactions μ, ν on A, B, C is an isomorphism $f : \mu \cong \nu$ in the category of event structures with symmetry, commuting with the display maps, *i.e.* $\partial_\nu \circ f = \partial_\mu$. The *interaction* between σ and τ is a pre-interaction whose configurations correspond *exactly* with pairs of causally compatible configurations $x^\sigma \in \mathcal{C}(\sigma), x^\tau \in \mathcal{C}(\tau)$:

Proposition 4.7. *There is a pre-interaction $\tau \circledast \sigma$, the **interaction** of σ and τ , with*

$$(- \circledast -) : \{(x^\tau, x^\sigma) \in \mathcal{C}(\tau) \times \mathcal{C}(\sigma) \mid x^\sigma \text{ and } x^\tau \text{ are causally compatible}\} \simeq \mathcal{C}(\tau \circledast \sigma)$$

an order-iso (with causally compatible pairs ordered by component-wise inclusion) satisfying

$$\partial_{\tau \circledast \sigma}(x^\tau \circledast x^\sigma) = x_A^\sigma \parallel x_B \parallel x_C^\tau$$

for all $x^\sigma \in \mathcal{C}(\sigma)$ and $x^\tau \in \mathcal{C}(\tau)$ causally compatible.

In particular, any $z \in \mathcal{C}(\tau \circledast \sigma)$ is written uniquely as $x^\tau \circledast x^\sigma$ for $x^\sigma \in \mathcal{C}(\sigma)$ and $x^\tau \in \mathcal{C}(\tau)$. Thus, $\mathcal{C}(\tau \circledast \sigma)$ may be regarded as the subset of $\mathcal{C}(\sigma) \times \mathcal{C}(\tau)$ restricted to those of the matching pairs which cause no deadlock. In fact, this property *almost* suffices




Figure 35: Example of an interaction

to characterise the interaction uniquely – to complete the picture, we must also consider *symmetries*. Because display maps preserve symmetry, for $\theta^\sigma \in \mathcal{S}(\sigma)$ and $\theta^\tau \in \mathcal{S}(\tau)$,

$$\partial_\sigma(\theta^\sigma) = \theta_A^\sigma \parallel \theta_B^\sigma, \quad \partial_\tau(\theta^\tau) = \theta_B^\tau \parallel \theta_C^\tau,$$

and we can say that θ^σ and θ^τ are **matching** if $\theta_B^\sigma = \theta_B^\tau$, and **causally compatible** if $\text{dom}(\theta^\sigma)$, $\text{dom}(\theta^\tau)$ are causally compatible – or, equivalently, $\text{cod}(\theta^\sigma)$ and $\text{cod}(\theta^\tau)$ are.

Taking into account symmetry, we may strengthen Proposition 4.7 to:

Proposition 4.8. *There is a pre-interaction $\tau \otimes \sigma$, unique up to iso, such that there are*

$$\begin{aligned} (- \otimes -) &: \{(x^\tau, x^\sigma) \in \mathcal{C}(\tau) \times \mathcal{C}(\sigma) \mid x^\sigma, x^\tau \text{ causally compatible}\} \simeq \mathcal{C}(\tau \otimes \sigma) \\ (- \otimes -) &: \{(\theta^\tau, \theta^\sigma) \in \mathcal{S}(\tau) \times \mathcal{S}(\sigma) \mid \theta^\sigma, \theta^\tau \text{ causally compatible}\} \simeq \mathcal{S}(\tau \otimes \sigma) \end{aligned}$$

order-isomorphisms commuting with dom and cod, and satisfying

$$\partial_{\tau \otimes \sigma}(\theta^\tau \otimes \theta^\sigma) = \theta_A^\sigma \parallel \theta_B \parallel \theta_C^\tau$$

for all $\theta^\sigma \in \mathcal{S}(\sigma)$ and $\theta^\tau \in \mathcal{S}(\tau)$ causally compatible.

Proof. Follows from the characterisation of the interaction as a *pullback*, whose projections

$$\sigma \parallel C \xleftarrow{\Pi_\sigma} \tau \otimes \sigma \xrightarrow{\Pi_\tau} A \parallel \tau$$

are maps of event structures with symmetry [CCW19] – see Appendix D.2.1 for details. \square

There is some redundancy in this statement: first, the action of $(-\otimes-)$ on configurations coincides with that on identity symmetries. Reciprocally, one can actually prove that the action of $(-\otimes-)$ on symmetries, *if it exists*, is uniquely determined by that on configurations – so the fact that $(-\otimes-)$ extends to symmetries is property rather than structure. Nevertheless, in the sequel, we often find convenient to perform the constructions on symmetries explicitly. Altogether, this characterises the interaction in terms of its *states* and *symmetries*.

But there is also an alternative, *event-based* view: an individual event $m \in |\tau \otimes \sigma|$ may be regarded as a synchronization between its projections $\Pi_\sigma(m) \in |\sigma \parallel C|$ and $\Pi_\tau(m) \in |A \parallel \tau|$. But we warn against the misleading idea that $m \in |\tau \otimes \sigma|$ is determined by these projections: intuitively, there is one event in $\tau \otimes \sigma$ for each pair (s, t) of synchronizable events, *and* each distinct way to reach s and t conjointly in σ and τ . A simple example appears in Figure 35 (ignoring for now the l/r annotation and the part in grey): the final two copies of \checkmark have the same projections, but a different causal history. Though we shall not unfold the concrete construction of the interaction [CCW19], it might nonetheless help the reader to have an idea of what its events are concretely defined to be. For $x^\sigma \in \mathcal{C}(\sigma)$ and $x^\tau \in \mathcal{C}(\tau)$ causally compatible, the reflexive transitive closure of \triangleleft (see Definition 4.5) yields a partial

order \leq_{x^σ, x^τ} on (the graph of) $\varphi_{x^\sigma, x^\tau}$. The *events* of $\tau \circledast \sigma$ are then precisely the causally compatible pairs (x^σ, x^τ) such that \leq_{x^σ, x^τ} has a top element: the pair $(\Pi_\sigma(m), \Pi_\tau(m))$.

In the sequel, we shall only reason on the interaction through the proxy of Proposition 4.8 and forthcoming lemmas characterizing immediate causality in the interaction. However, to ease the flow of the exposition, those are postponed to Section 4.2.6.

4.2.3. Composition. Following the traditional methodology of game semantics, composition is defined from interaction via *hiding*. We first briefly analyse the *components* of interactions.

The projections Π_σ and Π_τ project any event of $\tau \circledast \sigma$ to a matching pair of an event of $\sigma \parallel C$ and an event of $A \parallel \tau$. These projections help us classify every $p \in |\tau \circledast \sigma|$ into:

- (1) $\Pi_\sigma(p) = (1, m)$ with $m \in |\sigma|$, and $\Pi_\tau(p) = (1, a)$ with $a \in |A|$,
- (2) $\Pi_\sigma(p) = (1, m)$ with $m \in |\sigma|$, and $\Pi_\tau(p) = (2, n)$ with $n \in |\tau|$,
- (3) $\Pi_\sigma(p) = (2, c)$ with $c \in |C|$, and $\Pi_\tau(p) = (2, n)$ with $n \in |\tau|$.

In case (1), the only relevant projection is $\Pi_\sigma(p) = (1, m)$ as $\Pi_\tau(p) = \partial_\sigma(m)$. We write $p_\sigma = m$ and p_τ is undefined, and we say that p **occurs in** A . In case (3), the only relevant projection is $\Pi_\tau(p) = (2, n)$ as $\Pi_\sigma(p) = \partial_\tau(n)$. We write $p_\tau = n$ and p_σ is undefined, and we say that p **occurs in** C . Finally, in case (2) the two projections $\Pi_\sigma(p) = (1, m)$ and $\Pi_\tau(p) = (2, n)$ are relevant, but we must have $\partial_\sigma(m) = (2, b)$ and $\partial_\tau(p) = (1, b)$ for some $b \in |B|$. We write $p_\sigma = m$, $p_\tau = n$, we say that p **occurs in** B and that p is **synchronized**.

The definition of composition consists simply in removing all synchronized events:

Definition 4.9. The **composition** of $\sigma : A \vdash B$ and $\tau : B \vdash C$ comprises components:

$$\begin{aligned} |\tau \odot \sigma| &= \{p \in |\tau \circledast \sigma| \mid p \text{ occurs in } A \text{ or } C\}, \\ p_1 \leq_{\tau \odot \sigma} p_2 &\Leftrightarrow p_1 \leq_{\tau \circledast \sigma} p_2, \\ p_1 \#_{\tau \odot \sigma} p_2 &\Leftrightarrow p_1 \#_{\tau \circledast \sigma} p_2, \\ \theta : x \cong_{\tau \odot \sigma} y &\Leftrightarrow \exists \theta' \subseteq \theta' : x' \cong_{\tau \circledast \sigma} y'. \end{aligned}$$

with display map $\partial_{\tau \odot \sigma} : |\tau \odot \sigma| \rightarrow |A \vdash C|$ obtained as restriction of $\partial_{\tau \circledast \sigma}$.

The composition of prestrategies $\sigma : A \vdash B$ and $\tau : B \vdash C$ gives data

$$(|\tau \odot \sigma|, \leq_{\tau \odot \sigma}, \#_{\tau \odot \sigma}, \mathcal{S}(\tau \odot \sigma), \partial_{\tau \odot \sigma})$$

satisfying all the axioms of Definition 4.2 except, possibly, \sim -receptivity¹⁸. When composing prestrategies, we will check \sim -receptivity separately – this only occurs in Section 4.4.3.

However, if σ and τ are *strategies*, then so is $\tau \odot \sigma$ [CCW19]. Composition is associative up to iso (with **isomorphisms** between causal strategies defined as between pre-interactions above). In Figure 35, the composition simply keeps the events in black. This means that the composition has two conflicting positive events, both corresponding to \checkmark^+ : the model records the point of non-deterministic branching even when it brings no observable difference. Though this does not appear in pictures, we insist that events of $\tau \odot \sigma$ are certain events of $\tau \circledast \sigma$. Thus an event of the composition always carries a unique causal explanation: *itself*.

To parallel this event-based definition of composition, there is a state-based characterization. A causally compatible pair $x^\sigma \in \mathcal{C}(\sigma), x^\tau \in \mathcal{C}(\tau)$ is **minimal** if for all causally compatible $y^\sigma \in \mathcal{C}(\sigma), y^\tau \in \mathcal{C}(\tau)$ with $y^\sigma \subseteq x^\sigma, y^\tau \subseteq x^\tau$ with $x_A^\sigma = y_A^\sigma$ and $x_C^\tau = y_C^\tau$, then $x_B = y_B$. The same definition applies to causally compatible pairs of symmetries.

¹⁸However, \sim -receptivity of σ and τ is required for $\tau \circledast \sigma$ to form an *ess* [CCW19].

Proposition 4.10. *Consider $\sigma : A \vdash B$, and $\tau : B \vdash C$ causal strategies.*

There is a causal strategy $\tau \odot \sigma$, unique up to iso, s.t. there are order-isos:

$$\begin{aligned} (- \odot -) & : \{(x^\tau, x^\sigma) \in \mathcal{C}(\tau) \times \mathcal{C}(\sigma) \mid x^\sigma, x^\tau \text{ minimal causally compatible}\} \simeq \mathcal{C}(\tau \odot \sigma) \\ (- \odot -) & : \{(\theta^\tau, \theta^\sigma) \in \mathcal{S}(\tau) \times \mathcal{S}(\sigma) \mid \theta^\sigma, \theta^\tau \text{ minimal causally compatible}\} \simeq \mathcal{S}(\tau \odot \sigma) \end{aligned}$$

commuting with dom and cod; s.t., for $\theta^\sigma \in \mathcal{S}(\sigma), \theta^\tau \in \mathcal{S}(\tau)$ minimal causally compatible,

$$\partial_{\tau \odot \sigma}(\theta^\tau \odot \theta^\sigma) = \theta_A^\sigma \parallel \theta_C^\tau.$$

The *minimality* requirement amounts to asking the maximal events of x^σ and x^τ to occur in A or C . As events of $\tau \odot \sigma$ carry their causal witness, configurations of $\tau \odot \sigma$ are in one-to-one correspondence with those configurations of $\tau * \sigma$ whose maximal events occur in A or C – thus Proposition 4.10 follows from Proposition 4.8 (see Appendix D.2.2).

In fact, trailing Opponent moves do not matter as they are forced by *receptivity* and *courtesy* to behave as in the game. A configuration $x \in \mathcal{C}(\sigma)$ is **+covered** iff the top elements of x (for \leq_σ) are positive – we write $x \in \mathcal{C}^+(\sigma)$. Likewise, $\theta \in \mathcal{S}(\sigma)$ is **+covered** if $\text{dom}(\theta)$ (or, equivalently, $\text{cod}(\theta)$) is +covered – we write $\theta \in \mathcal{S}^+(\sigma)$. We have [dV20]:

Lemma 4.11. *Consider $\sigma, \tau : A$ two causal strategies. Assume there are*

$$\psi : \mathcal{C}^+(\sigma) \simeq \mathcal{C}^+(\tau) \quad \psi : \mathcal{S}^+(\sigma) \simeq \mathcal{S}^+(\tau)$$

order-isomorphisms compatible with dom, cod, and display maps.

Then, σ and τ are isomorphic.

See Appendix D.2.3 for the proof. Relying on this we can finally prove:

Proposition 4.12. *Consider $\sigma : A \vdash B$ and $\tau : B \vdash C$ causal strategies.*

Then, there is a strategy $\tau \odot \sigma : A \vdash C$, unique up to iso, such that there are order-isos:

$$\begin{aligned} (- \odot -) & : \{(x^\tau, x^\sigma) \in \mathcal{C}^+(\tau) \times \mathcal{C}^+(\sigma) \mid x^\sigma \text{ and } x^\tau \text{ causally compatible}\} \simeq \mathcal{C}^+(\tau \odot \sigma) \\ (- \odot -) & : \{(\theta^\tau, \theta^\sigma) \in \mathcal{S}^+(\tau) \times \mathcal{S}^+(\sigma) \mid \theta^\sigma \text{ and } \theta^\tau \text{ causally compatible}\} \simeq \mathcal{S}^+(\tau \odot \sigma) \end{aligned}$$

commuting with dom and cod, and s.t., for $\theta^\sigma \in \mathcal{S}^+(\sigma)$ and $\theta^\tau \in \mathcal{S}^+(\tau)$ causally compatible,

$$\partial_{\tau \odot \sigma}(\theta^\tau \odot \theta^\sigma) = \theta_A^\sigma \parallel \theta_C^\tau.$$

Proof. Relatively direct from Proposition 4.10 and Lemma D.14, see Appendix D.2.3. \square

This is convenient as +covered configurations of strategies often have a simpler description (see *e.g.* Lemma 4.18). Minimality also disappears as a causally compatible pair of +covered configurations is always minimal (indeed, a synchronized maximal event would be negative for one of the players). This final characterization will be used often to prove equalities between strategies. It is also of great use when linking with the relational model (see Section 7.1), but also for quantitative extensions (see *e.g.* [CCPW18, CdV20]).

4.2.4. *Congruence.* What is the right equivalence between causal (pre)strategies? There are a few options, several investigated in [CCW19]; here we use *positive isomorphism*:

Definition 4.13. Consider $\sigma, \tau : A$ two causal strategies on arena A .

A **positive isomorphism** $\varphi : \sigma \approx \tau$ is an isomorphism of ess satisfying

$$\partial_\tau \circ \varphi \sim^+ \partial_\sigma,$$

i.e. for all $x \in \mathcal{C}(\sigma)$, $\{(\partial_\sigma(s), \partial_\tau \circ \varphi(s)) \mid s \in x\} \in \mathcal{S}_+(A)$: the two maps are **positively symmetric**. In that case we say σ and τ are **positively isomorphic**, and write $\sigma \approx \tau$.

This means that σ and τ are the same up to renaming of their events. This renaming might cause a reindexing of positive events, but it must keep the copy indices of negative events unchanged. Crucially, positive isomorphism is preserved by composition [CCW19]:

Proposition 4.14. *Consider $\sigma, \sigma' : A \vdash B$, $\tau, \tau' : B \vdash C$ s.t. $\sigma \approx \sigma'$ and $\tau \approx \tau'$. Then, we have $\tau \odot \sigma \approx \tau' \odot \sigma'$.*

The proof is fairly elaborate. Without going into details, it will be useful to have in mind the first key step: showing that two (pre)strategies able to synchronize *up to symmetry*, always also have a synchronization *on the nose*. More precisely, we have the following:

Proposition 4.15. *Consider $\sigma : A \vdash B$ and $\tau : B \vdash C$ two causal (pre)strategies. For any $x^\sigma \in \mathcal{C}(\sigma)$, $x^\tau \in \mathcal{C}(\tau)$ and $\theta : x_B^\sigma \cong_B x_B^\tau$ s.t. the composite bijection is secured:*

$$x^\sigma \parallel x_C^\tau \stackrel{\partial_\sigma \parallel C}{\cong} x_A^\sigma \parallel x_B^\sigma \parallel x_C^\tau \stackrel{A \parallel \theta \parallel C}{\cong} x_A^\sigma \parallel x_B^\tau \parallel x_C^\tau \stackrel{A \parallel \partial_\tau^{-1}}{\cong} x_A^\sigma \parallel x^\tau,$$

then there are $y^\sigma \in \mathcal{C}(\sigma)$ and $y^\tau \in \mathcal{C}(\tau)$ causally compatible, along with symmetries

$$\varphi^\sigma : y^\sigma \cong_\sigma x^\sigma, \quad \varphi^\tau : y^\tau \cong_\tau x^\tau,$$

such that $\varphi_B^\tau \circ \theta = \varphi_B^\sigma$.

This follows from Lemma 3.23 in [CCW19]. Intuitively, we play $\mathcal{S}(\sigma)$ and $\mathcal{S}(\tau)$ against each other. By \sim -receptivity and *extension* they adjust their copy indices interactively until reaching an agreement. This is the first step to congruence, but not the only one: the requirement that we should get a global map $\varphi : \tau \odot \sigma \approx \tau' \odot \sigma'$ is in tension with the definition of isomorphism families, which only guarantees a more local bisimulation-like property. The mismatch is compensated by the uniqueness of extensions granted by *thin*, without which congruence fails. Details are out of scope for the present paper [CCW19].

If $\sigma \approx \tau$, there can be in principle multiple $\varphi : \sigma \approx \tau$. We leave these isomorphisms to the background, as we have not yet encountered a computational use for these. If they are retained, then arenas, causal strategies and positive morphisms form a bicategory [Pdq20].

Remark 4.16. In Definition 4.13, one could ask φ to preserve ∂ up to arbitrary symmetry (*weak isomorphism*) or even, to be itself invertible only up to symmetry (*weak equivalence*). This changes the mediating morphisms, but not the resulting equivalence relation between strategies (see Corollary 3.30 in [CCW19]). In this paper we choose positive isomorphism as it seems natural conceptually, and because the additional positivity constraint is useful.

4.2.5. *Copycat.* So as to complete the categorical structure, it remains to define *copycat*.

Definition 4.17. For each $--$ -arena A , the **copycat strategy** $\mathfrak{c}_A : A \vdash A$ is defined as:

$$\begin{aligned} |\mathfrak{c}_A| &= |A \vdash A| \\ \partial_{\mathfrak{c}_A}(i, a) &= (i, a) \\ (i, a) \leq_{\mathfrak{c}_A} (j, a') &\Leftrightarrow a <_A a'; \text{ or } a = a', \text{ pol}_{A \vdash A}(i, a) = - \text{ and } \text{pol}_{A^\perp \parallel A}(j, a') = + \\ (i, a) \#_{\mathfrak{c}_A} (j, a') &\Leftrightarrow a \#_A a', \end{aligned}$$

with symmetries those bijections of the form $\theta_1 \parallel \theta_2 : x_1 \parallel x_2 \cong_{\mathfrak{c}_A} y_1 \parallel y_2$ such that

$$\theta_1 : x_1 \cong_A y_1, \quad \theta_2 : x_2 \cong_A y_2, \quad \text{and} \quad \theta_1 \cap \theta_2 : x_1 \cap x_2 \cong_A y_1 \cap y_2.$$

This simplifies the usual definition [CCW19], exploiting the particular shape of arenas. Its immediate causal links import \rightarrow_A on either side, along with all the $(i, a) \rightarrow_{\mathbf{c}_A} (j, a)$ when $\text{pol}_{A \vdash A}(i, a) = -$ and $\text{pol}_{A \vdash A}(j, a) = +$. In other words, \mathbf{c}_A is an asynchronous forwarder: it is prepared to play any positive event on one side, under the condition that the corresponding negative event appears first on the other side. Its symmetries are inherited from $\mathcal{S}(A \vdash A)$, with the constraint that they should agree on events already forwarded.

Perhaps the simplest description of copycat is through its completely forwarded states:

Lemma 4.18. *Consider A any $--$ -arena. Then, we have:*

$$\begin{aligned} \mathcal{C}^+(\mathbf{c}_A) &= \{x_A \parallel x_A \in \mathcal{C}(A \parallel A) \mid x_A \in \mathcal{C}(A)\} \\ \mathcal{S}^+(\mathbf{c}_A) &= \{\theta_A \parallel \theta_A \in \mathcal{S}(A \parallel A) \mid \theta_A \in \mathcal{S}(A)\}. \end{aligned}$$

Proof. Straightforward. \square

This foreshadows the link with relational semantics in Section 7.1: when restricted to $+$ -covered configurations, copycat looks like the identity relation. We may deduce:

Proposition 4.19. *Composition is associative up to \approx on prestrategies. For any $\sigma : A \vdash B$,*

$$\mathbf{c}_B \odot \sigma \odot \mathbf{c}_A \approx \sigma,$$

so that $--$ -arenas and causal strategies form a category.

Proof. Associativity follows from Proposition 4.10 and a ternary version of causal compatibility – see also [CCW19] for a detailed proof via the universal property of the interaction pullback. For neutrality of copycat, there is an order-isomorphism preserving display maps

$$\begin{aligned} \mathcal{C}^+(\mathbf{c}_B \odot \sigma) &\simeq \{(x^{\mathbf{c}_B}, x^\sigma) \in \mathcal{C}^+(\mathbf{c}_B) \times \mathcal{C}^+(\sigma) \mid x^{\mathbf{c}_B} \text{ and } x^\sigma \text{ causally compatible}\} \\ &\simeq \{(x^{\mathbf{c}_B}, x^\sigma) \in \mathcal{C}^+(\mathbf{c}_B) \times \mathcal{C}^+(\sigma) \mid x^{\mathbf{c}_B} \text{ and } x^\sigma \text{ matching}\} \\ &\simeq \{(x_B^\sigma \parallel x_B^\sigma, x^\sigma) \mid x^\sigma \in \mathcal{C}^+(\sigma)\} \\ &\simeq \mathcal{C}^+(\sigma), \end{aligned}$$

using first Proposition 4.12; verifying directly that securedness always holds when composing with copycat; using Lemma 4.18. The same reasoning can be made with symmetries, concluding that $\mathbf{c}_B \odot \sigma$ and σ are isomorphic by uniqueness in Proposition 4.12. \square

Before we develop further this categorical structure, we introduce a few useful lemmas.

4.2.6. Immediate causality in interactions. Later on, we will need some tools to reason on the causality in $\tau \circledast \sigma$ and how it relates to that in σ and τ .

Lemma 4.20. *For $\sigma : A \vdash B, \tau : B \vdash C$ causal prestrategies, for $m, m' \in |\tau \circledast \sigma|$, if $m \rightarrow_{\tau \circledast \sigma} m'$, then $m_\sigma \rightarrow_\sigma m'_\sigma$, or $m_\tau \rightarrow_\tau m'_\tau$, where m_σ, m_τ are defined whenever used.*

The proof is in Appendix D.3. So in the event-based view of interaction, immediate causal links originate in one of the components. For σ and τ strategies, one can track down the responsible component via a polarity analysis. Of course, it is usual in game semantics that events of $\tau \circledast \sigma$ cannot sensibly be assigned a polarity in $\{-, +\}$, because σ and τ disagree on B . A more useful notion of polarity is $\text{pol}_{\tau \circledast \sigma} : |\tau \circledast \sigma| \rightarrow \{-, l, r\}$ given by:

$$\begin{aligned} \text{pol}_{\tau \circledast \sigma}(m) &= l && \text{if } m_\sigma \text{ is defined and } \text{pol}_\sigma(m_\sigma) = +, \\ \text{pol}_{\tau \circledast \sigma}(m) &= r && \text{if } m_\tau \text{ is defined and } \text{pol}_\tau(m_\tau) = +, \\ \text{pol}_{\tau \circledast \sigma}(m) &= - && \text{otherwise.} \end{aligned}$$

As an example, we show in Figure 35 the polarities arising from this definition. Then:

Lemma 4.21. Consider $\sigma : A \vdash B$ and $\tau : B \vdash C$ strategies, and $m \rightarrow_{\tau \otimes \sigma} m'$. Then,

- (1) if $\text{pol}_{\tau \otimes \sigma}(m') = l$, then $m_\sigma \rightarrow_\sigma m'_\sigma$,
- (2) if $\text{pol}_{\tau \otimes \sigma}(m') = r$, then $m_\tau \rightarrow_\tau m'_\tau$,
- (3) if $\text{pol}_{\tau \otimes \sigma}(m') = -$, then $\partial_{\tau \otimes \sigma}(m) \rightarrow_{A \parallel B \parallel C} \partial_{\tau \otimes \sigma}(m')$.

Proof. (1) By Lemma 4.20, m_σ, m'_σ defined and $m_\sigma \rightarrow_\sigma m'_\sigma$ – in which case we are done; or m_τ, m'_τ defined and $m_\tau \rightarrow_\tau m'_\tau$. Since $\text{pol}_{\tau \otimes \sigma}(m') = l$, $\text{pol}_\tau(m'_\tau) = -$. By courtesy, $\partial_\tau(m_\tau) \rightarrow_{B \vdash C} \partial_\tau(m'_\tau)$; hence m occurs in B and $\partial_\sigma(m_\sigma) \rightarrow_{A \vdash B} \partial_\sigma(m'_\sigma)$. By Lemma A.2, $m_\sigma <_\sigma m'_\sigma$, and the causality must be immediate by Lemma D.17. (2) is symmetric.

(3) Assume m' occurs in A , the other case is symmetric. In that case only m'_σ is defined, so Lemma 4.20 entails that m_σ is defined and $m_\sigma \rightarrow_\sigma m'_\sigma$. But $\text{pol}_\sigma(m'_\sigma) = -$, so by courtesy $\partial_\sigma(m_\sigma) \rightarrow_{A \vdash B} \partial_\sigma(m'_\sigma)$, from which the conclusion follows. \square

4.3. Seely Category. Now, we turn to the different components of a Seely category.

4.3.1. Symmetric monoidal category with products. On $--$ -arenas, we keep the definitions for $\downarrow\uparrow\text{-Strat}$, enriched as in Section 4.1.2. The tensor of causal strategies is defined below:

Definition 4.22. For $\sigma_1 : A_1 \vdash B_1$, $\sigma_2 : A_2 \vdash B_2$ causal strategies between $--$ -arenas, then

$$\sigma_1 \otimes \sigma_2 : A_1 \otimes A_2 \vdash B_1 \otimes B_2$$

is defined as the ess $\sigma_1 \parallel \sigma_2$ along with display map $\partial_{\sigma_1 \otimes \sigma_2}(i, s) = (j, (i, a))$ if $\sigma_i(s) = (j, a)$.

Bifactoriality is direct via Proposition 4.12. The symmetric monoidal structural isomorphisms are provided by copycat strategies, only changing display maps:

$$\begin{array}{ll} \alpha_{A,B,C} : (A \otimes B) \otimes C \cong A \otimes (B \otimes C) & \rho_A : A \otimes \mathbf{1} \cong A \\ s_{A,B} : A \otimes B \cong B \otimes A & \lambda_A : \mathbf{1} \otimes A \cong A \end{array}$$

satisfying up to positive iso the expected naturality and coherence laws [CCW19]. For cartesian products, the *projections* $\pi_1 : A_1 \& A_2 \vdash A_1$ and $\pi_2 : A_1 \& A_2 \vdash A_2$ are relabeled copycat strategies, while the *pairing* of causal strategies is defined similarly to the tensor:

Definition 4.23. Consider $\sigma_1 : A \vdash B_1$ and $\tau : A \vdash B_2$ causal strategies between $--$ -arenas. Then, $\langle \sigma_1, \sigma_2 \rangle : A \vdash B_1 \& B_2$ is defined as having ess $\sigma_1 \& \sigma_2$, along with

$$\begin{array}{ll} \partial_{\langle \sigma_1, \sigma_2 \rangle}(i, s) = (1, a) & \text{if } \partial_{\sigma_i}(s) = (1, a) \\ \partial_{\langle \sigma_1, \sigma_2 \rangle}(i, s) = (2, (i, b)) & \text{if } \partial_{\sigma_i}(s) = (2, b). \end{array}$$

It follows from Proposition 4.12 and direct verifications that this yields binary products.

4.3.2. Monoidal closed structure. We now describe the monoidal closure.

On objects, the closure is the arrow $A \multimap B$ from $\downarrow\uparrow\text{-Strat}$. However, for now, the strategies on $A \vdash B$ and on $A \multimap B$ are *not* in one-to-one correspondence. Indeed strategies in $A \multimap B$ include a *pointer* for initial moves in A , while strategies in $A \vdash B$ do not. This pointer is not always unique, as illustrated in Figure 36. To cope with this we could have, as for the play-based strategies of the previous sections, set the morphisms of our category directly as strategies on $A \multimap B$; but in this causal setting that obfuscates composition.

Instead, we restrict to strategies for which this pointer reconstruction is *unique*:




Figure 36: Non-uniqueness of the threading pointer

Definition 4.24. A causal strategy $\sigma : A$ on arena A is **pointed** if for each $s \in |\sigma|$ there is a unique event $\text{init}(s) \in |\sigma|$ which is minimal for \leq_σ and such that $\text{init}(s) \leq_\sigma s$.

Copycat strategies are pointed (as arenas are forestial), and pointed strategies are stable under composition and the other operations on strategies. From now on, we consider that all causal strategies are pointed. We write $\rightarrow\text{-Strat}$ for the category having $--$ -arenas as objects, and as morphisms from A to B , the pointed causal strategies on $A \vdash B$.

For pointed strategies, the missing pointer can always be recovered uniquely:

Lemma 4.25. *Let A, B and C be $--$ -arenas. Then, we have a bijection:*

$$\Lambda_{A,B,C} : \rightarrow\text{-Strat}(A \otimes B, C) \simeq \rightarrow\text{-Strat}(A, B \multimap C)$$

Proof. The bijection only affects the display map, leaving the other components unchanged.

The non-trivial direction is from left to right. Consider $\sigma : A \otimes B \vdash C$. We set:

$$\partial_{\Lambda(\sigma)}(s) = \begin{cases} (1, a) & \text{if } \partial_\sigma(s) = (1, (1, a)), \\ (2, (2, c)) & \text{if } \partial_\sigma(s) = (2, c), \\ (2, (1, (c, b))) & \text{if } \partial_\sigma(s) = (1, (2, b)) \text{ and } \partial_\sigma(\text{init}(s)) = (2, c). \end{cases}$$

It is a direct verification that this yields a bijection as claimed. \square

From this point, we may now easily wrap up the symmetric monoidal closed structure.

Proposition 4.26. *The category $\rightarrow\text{-Strat}$ is symmetric monoidal closed.*

Proof. First, for any $--$ -arenas A and B , we have $A \multimap B$ and an **evaluation**

$$\text{ev}_{A,B} = \Lambda_{A \multimap B, A, B}^{-1}(\mathbf{c}_{A \multimap B}) : (A \multimap B) \otimes A \vdash B,$$

and given $\sigma : A \otimes B \vdash C$, $\text{ev}_{B,C} \odot (\Lambda_{A,B,C}(\sigma) \otimes B) \approx \sigma$ follows from a variation over the neutrality of copycat for composition. From there, the universal property is routine. \square

4.3.3. *The exponential.* The first step is to introduce a functor $! : \rightarrow\text{-Strat} \rightarrow \rightarrow\text{-Strat}$.

Definition 4.27. Consider A and B two $--$ -arenas, and $\sigma : A \vdash B$ a causal strategy.

We define a strategy $!\sigma : !A \vdash !B$ with $!\sigma$ as event structure with symmetry and:

$$\partial_{!\sigma}(i, m) = \begin{cases} (1, (i, a)) & \text{if } \partial_\sigma(m) = (1, a), \\ (2, (i, b)) & \text{if } \partial_\sigma(m) = (2, b). \end{cases}$$

It is a direct verification that this defines a causal strategy, and functoriality is proved as for the tensor product of strategies. To complete the categorical structure, we have:

Proposition 4.28. *The category $\rightarrow\text{-Strat}$ is a Seely category.*

Proof. The structure presented above is completed by structural natural families of strategies:

$$\begin{array}{ll} \mathbf{dig}_A : !A \rightarrow !!A & \mathbf{mon}_{A,B}^2 : !A \otimes !B \cong !(A \& B) \\ \mathbf{der}_A : !A \rightarrow A & \mathbf{mon}^0 : \mathbf{1} \cong !\mathbf{1}. \end{array}$$

making $(!, \mathbf{dig}, \mathbf{der})$ a comonad along with the *Seely isomorphisms*. Those are all relabeled copycat strategies: for instance, \mathbf{dig}_A is $\mathbf{c}_{!!A}$ relabeled on the left hand side following a bijection $\mathbb{N} \times \mathbb{N} \simeq \mathbb{N}$, \mathbf{der}_A is \mathbf{c}_A relabeled to set events on the left hand side to copy index 0, etc. The naturality and coherence are easily verified, exploiting again Proposition 4.12. \square

4.3.4. *Extracting plays.* In Section 4.1.4, we unfolded causal strategies to non-alternating strategies. Here, we show that this is compatible with the categorical operations.

First, we extend the definition in Proposition 4.3 for causal strategies *from A to B*.

Definition 4.29. For A and B two $--$ arenas, and $\sigma : A \vdash B$ a causal strategy, we define

$$\mathcal{U}\text{-Unf}(\sigma) = \partial_{\Lambda(\sigma)}(\mathcal{U}\text{-Plays}(\sigma)) \in \mathcal{U}\text{-Strat}(A, B)$$

exploiting that σ and $\Lambda(\sigma)$ only differ via their display map.

This matches applying Proposition 4.3 to $\Lambda(\sigma) : A \multimap B$ obtained by monoidal closure.

Proposition 4.30. *There is a symmetric monoidal closed $\mathcal{U}\text{-Unf}(-) : \rightarrow\text{-Strat} \rightarrow \mathcal{U}\text{-Strat}$.*

Proof. For identities, the definition of plays of the asynchronous copycat (Definition C.4) follows the characterisation of configurations of \mathbf{c}_A found *e.g.* in Lemma 3.11 in [CCRW17].

For composition, take $\sigma : A \vdash B$ and $\tau : B \vdash C$. Though $\tau \circledast \sigma$ is not an esp, Definition 3.18 generalizes transparently to $\mathcal{U}\text{-Plays}(\tau \circledast \sigma)$. There are two inclusions to check:

$\mathcal{U}\text{-Unf}(\tau \circledast \sigma) \subseteq \mathcal{U}\text{-Unf}(\tau) \circledast \mathcal{U}\text{-Unf}(\sigma)$: any $s \in \mathcal{U}\text{-Unf}(\tau \circledast \sigma)$ has the form $\partial_{\Lambda(\tau \circledast \sigma)}(t)$ for $t \in \mathcal{U}\text{-Plays}(\tau \circledast \sigma)$, which in turn can be completed to $v \in \mathcal{U}\text{-Plays}(\tau \circledast \sigma)$. Then v may be displayed to $u \in \mathcal{U}\text{-Unf}(\tau) \circledast \mathcal{U}\text{-Unf}(\sigma)$, witnessing $s \in \mathcal{U}\text{-Unf}(\tau) \circledast \mathcal{U}\text{-Unf}(\sigma)$.

$\mathcal{U}\text{-Unf}(\tau) \circledast \mathcal{U}\text{-Unf}(\sigma) \subseteq \mathcal{U}\text{-Unf}(\tau \circledast \sigma)$: any $s \in \mathcal{U}\text{-Unf}(\tau) \circledast \mathcal{U}\text{-Unf}(\sigma)$ has a witness $u \in \mathcal{U}\text{-Unf}(\tau) \circledast \mathcal{U}\text{-Unf}(\sigma)$, projecting to $u \upharpoonright A, B \in \mathcal{U}\text{-Unf}(\sigma)$ and $u \upharpoonright B, C \in \mathcal{U}\text{-Unf}(\tau)$. Those are respectively $\partial_{\Lambda(\sigma)}(s^\sigma)$ and $\partial_{\Lambda(\tau)}(s^\tau)$ for $s^\sigma \in \mathcal{U}\text{-Plays}(\sigma)$ and $s^\tau \in \mathcal{U}\text{-Plays}(\tau)$. Then $x^\sigma := |s^\sigma|$ and $x^\tau := |s^\tau|$ are causally compatible as u induces a linearization of the corresponding bijection. By construction, $x^\tau \circledast x^\sigma$ has a linearization v that displays to u ; and its restriction to visible events yields $t \in \mathcal{U}\text{-Plays}(\tau \circledast \sigma)$ that displays to s .

The preservation of the monoidal structure is direct; the functor is strict monoidal. \square

As $\mathcal{U}\text{-Strat}$ does not handle symmetry, it supports no equivalence relation corresponding to \approx . Nevertheless, this lets us import the stability of well-bracketing under composition. We first generalize Definition 4.4 to well-bracketed causal strategies between $--$ arenas:

Definition 4.31. Consider $--$ arenas A, B , and $\sigma : A \vdash B$ a causal strategy.

We say that $\sigma : A \vdash B$ is **well-bracketed** iff $\mathcal{U}\text{-Unf}(\sigma)$ is well-bracketed.

From Propositions 3.22 and 4.30, there is an smcc with products $\rightarrow\text{-Strat}^{\text{wb}}$ of $--$ arenas and well-bracketed causal strategies. Finally, $!(-)$ preserves well-bracketing and the other components for the exponential are well-bracketed, so $\rightarrow\text{-Strat}^{\text{wb}}$ extends to a Seely category.

4.4. **Interpretation of \mathbf{IA}_{\parallel} .** We now describe the interpretation of \mathbf{IA}_{\parallel} in $\rightarrow\text{-Strat}_!^{\text{wb}}$. First the *types* of \mathbf{IA}_{\parallel} are the same as those of \mathbf{IA} ; their interpretation does not change.

4.4.1. *Interpretation of core PCF.* We focus on the core PCF primitives, postponing **let**.

The λ -calculus primitives are interpreted following the cartesian closed structure of $\rightarrow\text{-Strat}_!^{\text{wb}}$. For constants, we use the obvious strategies returning the corresponding value. For basic PCF combinators, there are obvious causal strategies corresponding to Figures 12 and 14. For recursion, we must first define a partial order on causal strategies:

Definition 4.32. Consider A an arena, and $\sigma, \tau : A$ causal (pre)strategies.

We write $\sigma \leq \tau$ iff $\mathcal{C}(\sigma) \subseteq \mathcal{C}(\tau)$ – so $|\sigma| \subseteq |\tau|$ as well – with, additionally:

- (1) for all $s_1, s_2 \in |\sigma|$, $s_1 \leq_{\sigma} s_2$ iff $s_1 \leq_{\tau} s_2$,
- (2) for all $s_1, s_2 \in |\sigma|$, $s_1 \#_{\sigma} s_2$ iff $s_1 \#_{\tau} s_2$,
- (3) for all $x, y \in \mathcal{C}(\sigma)$ and bijection $\theta : x \simeq y$, we have $\theta \in \mathcal{S}(\sigma)$ iff $\theta \in \mathcal{S}(\tau)$,
- (4) for all $s \in |\sigma|$, $\partial_{\sigma}(s) = \partial_{\tau}(s)$,

i.e. all components compatible with the inclusion.

Causal strategies on A , ordered by \leq , form a *directed complete partial order*; however without a least element. Indeed, strategies minimal for \leq still have – by receptivity – events corresponding to the minimal events of A , but those are *named* arbitrarily. We solve this as in [CCW19]: we choose one minimal causal strategy $\perp_A : A$ with events *exactly* those negative minimal in A ; induced causality, conflict, and isomorphism family; and as display map the identity. For any $\sigma : A$, we pick $\sigma^b \cong \sigma : A$ such that $\perp_A \leq \sigma^b$, obtained by renaming minimal events. We write \mathcal{D}_A for the pointed dcpo of causal strategies above \perp_A .

All operations on strategies examined so far are continuous. So, the operation

$$\begin{aligned} F & : \mathcal{D}_{\mathbf{1} \vdash (A \rightarrow A) \rightarrow A} \rightarrow \mathcal{D}_{\mathbf{1} \vdash (A \rightarrow A) \rightarrow A} \\ & \sigma \mapsto (\lambda f^{A \rightarrow A}. f(\sigma f))^b, \end{aligned}$$

in λ -calculus syntax following the cartesian closed structure of $\rightarrow\text{-Strat}_!^{\text{wb}}$, is continuous. Thus it has a least fixed point $\mathcal{Y}_A \in \mathcal{D}_{\mathbf{1} \vdash (A \rightarrow A) \rightarrow A}$, *i.e.* such that $\mathcal{Y}_A = F(\mathcal{Y}_A)$.

4.4.2. *Interpretation of **let**.* We give the interpretation of parallel **let**, as

$$\llbracket \Gamma \vdash \mathbf{let} \begin{pmatrix} x_1 & = & N_1 \\ x_2 & = & N_2 \end{pmatrix} \mathbf{in} M : \mathbb{Y} \rrbracket = \mathbf{plet}_{\mathbb{X}, \mathbb{Y}} \odot! \langle \llbracket N_1 \rrbracket, \llbracket N_2 \rrbracket, \Lambda_{\mathbb{X} \& \mathbb{X}}^!(\llbracket M \rrbracket) \rangle.$$

where $\mathbf{plet}_{\mathbb{X}, \mathbb{Y}} \in \rightarrow\text{-Strat}_!^{\text{wb}}(\mathbb{X} \& \mathbb{X} \& ((\mathbb{X} \& \mathbb{X}) \rightarrow \mathbb{Y}), \mathbb{Y})$ first evaluates its two arguments \mathbb{X} in parallel. Once they both terminate on v and w , it calls its function argument, with $\langle v, w \rangle$.

More formally, we first define a prestrategy “forcing” evaluation to v, w , *i.e.*

$$\mathbf{force}_{v, w} : !(\mathbb{X} \& \mathbb{X} \& ((\mathbb{X} \& \mathbb{X}) \rightarrow \mathbb{Y})) \vdash \mathbb{Y}$$



as in Figure 37 – only a prestrategy, not receptive to other values. Likewise,

$$\mathbf{eval}_{v, w} : !(\mathbb{X} \& \mathbb{X} \& ((\mathbb{X} \& \mathbb{X}) \rightarrow \mathbb{Y})) \vdash \mathbb{Y}$$

is defined by first using the cartesian closed structure of $\rightarrow\text{-Strat}_!^{\text{wb}}$ to obtain a strategy

$$\lambda \langle x_1, x_2, f \rangle. f \langle v, w \rangle \in \rightarrow\text{-Strat}_!^{\text{wb}}(\mathbb{X} \& \mathbb{X} \& ((\mathbb{X} \& \mathbb{X}) \rightarrow \mathbb{Y}), \mathbb{Y}),$$

evaluating f on $\langle v, w \rangle$, from which $\mathbf{eval}_{v, w}$ is obtained simply by removing the initial \mathbf{q}^- . Up to reindexing, we assume $\mathbf{eval}_{v, w}$ does not use copy indices 0 and 1 on the calls to context.

Figure 37: The prestrategy $\text{force}_{v,w}$ Figure 38: The strategy $\text{plet}_{X,Y}$

Finally, this lets us define plet via the expression (the supremum refers to \leq):

$$\text{plet}_{X,Y} = \bigvee_{v,w \in X} \text{force}_{v,w} \cdot \text{eval}_{v,w} \in \rightarrow\text{-Strat}_!^{\text{wb}}(X \& X \& ((X \& X) \rightarrow Y), Y)$$

where the concatenation \cdot sets both maximal (negative) events of $\text{force}_{v,w}$ as dependencies for $\text{eval}_{v,w}$. The full strategy is represented in Figure 38. Unlike for Figure 37, the picture in Figure 38 is to be read as a symbolic representation. The concrete strategy $\text{plet}_{X,Y}$ has patterns as in Figure 38 for all concrete values for v, w, u and copy indices i and j .

The interpretation of the sequential let can be obtained as a simplification.

4.4.3. Interpretation of interference. Now, we interpret shared state and semaphores. By and large, it closely follows the sequential interpretation of Sections 3.4.2 and 3.4.3.

For the primitives interacting with memory and semaphores, we use the same definitions as in Section 3.4.2, with the obvious causal strategies matching Figures 24, 25, 26, and 27.

For new references and semaphores, we regard the alternating strategy cell_n from Section 3.4.3 as a (sequential) event structure: its events $|\text{cell}_n|$ are the non-empty plays, the causal order is given by prefix, the conflicting pairs are all non-comparable plays. The display map

$$\begin{aligned} \partial_{\text{cell}_n} : |\text{cell}_n| &\rightarrow !\text{ref} \\ sa &\mapsto a \end{aligned}$$



keeps the last move. As configurations are sets of prefixes of a given play, we set $\mathcal{S}(\text{cell}_n)$ to comprise those bijections induced by plays $s_1 \cong_A s_2$ symmetric on A (see Definition 2.11).

We display in Figure 39 a few early moves of cell_0 , with the convention that all moves in the same row are in pairwise conflict. In this diagram we observe that cell_0 fails *courtesy*: indeed, the immediate causal link $\checkmark_2^+ \rightarrow r_8^-$, for instance, would be illegal for a strategy.

Note that although $\text{cell}_n : !\text{ref}$ is only a prestrategy and not a strategy, we have:

Proposition 4.33. *For all $\sigma : !\text{ref} \vdash A$, for all $n \in \mathbb{N}$, $\sigma \odot \text{cell}_n : A$ is a strategy.*

Proof. First, \sim -receptivity is established directly, exploiting that σ and cell_n do not have non-courteous immediate causal links across components – they are *componentwise courteous* in the sense of [CCW19]. Lemma 3.36 from [CCW19] ensures that $\sigma \odot \text{cell}_n$ is \sim -receptive.

Figure 39: Beginning of \mathbf{cell}_0 Figure 40: Beginning of \mathbf{lock}_0

It remains to prove $\sigma \odot \mathbf{cell}_n$ receptive and courteous. But those properties are independent of symmetry. Without symmetry, strategies are exactly those prestrategies invariant (up to iso) under composition with copycat [CCRW17]. But then, we compute:

$$\begin{aligned} \mathfrak{c}_A \odot (\sigma \odot \mathbf{cell}_n) &\cong (\mathfrak{c}_A \odot \sigma) \odot \mathbf{cell}_n \\ &\cong \sigma \odot \mathbf{cell}_n, \end{aligned}$$

so by Theorem 3.20 of [CCRW17], $\sigma \odot \mathbf{cell}_n$ is receptive and courteous as required. \square

For semaphores, using the same recipe we obtain a prestrategy $\mathbf{lock}_n : \mathbf{!sem}$ for each $n \in \mathbb{N}$. A beginning of \mathbf{lock}_0 is displayed in Figure 40. Note that in this diagram, trailing Opponent moves are indeed maximal, Player has no response – an attempt to release a lock that has not been grabbed, or to grab a lock that has not been released, triggers no response.

The interpretation of new references and semaphores is defined as in Section 3.4.3.

4.5. Adequacy. Adequacy could be deduced from the connection with [GM08] – see Section 4.6.2. Instead we give an independent proof, as we believe it helps build the operational intuitions for the model. Rather than as usual relying on logical relations, we follow an alternative route, proving first adequacy for certain finitary terms in which the correspondence between operational and game semantics is more concrete.

4.5.1. Canonical adequacy. The sharpest link between operational and game semantics holds factoring out recursion, higher-order, and dynamic generation of semaphores or references.

We temporarily extend $\mathbf{IA}_{//}$ with an explicit $\perp_A : A$ for every type A , interpreted by a minimal strategy with no Player move. Take Σ an **interference context**, *i.e.* of the form

$$\Sigma = \ell_1 : \mathbf{ref}, \dots, \ell_i : \mathbf{ref}, \ell_{i+1} : \mathbf{sem}, \dots, \ell_n : \mathbf{sem},$$

and $\Sigma \vdash M : \mathbb{X}$ where $\mathbb{X} \in \{\mathbb{U}, \mathbb{B}, \mathbb{N}\}$. We say M is in **canonical form** if it contains no fixpoint, subterm of higher type, bad references or semaphores, **newref** or **newsem**, with no store-independent reductions available (other than the *interfering reductions* of Figure 2).

To help us link operational and denotational semantics for canonical forms, we introduce a few concepts. First, if s is a store with $\text{dom}(s) = \Sigma$ as above, we interpret it as

$$\llbracket s \rrbracket = \left(\bigotimes_{1 \leq k \leq i} \mathbf{cell}_{s(\ell_k)} \right) \otimes \left(\bigotimes_{i+1 \leq k \leq n} \mathbf{lock}_{s(\ell_k)} \right) : \left(\bigotimes_{1 \leq k \leq i} \mathbf{!ref} \right) \otimes \left(\bigotimes_{i+1 \leq k \leq n} \mathbf{!sem} \right) \cong \llbracket \Sigma \rrbracket.$$

To track evolution of an interaction we use the notion of *residuals*.

Definition 4.34. If E is an ess and $x \in \mathcal{C}(E)$, the **residual** E/x has

$$\begin{aligned} |E/x| &= \{e \in |E| \setminus x \mid \forall e' \in x, \neg(e \#_E e')\}, \\ e_1 \leq_{E/x} e_2 &\Leftrightarrow e_1 \leq_E e_2, \\ e_1 \#_{E/x} e_2 &\Leftrightarrow e_1 \#_E e_2, \\ \theta : y \cong_{E/x} z &\Leftrightarrow \theta \cup \text{id}_x : x \cup y \cong_E x \cup z. \end{aligned}$$

If E has polarities they are preserved as well and for A an arena, A/x is an arena. In particular for $!\llbracket \Sigma \rrbracket$, by definition and playing Hilbert's hotel, for any $x \in \mathcal{C}(!\llbracket \Sigma \rrbracket)$ with as many Player as Opponent moves, we have $(!\llbracket \Sigma \rrbracket)/x \cong !\llbracket \Sigma \rrbracket$. Residuals also apply to causal (pre)strategies: if $\sigma : A$, then for each $x \in \mathcal{C}(\sigma)$, we have $\sigma/x : A/\partial_\sigma(x)$ a prestrategy.

Take $\Sigma \vdash M : \mathbb{X}$ in canonical form. Necessarily $\llbracket \mathbb{X} \rrbracket$ – also written \mathbb{X} – has a unique minimal move \mathbf{q}^- and $\llbracket M \rrbracket$ has a unique matching minimal move, also written \mathbf{q}^- . We set

$$\langle M \rangle = \llbracket M \rrbracket / \{\mathbf{q}^-\} : \Sigma \vdash \langle \mathbb{X} \rangle,$$

where $\langle \mathbb{X} \rangle = \llbracket \mathbb{X} \rrbracket / \{\mathbf{q}^-\}$, yielding a causal prestrategy in the sense of Definition 4.2.

Now, we are equipped to state the most central ingredient of our proof of adequacy.

Lemma 4.35. Consider $\Sigma \vdash M : \mathbb{X}$ in canonical form, and s a store with $\text{dom}(s) = \Sigma$.

Then, there is a (necessarily interfering) one-step reduction

$$\langle M, s \rangle \rightsquigarrow \langle M', s' \rangle$$

iff there are matching $x \in \mathcal{C}(\langle M \rangle)$ and $y \in \mathcal{C}(\llbracket s \rrbracket)$ with two elements each, such that

$$\langle M \rangle / x \approx \langle M' \rangle : !\llbracket \Sigma \rrbracket \vdash \langle \mathbb{X} \rangle, \quad \llbracket s \rrbracket / y \approx \llbracket s' \rrbracket : !\llbracket \Sigma \rrbracket.$$

Proof. For interfering operations $\ell := v, !\ell, \mathbf{grab}(\ell)$ or $\mathbf{release}(\ell)$, it is a direct verification by definition of the interpretation. The result then follows by induction on M . \square

This identifies *store operations* that a canonical M may perform immediately with store s , with the minimal events of $\langle M \rangle \otimes \llbracket s \rrbracket$ that occur in Σ . It almost suffices to iterate this to obtain adequacy for canonical terms; however interfering reductions might yield non-canonical terms, so state operations must be interleaved with pure reductions. Write $M \triangleright N$ for *pure* reductions, *i.e.* the context closure of interference-independent reductions in Figure 2 – these reductions leave invariant the interpretation as causal strategies. Moreover:

Lemma 4.36. Consider $\Sigma \vdash M : \mathbb{X}$ without fixpoint, subterm of higher type, bad references or semaphores, **newref** or **newsem**. Then, there exists $\Sigma \vdash N : \mathbb{X}$ canonical with $M \triangleright^* N$.

Moreover, for any store s with $\text{dom}(s) = \Sigma$, there is $\langle M, s \rangle \rightsquigarrow^* \langle \mathbf{skip}, s' \rangle$ iff there is $\langle N, s \rangle \rightsquigarrow^* \langle \mathbf{skip}, s' \rangle$, and the correspondence preserves the number of interfering operations.

Proof. A routine standardization argument. \square

Using this, we can prove adequacy for canonical terms. Intuitively, a sequence $\langle M, s \rangle \rightsquigarrow^* \langle \mathbf{skip}, s' \rangle$ can be reproduced semantically: each interfering reduction yields by Lemma 4.35 a pair of events in $\langle M \rangle \otimes \llbracket s \rrbracket$, while state-free reductions leave the interpretation unchanged. Reciprocally, a successful interaction in $\langle M \rangle \odot \llbracket s \rrbracket$ is a partially ordered set of memory operations, which may be linearized by securedness; informing a reduction sequence.

Proposition 4.37. Consider $\Sigma \vdash M : \mathbb{U}$ in canonical form, and s a store with $\text{dom}(s) = \Sigma$.

Then, $\langle M, s \rangle \rightsquigarrow^* \langle \mathbf{skip}, s' \rangle$ iff $\langle M \rangle \odot \llbracket s \rrbracket$ has a positive move.

Proof. If. Taking $x \odot y \in \langle M \rangle \otimes \llbracket s \rrbracket$, we build the sequence by induction on the size of $x \otimes y$.

If $x \otimes y$ has exactly one event, it must match \checkmark^+ in $\langle \mathbb{U} \rangle$. Since M is canonical, by a case inspection on the shape of M , the only case for which $\langle M \rangle$ has a minimal positive event in $\langle \mathbb{U} \rangle$ is a value **skip**. Otherwise, since x and y are causally compatible, there is a sequence

$$(\emptyset, \emptyset) \prec (x_1, y_1) \prec \dots \prec (x_n, y_n) = (x, y)$$

of one-step extensions (for the product inclusion order) of matching $x_i \in \mathcal{C}(\langle M \rangle)$ and $y_i \in \mathcal{C}(\llbracket s \rrbracket)$, obtained by linearization from the acyclicity of causal compatibility of x and y . In particular, $x_1 = \{m_1^+\}$ and $y_1 = \{n_1^-\}$ singleton sets. By Proposition 4.12, y is $+$ -covered, so there is a (unique, by definition of $\llbracket s \rrbracket$) $n_1^- \rightarrow_{\llbracket s \rrbracket} n_2^+$, and by Lemma A.5 there is a unique matching $m_1^+ \rightarrow_{\langle M \rangle} m_2^-$. Since n_2 only depends on n_1 and m_2 only depends on m_1 , *w.l.o.g.* $x_2 = \{m_1, m_2\}$ and $y_2 = \{n_1, n_2\}$. So, by Lemma 4.35, there is a one-step $\langle M, s \rangle \rightsquigarrow \langle M', s' \rangle$ s.t. $\langle M \rangle / x_2 \approx \langle M' \rangle$ and $\llbracket s \rrbracket / y_2 \approx \llbracket s' \rrbracket$. We may now use Lemma 4.36 to obtain $M' \triangleright^* M''$ with M'' canonical and $\langle M' \rangle \approx \langle M'' \rangle$. Now, removing x_2 and y_2 to the sequence above yields

$$(\emptyset, \emptyset) \prec (x'_3, y'_3) \prec \dots \prec (x'_n, y'_n) = (x', y')$$

with $x' = x \setminus \{m_1, m_2\}$ and $y' = y \setminus \{n_1, n_2\}$, witnessing that matching $x' \in \mathcal{C}(\langle M'' \rangle)$ and $y' \in \mathcal{C}(\llbracket s' \rrbracket)$ are causally compatible, so that $x' \otimes y' \in \mathcal{C}(\langle M'' \rangle \otimes \llbracket s' \rrbracket)$. But x' and y' are $+$ -covered, so $x' \odot y' \in \mathcal{C}(\langle M'' \rangle \odot \llbracket s' \rrbracket)$, still with a positive move in \mathbb{U} . So $\langle M'', s' \rangle \rightsquigarrow^* \langle \mathbf{skip}, s'' \rangle$ by induction hypothesis, so $\langle M', s' \rangle \rightsquigarrow^* \langle \mathbf{skip}, s'' \rangle$ by Lemma 4.36, so $\langle M, s \rangle \rightsquigarrow^* \langle \mathbf{skip}, s'' \rangle$.

Only if. By induction on the number of interfering operations in $\langle M, s \rangle \rightsquigarrow^* \langle \mathbf{skip}, s'' \rangle$.

If $M = \mathbf{skip}$ it is immediate. Otherwise, consider $\langle M, s \rangle \rightsquigarrow \langle M', s' \rangle \rightsquigarrow^* \langle \mathbf{skip}, s'' \rangle$. By Lemma 4.35, there are $x \in \mathcal{C}(\langle M \rangle)$ and $y \in \mathcal{C}(\llbracket s \rrbracket)$ matching with two elements each, with

$$\langle M \rangle / x \approx \langle M' \rangle \quad \llbracket s \rrbracket / y \approx \llbracket s' \rrbracket.$$

Now, to use the induction hypothesis we convert $M' \triangleright^* M''$ to canonical form; by Lemma 4.36 we have $\langle M'', s' \rangle \rightsquigarrow^* \langle \mathbf{skip}, s'' \rangle$ with the same number of interfering operations; and $\langle M' \rangle \approx \langle M'' \rangle$. So by induction hypothesis, there are $+$ -covered and causally compatible

$$x' \in \mathcal{C}(\langle M'' \rangle), \quad y' \in \mathcal{C}(\llbracket s' \rrbracket),$$

where x' has an occurrence of the positive event of \mathbb{U} ; so up to renaming, $x' \in \mathcal{C}(\langle M \rangle / x)$ and $y' \in \mathcal{C}(\llbracket s \rrbracket / y)$. Therefore, by definition of residuals, one may add back x and y to obtain

$$x \cup x' \in \mathcal{C}(\langle M \rangle), \quad y \cup y' \in \mathcal{C}(\llbracket s \rrbracket),$$

$+$ -covered, causally compatible with a positive move in \mathbb{U} , which concludes the proof. \square

4.5.2. *Finitary adequacy.* We deduce recursion-free adequacy from the canonical case.

To convert terms to canonical form, we perform state-free reductions while pushing declarations of new references or semaphores outside. The latter is done by the commutation rules of Figure 41, from which are missing the three rules for the parallel let, and matching commutations for new semaphores. It is direct that these rules leave the game semantics invariant, and preserve and reflect infinite reduction chains in the operational semantics.

Writing \equiv for the congruence closure of state-free reductions and commutations above:

Lemma 4.38. *Consider $\vdash M : \mathbb{X}$ a recursion-free term of $\mathbf{IA}_{//}$. Then, there exists*

$$M \equiv \mathbf{new} \ x_1 := n_1 \ \mathbf{in} \ \dots \ \mathbf{new} \ x_p := n_p \ \mathbf{in} \ N$$

where each **new** is either **newref** or **newsem**, and $\Sigma \vdash N : \mathbb{X}$ is canonical.

$$\begin{array}{lll}
 (\mathbf{newref} \ x:=n \ \mathbf{in} \ M); N & \rightarrow & \mathbf{newref} \ x:=n \ \mathbf{in} \ M; N & (x \notin \mathbf{fv}(N)) \\
 M; (\mathbf{newref} \ x:=n \ \mathbf{in} \ N) & \rightarrow & \mathbf{newref} \ x:=n \ \mathbf{in} \ M; N & (x \notin \mathbf{fv}(M)) \\
 \mathbf{if} \ (\mathbf{newref} \ x:=n \ \mathbf{in} \ M) \ N_1 \ N_2 & \rightarrow & \mathbf{newref} \ x:=n \ \mathbf{in} \ \mathbf{if} \ M \ N_1 \ N_2 & (x \notin \mathbf{fv}(N_i)) \\
 \mathbf{if} \ N_1 \ (\mathbf{newref} \ x:=n \ \mathbf{in} \ M) \ N_2 & \rightarrow & \mathbf{newref} \ x:=n \ \mathbf{in} \ \mathbf{if} \ N_1 \ M \ N_2 & (x \notin \mathbf{fv}(N_i)) \\
 \mathbf{if} \ N_1 \ N_2 \ (\mathbf{newref} \ x:=n \ \mathbf{in} \ M) & \rightarrow & \mathbf{newref} \ x:=n \ \mathbf{in} \ \mathbf{if} \ N_1 \ N_2 \ M & (x \notin \mathbf{fv}(N_i)) \\
 \mathbf{let} \ y = (\mathbf{newref} \ x:=n \ \mathbf{in} \ M) \ \mathbf{in} \ N & \rightarrow & \mathbf{newref} \ x:=n \ \mathbf{in} \ (\mathbf{let} \ y = M \ \mathbf{in} \ N) & (x \notin \mathbf{fv}(N)) \\
 \mathbf{let} \ y = M \ \mathbf{in} \ (\mathbf{newref} \ x:=n \ \mathbf{in} \ N) & \rightarrow & \mathbf{newref} \ x:=n \ \mathbf{in} \ (\mathbf{let} \ y = M \ \mathbf{in} \ N) & (x \notin \mathbf{fv}(M)) \\
 f \ (\mathbf{newref} \ x:=n \ \mathbf{in} \ M) & \rightarrow & \mathbf{newref} \ x:=n \ \mathbf{in} \ (f \ M) &
 \end{array}$$

where $f \in \{\mathbf{succ}, \mathbf{pred}, \mathbf{iszero}\}$.

Figure 41: Commutation rules for **newref**

Proof. Consider the reduction \rightarrow comprising the (context closure of) the commutations above with β -reduction and the state-free reductions for **mkvar** and **mksem**. Treating **ref** and **sem** as product types, it is easy to prove from the strong normalization of the simply-typed λ -calculus with products that \rightarrow terminates. Moreover, as M has type $\mathbb{X} \in \{\mathbb{U}, \mathbb{B}, \mathbb{N}\}$, a \rightarrow -normal form $M \rightarrow^* M'$ has no abstraction, bad variable or semaphore subterm. Thus

$$M' = \mathbf{new} \ x_1:=n_1 \ \mathbf{in} \ \dots \ \mathbf{new} \ x_p:=n_p \ \mathbf{in} \ N'$$

with $\Sigma \vdash N' : \mathbb{X}$ without recursion, subterm of higher type, bad reference or semaphores, and reference and semaphore initialization. Finally, we conclude by Lemma 4.36. \square

The semantics enjoy *finitary adequacy*:

Proposition 4.39. *For any recursion-free $\vdash M : \mathbb{U}$, we have $M \Downarrow$ iff $\llbracket M \rrbracket \Downarrow$.*

Proof. Immediate consequence of Proposition 4.37 and Lemma 4.38. \square

4.5.3. *Deducing adequacy.* Finally, we extend the above with recursion.

Theorem 4.40 (Adequacy). *For $\vdash M : \mathbb{U}$ any term of $\mathbf{IA}_{//}$, $M \Downarrow$ iff $\llbracket M \rrbracket \Downarrow$.*

Proof. As expected, we simply reason by continuity. For all type A and $n \in \mathbb{N}$, we set

$$\mathcal{Y}_A^0 = \lambda f^{A \rightarrow A}. \perp_A \quad \mathcal{Y}_A^{n+1} = \lambda f^{A \rightarrow A}. f(\mathcal{Y}_A^n f),$$

yielding $\vdash \mathcal{Y}_A^n : (A \rightarrow A) \rightarrow A$. The n -th **approximation** $\Gamma \vdash M_n : A$ of any term $\Gamma \vdash M : A$ of $\mathbf{IA}_{//}$ is obtained by replacing each $\mathcal{Y} N$ with $\mathcal{Y}_n N$. It is then routine to show that $M \Downarrow$ iff there is $n \in \mathbb{N}$ such that $M_n \Downarrow$. Likewise, by definition of the interpretation of recursion (see Section 4.4.1) and continuity of the interpretation, $\llbracket M \rrbracket = \bigvee_{n \in \mathbb{N}} \llbracket M_n \rrbracket$. Now:

$$\begin{aligned}
 M \Downarrow & \Leftrightarrow \exists n \in \mathbb{N}, M_n \Downarrow, \\
 & \Leftrightarrow \exists n \in \mathbb{N}, \llbracket M_n \rrbracket \text{ has a positive move,} \\
 & \Leftrightarrow \llbracket M \rrbracket \text{ has a positive move,}
 \end{aligned}$$

using the above along with Proposition 4.39. \square

This continuity argument only works for *may*-convergence: for extensions such as *must* or *fair*-convergence we would have to formulate a more complete correspondence between operational and game semantics – see *e.g.* [Cas17] for an adequacy result for non-deterministic PCF *w.r.t.* must-convergence. However, we leave this out of this paper.




Figure 42: $\lambda x^{\mathbb{U}} y^{\mathbb{U}}. \text{newref } r := 0 \text{ in}$
 $\quad \text{assume (iszero !}r\text{)}; x; r := \text{tt}$
 $\quad \parallel \text{assume (not (iszero !}r\text{))}; y$




Figure 43: $\lambda f^{\mathbb{U} \rightarrow \mathbb{U} \rightarrow \mathbb{U}}. \text{newref } r := 0 \text{ in}$
 $\quad f (r := 1)$
 $\quad (\text{skip } \odot (\text{assume !}r))$

4.6. **Full Abstraction.** Now that we have established \rightarrow -Strat as an adequate model of IA_{\parallel} , we explore a bit further its observable operational content and prove full abstraction.

4.6.1. *Observable behaviour of causal strategies.* It is clear that without quotient, the interpretation of IA_{\parallel} in \rightarrow -Strat will *not* be fully abstract: the model records very intensional information that is typically not observable. We give two examples in Figures 42 and 43.

We introduce a few additional pieces of syntactic sugar. First, $\text{assume} : \mathbb{B} \rightarrow \mathbb{U}$ is $\lambda x^{\mathbb{B}}. \text{if } x \text{ skip } \perp$ which terminates on tt and diverges otherwise. We also define $\text{not} : \mathbb{B} \rightarrow \mathbb{B}$ as the obvious program. Finally, for any $\Gamma \vdash M : A$ and $\Gamma \vdash N : A$ of IA_{\parallel} , we set

$$M \odot N = \text{newref } x := 0, y := 0 \text{ in}$$

$$(x := 1 \parallel y := !x);$$

$$\text{if (iszero (!}y\text{)) } M N$$

a non-deterministic sum $\Gamma \vdash M \odot N : A$, behaving non-deterministically as M or N .

First, Figure 42 represents the semantics of an encoding of sequential composition via parallel composition plus shared state. In Ghica and Murawski's model, this program has the same interpretation as sequential composition, showing their equivalence with respect to may-testing. In contrast, \rightarrow -Strat also gives account of the limitation of this encoding: the greyed out branch on the left corresponds to the bottom read winning the race, causing divergence¹⁹. Likewise, Figure 43 also has an unobservable branch greyed out. The model shows that the program may provide values for the two arguments of f independently, but it may also provide a value for the second argument of f *because* f called its first argument. In a play, an occurrence of \checkmark^+ corresponding to a value for the second argument may be causally explained by either of the two moves, but the distinction is un-observable.

For full abstraction only the observable behaviour matters, and \rightarrow -Strat clearly records more than necessary. So we ask: what parts of a concurrent strategy are observable?

¹⁹This shows our model remembers *some* divergences, though not enough to get adequacy for *must*: some divergences are lost through hiding. This can be addressed by tweaking hiding to retain those events dubbed *essential* that carry divergences, see [CCHW18, Cas17] – but we shall not take this route in this paper.

4.6.2. *Non-alternating plays with pointers.* We approach this question in terms which are no surprise to the reader familiar with Ghica and Murawski's non-alternating games: the observable behaviour is exactly captured by certain *non-alternating plays with pointers*:

Definition 4.41. A **non-alternating play with pointers** on A is a $s_1 \dots s_n$ on $|A|$ s.t.

$$\text{negative: } n \geq 1 \implies \text{pol}(s_1) = -,$$

with, for all $1 \leq j \leq n$ s.t. s_j is non-minimal in A , a **pointer** to some earlier s_i such that $s_i \rightarrow_A s_j$. We write $\mathcal{P}\text{-}\mathcal{U}\text{-Plays}(A)$ for the set of non-alternating plays with pointers on A .

Ghica and Murawski's model is an analogue of $\mathcal{U}\text{-Strat}$ based on non-alternating plays with pointers. A **GM-strategy** on meager A is a non-empty set of well-bracketed non-alternating plays with pointers (with well-bracketing defined as in Definition 3.20) satisfying conditions analogous to Definition 3.19 (see Definitions 4 and 13 in [GM08]) and additionally *thread-independent* (see Definition 17 of [GM08]). There is a cartesian closed category **GM** of meager arenas and GM-strategies, supporting the interpretation of \mathbf{IA}_{\parallel} .

We shall build a functorial bridge between $\rightarrow\text{-Strat}$ and **GM**, as in Proposition 4.30, but restricted to the cartesian closed structure as **GM** has no linear decomposition. We use the *concrete arenas* of Section 2.20, extended in the obvious way with Question/Answer labeling.

The proof of Proposition 2.22 applies unchanged to prove:

Proposition 4.42. *For any concrete arena A , there is an injective function*

$$\mathcal{P} : \mathcal{U}\text{-Plays}(A)/\cong \rightarrow \mathcal{P}\text{-}\mathcal{U}\text{-Plays}(A^0)$$

preserving length, prefix, justifiers, and reflecting and preserving well-bracketing.

We consider the cartesian closed category $\rightarrow\text{-Strat}_{\mathcal{C}}$ with objects concrete arenas, obtained from $\rightarrow\text{-Strat}_{\dagger}^{\text{wb}}$ by replacing all operations on arenas with those on concrete arenas.

Proposition 4.43. *Consider concrete $--$ -arenas A, B , and $\sigma \in \rightarrow\text{-Strat}_{\dagger}^{\text{wb}}(A, B)$. Then,*

$$\mathcal{P}\text{-}\mathcal{U}\text{-Unf}(\sigma) = \{\mathcal{P}(s) \mid s \in \mathcal{U}\text{-Unf}(\sigma^{\dagger}) \text{ well-bracketed}\}$$

is a GM-strategy on $A^0 \multimap B^0$. Moreover, $\mathcal{P}\text{-}\mathcal{U}\text{-Unf}$ extends to a cartesian closed functor

$$\mathcal{P}\text{-}\mathcal{U}\text{-Unf} : \rightarrow\text{-Strat}_{\mathcal{C}} \rightarrow \mathbf{GM}$$

preserving the interpretation of \mathbf{IA}_{\parallel} .

Proof. We detail the two critical points: preservation of symmetry, and composition.

Symmetry. Consider $\sigma, \tau \in \rightarrow\text{-Strat}_{\dagger}^{\text{wb}}(A, B)$ s.t. $\sigma \approx \tau$. Then we also have $\Lambda(\sigma^{\dagger}) \approx \Lambda(\tau^{\dagger})$, i.e. there is an iso $\varphi : \sigma^{\dagger} \cong \tau^{\dagger}$ s.t. $\partial_{\Lambda(\tau^{\dagger})} \circ \varphi \sim^+ \partial_{\Lambda(\sigma^{\dagger})}$. Consider $\mathcal{P}(s) \in \mathcal{P}\text{-}\mathcal{U}\text{-Unf}(\sigma)$ for some $s \in \mathcal{U}\text{-Unf}(\Lambda(\sigma^{\dagger}))$. This means there is $t \in \mathcal{U}\text{-Plays}(\sigma^{\dagger})$ s.t. $s = \partial_{\Lambda(\sigma^{\dagger})}(t)$. But then, $\varphi(t) \in \mathcal{U}\text{-Plays}(\tau^{\dagger})$, and from $\partial_{\Lambda(\tau^{\dagger})} \circ \varphi \sim^+ \partial_{\Lambda(\sigma^{\dagger})}$ it is direct that $s = \partial_{\Lambda(\sigma^{\dagger})}(t) \cong_{A \multimap B} \partial_{\Lambda(\tau^{\dagger})}(\varphi(t))$. By Proposition 4.42, $\mathcal{P}(\partial_{\Lambda(\sigma^{\dagger})}(t)) = \mathcal{P}(\partial_{\Lambda(\tau^{\dagger})}(\varphi(t)))$, so that $\mathcal{P}(s) \in \mathcal{P}\text{-}\mathcal{U}\text{-Unf}(\tau)$. The other inclusion is symmetric, so $\mathcal{P}\text{-}\mathcal{U}\text{-Unf}(\sigma) = \mathcal{P}\text{-}\mathcal{U}\text{-Unf}(\tau)$.

Composition. For $\sigma \in \rightarrow\text{-Strat}_{\dagger}^{\text{wb}}(A, B)$, $\tau \in \rightarrow\text{-Strat}_{\dagger}^{\text{wb}}(B, C)$, we must show

$$\mathcal{P}\text{-}\mathcal{U}\text{-Unf}(\tau \odot_{\dagger} \sigma) = \mathcal{P}\text{-}\mathcal{U}\text{-Unf}(\tau) \odot \mathcal{P}\text{-}\mathcal{U}\text{-Unf}(\sigma),$$

there are two inclusions to prove:

\subseteq . Consider $\mathcal{P}(s) \in \mathcal{P}\text{-}\mathcal{U}\text{-Unf}(\tau \odot_{\dagger} \sigma)$. Since $\mathcal{P}\text{-}\mathcal{U}\text{-Unf}(-)$ preserves positive isomorphism and by the laws of Seely categories, $\mathcal{P}(s) \in \mathcal{P}\text{-}\mathcal{U}\text{-Unf}(\tau^{\dagger} \odot \sigma^{\dagger})$ for some $s \in \mathcal{U}\text{-Unf}(\tau^{\dagger} \odot \sigma^{\dagger})$. By Proposition 4.30, the latter is $\mathcal{U}\text{-Unf}(\tau^{\dagger}) \odot \mathcal{U}\text{-Unf}(\sigma^{\dagger})$. Now,

if $u \in \mathcal{U}\text{-Unf}(\tau^\dagger) \otimes \mathcal{U}\text{-Unf}(\sigma^\dagger)$ is a witness for $s \in \mathcal{U}\text{-Unf}(\tau^\dagger) \odot \mathcal{U}\text{-Unf}(\sigma^\dagger)$, then as σ^\dagger and τ^\dagger are well-bracketed, it is direct that $u \upharpoonright A, B \in \mathcal{U}\text{-Unf}(\sigma^\dagger)$ and $u \upharpoonright B, C \in \mathcal{U}\text{-Unf}(\tau^\dagger)$ are well-bracketed as well. Therefore, $\mathcal{P}(u \upharpoonright A, B) \in \mathcal{U}\text{-Unf}(\sigma^\dagger)$ and $\mathcal{P}(u \upharpoonright B, C) \in \mathcal{U}\text{-Unf}(\tau^\dagger)$, hence $\mathcal{P}(u)$ is a witness in the sense of [GM08] for $\mathcal{P}(s) \in \mathcal{P}\text{-}\mathcal{U}\text{-Unf}(\tau) \odot \mathcal{P}\text{-}\mathcal{U}\text{-Unf}(\sigma)$.

\ni . Consider now $s \in \mathcal{P}\text{-}\mathcal{U}\text{-Unf}(\tau) \odot \mathcal{P}\text{-}\mathcal{U}\text{-Unf}(\sigma)$. There is a witness u , a sequence with pointers on $(A' \multimap B') \multimap C'$, with restrictions $u \upharpoonright A', B' \in \mathcal{P}\text{-}\mathcal{U}\text{-Unf}(\sigma)$ and $u \upharpoonright B', C' \in \mathcal{P}\text{-}\mathcal{U}\text{-Unf}(\tau)$ – we refer to [GM08] for the definitions of restrictions of plays with pointers. Write $u \upharpoonright A', B' = \mathcal{P}(t^\sigma)$ and $u \upharpoonright B', C' = \mathcal{P}(t^\tau)$ for $t^\sigma \in \mathcal{U}\text{-Unf}(\sigma^\dagger)$ and $t^\tau \in \mathcal{U}\text{-Unf}(\tau^\dagger)$. Though t^σ and t^τ yield plays with pointers compatible in B' , they might not match in B on the nose. But by Proposition 4.42, $t^\sigma \upharpoonright B \cong_{!B} t^\tau \upharpoonright B$ match *up to symmetry*. So, writing

$$t^\sigma = \partial_{\Lambda(\sigma^\dagger)}(v^\sigma) \quad t^\tau = \partial_{\Lambda(\tau^\dagger)}(v^\tau),$$

for $v^\sigma \in \mathcal{U}\text{-Plays}(\sigma^\dagger)$ and $v^\tau \in \mathcal{U}\text{-Plays}(\tau^\dagger)$, writing $x^\sigma = |v^\sigma| \in \mathcal{C}(\sigma^\dagger)$ and $x^\tau = |v^\tau| \in \mathcal{C}(\tau^\dagger)$ the symmetry $t^\sigma \upharpoonright B \cong_{!B} t^\tau \upharpoonright B$ induces $\theta : x_B^\sigma \cong_{!B} x_B^\tau$. Moreover,

$$x^\sigma \parallel x_C^\tau \stackrel{\partial_{\sigma^\dagger} \parallel !C}{\cong} x_A^\sigma \parallel x_B^\sigma \parallel x_C^\tau \stackrel{!A \parallel \theta \parallel !C}{\cong} x_A^\sigma \parallel x_B^\tau \parallel x_C^\tau \stackrel{!A \parallel \partial_\tau^{-1}}{\cong} x_A^\sigma \parallel x^\tau,$$

is secured, as u directly informs a total ordering of its graph compatible with σ^\dagger and τ^\dagger . Hence, by Proposition 4.15, there are $\varphi^\sigma : x^\sigma \cong_{\sigma^\dagger} y^\sigma$ and $\varphi^\tau : x^\tau \cong_{\tau^\dagger} y^\tau$ such that $y_B^\sigma = y_B^\tau$.

Transporting t^σ and t^τ through φ^σ and φ^τ , we get $w^\sigma \in \mathcal{U}\text{-Plays}(\sigma^\dagger)$ and $w^\tau \in \mathcal{U}\text{-Plays}(\tau^\dagger)$ s.t. we still have $\mathcal{P}(w^\sigma) = u \upharpoonright A', B'$ and $\mathcal{P}(w^\tau) = u \upharpoonright B', C'$; but this time $w^\sigma \upharpoonright B = w^\tau \upharpoonright B$. Zipping them following u we obtain $w \in \mathcal{U}\text{-Plays}(\tau^\dagger) \otimes \mathcal{U}\text{-Plays}(\sigma^\dagger)$ such that $\mathcal{P}(w \upharpoonright A, C) = s$. But then $w \upharpoonright A, C \in \mathcal{U}\text{-Plays}(\tau^\dagger) \odot \mathcal{U}\text{-Plays}(\sigma^\dagger)$ by construction, so in $\mathcal{U}\text{-Plays}((\tau \odot \sigma)^\dagger)$ by Proposition 4.30, hence $s \in \mathcal{P}\text{-}\mathcal{U}\text{-Unf}(\tau \odot \sigma)$ as required. \square

So there is a functorial unfolding from $\multimap\text{-Strat}$ to GM. To further factor out non-observable behaviour, one can restrict to *complete* plays:

Definition 4.44. Consider A a meager \multimap -arena, and $s \in \mathcal{P}\text{-}\mathcal{U}\text{-Plays}(A)$.

We say that s is **complete** iff it is well-bracketed and all its questions have an answer.

If $\sigma \in \multimap\text{-Strat}_!^{\text{wb}}(A, B)$, $\text{comp}(\sigma)$ is the subset of $\mathcal{P}\text{-}\mathcal{U}\text{-Unf}(\sigma)$ comprising complete plays only. The proposition above allows us to deduce (with \sim defined in Section 3.3.2):

Proposition 4.45. Consider A concrete and $\sigma, \tau : A$ well-bracketed causal strategies. Then,


$$\text{comp}(\sigma) = \text{comp}(\tau) \quad \implies \quad \sigma \sim \tau.$$

Proof. Consider $\alpha \in \multimap\text{-Strat}_!^{\text{wb}}(A, \mathbb{U})$ s.t. $\alpha \odot \sigma \Downarrow$. By Proposition 4.43, $\mathbf{q}^- \checkmark^+ \in \mathcal{P}\text{-}\mathcal{U}\text{-Unf}(\alpha) \odot \mathcal{P}\text{-}\mathcal{U}\text{-Unf}(\sigma)$. Considering $u \in \mathcal{P}\text{-}\mathcal{U}\text{-Unf}(\alpha) \otimes \mathcal{P}\text{-}\mathcal{U}\text{-Unf}(\sigma)$, $u \upharpoonright A^0, \mathbb{U} \in \mathcal{P}\text{-}\mathcal{U}\text{-Unf}(\sigma)$. From well-bracketing of σ and α , $u \upharpoonright A^0$ is well-bracketed and complete, so $u \upharpoonright A^0 \in \text{comp}(\tau)$, so $\mathcal{P}\text{-}\mathcal{U}\text{-Unf}(\alpha) \odot \mathcal{P}\text{-}\mathcal{U}\text{-Unf}(\tau) \Downarrow$, so $\alpha \odot \tau \Downarrow$ by Proposition 4.43. \square

To prove the converse and link it to syntactic equivalence, we examine *definability*.

4.6.3. Definability of plays with pointers. As in Sections 3.3.2 and 3.4.4, full abstraction relies on *definability*. While definability in \mathbf{IA} rests on definability for finite innocent strategies, Ghica and Murawski's definability for \mathbf{IA}_\parallel gives directly terms realizing *individual plays*. For conciseness we omit the full development, but illustrate it on a representative example.

Consider s on the left hand side of Figure 44. Naively, we want a term whose only execution is s . But strategies satisfy *courtesy*, so one realizing s must also realize all


 Figure 44: Example of definability of plays in IA_{\parallel}

plays obtainable by adding asynchronous delays. The information from s that survives asynchronous delays is that certain positive moves appear *after* earlier negative moves in s . Together with the static causality from the arena, this yields a diagram as in the right hand side of Figure 44, very much like a causal strategy. It is directly this diagram that Ghica and Murawski's finite definability reproduces. First, we ignore \rightarrow -links and start with a term

$$\lambda x^{\mathbb{U}_1} f^{\mathbb{U}_2 \rightarrow \mathbb{U}_3} y^{\mathbb{U}_4}. (x \parallel f \text{ skip} \parallel y); \perp,$$

that performs all computational events available in s in a maximally parallel fashion, with only causal dependency enforced by the game. The \rightarrow -links are restored through the memory. For that we define two helper functions. If $M : \text{ref}$, we write $\text{set}(M) : \mathbb{U}$ for $M := 1$, and $\text{test}(M) : \mathbb{U}$ for $\text{assume}(\text{not}(\text{iszero} !M))$ which converges iff $!M$ is non-zero. Then:

$$\lambda x^{\mathbb{U}_1} f^{\mathbb{U}_2 \rightarrow \mathbb{U}_3} y^{\mathbb{U}_4}. \left(\left(\begin{array}{c} x; \\ \text{set}(\sqrt{1}^-) \end{array} \right) \parallel \left(\begin{array}{c} f(\text{set}(\mathbf{q}_2^-); \\ \text{test}(\sqrt{4}^-); \\ \text{grab}(\sqrt{2}^+); \\ \text{skip}) \end{array} \right) \parallel \left(\begin{array}{c} \text{test}(\sqrt{1}^-); \\ \text{test}(\mathbf{q}_2^-); \\ y; \\ \text{set}(\sqrt{4}^-) \end{array} \right) \right); \perp$$

borrow the shape of the first term, signaling the \rightarrow -links through memory. We use one fresh reference (initialized to 0) for each Opponent move, which gets set to 1 when the Opponent move occurs. Finally, we use semaphores to ensure that Opponent replications does not cause a duplication of Player moves by prompting re-evaluation of the corresponding subterms – so that we only obtain linearizations of the diagram on the right hand side of Figure 44.

Done systematically for arbitrary plays, this establishes [GM08]:

Proposition 4.46 (Ghica, Murawski). *Consider A a type and $s \in \mathcal{P}\text{-}\mathcal{U}\text{-Plays}([A])$ complete. Then there is $\vdash M_s : A \rightarrow \mathbb{U}$ such that for all $\vdash N : A$,*

$$s \in \llbracket N \rrbracket_{\text{GM}} \quad \Leftrightarrow \quad M_s N \Downarrow$$

In particular, as $\llbracket N \rrbracket_{\text{GM}}$ is courteous, any $t \in \llbracket N \rrbracket_{\text{GM}}$ tracing a successful interaction with M_s can be converted to s through permutations whose correctness is granted by courtesy.

Theorem 4.47. *The model $\rightarrow\text{-Strat}_{\parallel}^{\text{wb}}$ is intensionally fully abstract for IA_{\parallel} .*

Proof. Consider $\vdash M, N : A$. If $\llbracket M \rrbracket \sim \llbracket N \rrbracket$, using Theorem 4.40, $M \sim N$. Reciprocally, assume $M \sim N$. Seeking a contradiction, assume $\llbracket M \rrbracket \not\sim \llbracket N \rrbracket$. By Proposition 4.45, there is *w.l.o.g.* $s \in \text{comp}(\llbracket M \rrbracket)$ where $s \notin \text{comp}(\llbracket N \rrbracket)$. So, by Proposition 4.43, $s \in \llbracket M \rrbracket_{\text{GM}}$ while $s \notin \llbracket N \rrbracket_{\text{GM}}$. Finally, by Proposition 4.46, we have $M_s M \Downarrow$ while $M_s N \Uparrow$, contradiction. \square

As for Theorem 3.14, the resulting quotient is effective and easily described: for $\sigma, \tau : A$,

$$\sigma \sim \tau \quad \Leftrightarrow \quad \text{comp}(\sigma) = \text{comp}(\tau).$$

Observational equivalence is undecidable even for the second-order fragment and without recursion [GMO06]. Note also that without semaphores, full abstraction fails [Mur10] as terms are closed under a stuttering behaviour which reduces their observational power.

4.7. Parallel Innocence and Sequentiality. We resume the discussion left at Section 3.6: can we find *parallel innocence* and *sequentiality* disentangling parallelism and interference?

A subtlety is that while $\text{IA}_{//}$ is non-deterministic, $\text{PCF}_{//}$ and IA are both deterministic – albeit in different senses. Non-determinism arises in $\text{IA}_{//}$ from the interaction of parallelism and interference so, removing either of these causes, determinism has to be reimposed as well. Accordingly, sequentiality and parallel innocence will include determinism.

5. PARALLEL INNOCENCE

We capture the causal patterns definable with pure parallel higher-order programming.

5.1. Causal determinism. The sense in which $\text{PCF}_{//}$ is deterministic is fairly different from Definition 2.8. For instance, after the first Opponent move, the strategy of Figure 38 has two available Player moves; but the order in which these moves are played does not matter and will eventually reach the same result: the program is *deterministic up to the choice of the scheduler*. If E is an event structure, write Con_E for the set of finite **consistent** sets of events, *i.e.* for $X \subseteq_f |E|$, $X \in \text{Con}_E$ iff for all $e_1, e_2 \in X$, we have $\neg(e_1 \#_{\sigma} e_2)$.

We use Winskel’s definition of *determinism* for concurrent strategies [Win12]:

Definition 5.1. A causal strategy $\sigma : A$ on arena A is **causally deterministic** if:

causal determinism: assume $X \subseteq_f |\sigma|$ is *negatively compatible*,
i.e. $X_- = \{s \in X \mid \text{pol}_{\sigma}(s) = -\} \in \text{Con}_{\sigma}$. Then, $X \in \text{Con}_{\sigma}$ as well.

This ensures that Player branching only spawns parallel threads: only Opponent may initiate conflict. Copycat is deterministic, and deterministic strategies compose [Win12]. All other constructions in the Seely category structure of $\rightarrow\text{-Strat}$ preserve determinism.

5.2. Parallel Innocence. What causal shapes are distinctive of pure parallel computation?

5.2.1. Pre-innocence. Pure parallel programs may spawn parallel threads, which must remain independent in the absence of interference. Once they both terminate the program may take new actions that depend on their results, causally “merging” them. A typical causal strategy featuring this behaviour, for $x : \mathbb{U}, y : \mathbb{U} \vdash x \parallel y : \mathbb{U}$, appears in Figure 46. The slogan is:

“Player may merge threads than he himself has spawned”.

In contrast, both diagrams of Figure 32 bear signs of interference. In the first, the answer 1^+ depends on \mathbf{q}^- : the program somehow observes if the function has called its argument, which is only possible if the argument performs some side-effect that the program observes. In the second, \checkmark^+ depends on \checkmark^- ; but likewise this can only occur if the termination of the function triggers a side-effect. In both cases, this is witnessed by Player “merging” causal chains which forked at Opponent moves. To ban such interference, the slogan is:




Figure 45: Maximal grounded causal chains of a causal strategy




Figure 46: A typical pre-innocent strategy




Figure 47: Partiality of views

“Player may not merge threads spawned by Opponent”.

To define parallel innocence, our first step is to introduce a formal notion of “thread”:

Definition 5.2. Consider A an arena, and $\sigma : A$ a causal strategy.

A **grounded causal chain (gcc)** in σ is $\rho = \{\rho_1, \dots, \rho_n\} \subseteq |\sigma|$ forming

$$\rho_1 \rightarrow_{\sigma} \dots \rightarrow_{\sigma} \rho_n$$

a chain with ρ_1 minimal with respect to \leq_{σ} . We write $\text{gcc}(\sigma)$ for the gccs in σ .

A gcc is just a set, but we write $\rho = \rho_1 \rightarrow_{\sigma} \dots \rightarrow_{\sigma} \rho_n \in \text{gcc}(\sigma)$ to insist on the causal ordering inherited from \leq_{σ} . If also $\rho_1 \rightarrow_{\sigma} \dots \rightarrow_{\sigma} \rho_n \rightarrow_{\sigma} m \in \text{gcc}(\sigma)$, then we write $\rho \rightarrow m = \rho \cup \{m\}$. Gccs are not necessarily down-closed: we show in Figure 45 all maximal gccs of a causal strategy. Of those, the second and third omit some dependencies of 1^+ .

We may now make formal the idea of a strategy “only merging threads forked by Player”.

Definition 5.3. Consider A an arena. A causally deterministic $\sigma : A$ is **pre-innocent** iff

pre-innocence: If $m^+ \in |\sigma|$ and $\rho_1 \rightarrow m, \rho_2 \rightarrow m \in \text{gcc}(\sigma)$ are distinct, then their least distinct moves are positive.

As causal strategies are *pointed*, ρ_1 and ρ_2 necessarily share the same initial move. The strategy of Figure 46 is pre-innocent. In contrast, that of Figure 31 is not – both augmentations of Figure 32 fail pre-innocence. For instance, the second and third gccs of Figure 45 arrive at 1^+ but before that, the greatest common event is q^+ , which is positive: Player is merging (via 1^+) two gccs forked by Opponent, which is forbidden by pre-innocence.

It will follow later on that the *sequential pre-innocent* causal strategies exactly match the standard alternating innocent strategies of Definition 3.6: *sequentiality* entails that there is no Player branching. Thus separate branches always correspond to threads spawned by




Figure 48: A gcc of a non-visible strategy, losing its pointer

Opponent, which by pre-innocence cannot interfere. The causal structure is then a forest, matching that of *P-views* of Section 3.2. We postpone the details to Section 6.3.

However, it turns out that pre-innocence is incomplete for parallel strategies.

5.2.2. *Visibility.* The problem arises as non-stability of pre-innocence under composition. A counter-example appears in Figure 49, examined when proving compositionality of innocence.

But we can explain the issue intuitively: the definition of pre-innocence relies on gccs which formalize a notion of *thread*. If that intuition is to be taken seriously, gccs should be valid executions of standalone sequential programs. But this is not the case: Figure 48 shows a gcc where the last move answers a question that *was not asked* within this gcc. This could not be a valid state of a sequential program, because the last move *loses its pointer*.

Visible strategies are simply those such that this does not happen.

Definition 5.4. A causal strategy $\sigma : A$ is **visible** if for all $\rho \in \text{gcc}(\sigma)$, $\partial_\sigma(\rho) \in \mathcal{C}(A)$.

In other words, every move in ρ *points* within ρ . This phrasing highlights the analogy with Definition 3.5, *i.e.* “Player always points in the P-view”. It is indeed this analogy that inspired the name²⁰. But one must be wary: the alternating interpretation of sequential programs with state yields sequential P-visible strategies, but their causal interpretation (as in Figure 31) is *not* necessarily visible. Visibility is very restrictive, it is not clear what would be a sensible programming primitive that would satisfy visibility but not pre-innocence.

We regard visibility as a key contribution. It has far-reaching consequences – some of which will be introduced in the course of this paper. In fact, visibility is used more than parallel innocence in further developments in this line of work [CCPW18, CdV20].

The following lemma captures how a gcc may be regarded as a standalone thread.

Lemma 5.5. Consider $\sigma : A$ a visible causal strategy.

Then for any $\rho = \rho_1 \rightarrow_\sigma \dots \rightarrow_\sigma \rho_n \in \text{gcc}(\sigma)$, $\partial_\sigma(\rho) = \partial_\sigma(\rho_1) \dots \partial_\sigma(\rho_n)$ is a P-view.

Proof. By Lemma A.3, $\partial_\sigma(\rho)$ is an alternating sequence. By visibility, its prefixes are configurations of A . So, $\partial_\sigma(\rho) \in \downarrow\uparrow\text{-Plays}(A)$. By Lemma A.4, the predecessor of $a^- \in \partial_\sigma(\rho)$ for \rightarrow_A is its predecessor in $\partial_\sigma(\rho)$, *i.e.* Opponent always points to the previous move. \square

We may now define *parallel innocent* causal strategies, or just *innocent* for short.

Definition 5.6. Consider $\sigma : A$ causally deterministic on arena A .

It is **parallel innocent** if it is pre-innocent and **visible**.

²⁰This, plus as in traditional game semantics, visibility is a prerequisite for a working notion of innocence.

A standard innocent strategy as in Section 3.2, under its “causal” presentation, is a forest of P-views (see Proposition 3.7), *i.e.* a forest of (displayed) gccs. In that light the definition of parallel innocent strategies seems natural: they are generated no longer by a forest of P-views, but by a directed acyclic graph of P-views with additional conflict relation. This graph describes how threads are spawned, and then may merge, following the innocence discipline ensuring that Player may not create interference between Opponent’s threads.

One of the main hurdles, in traditional game semantics, is to prove that innocent strategies compose. We now tackle this problem for parallel innocent strategies.

5.3. Composition of Visibility. First, we establish compositionality of visibility.

5.3.1. Justifiers in causal strategies. We introduce some machinery on *justifiers*. If $\sigma : A$ is a causal strategy on A some $--$ -arena, then as for plays, the immediate causality in A endows moves in $|\sigma|$ with a notion of *justifier*. This extends to $\sigma : A \vdash B$ with A and B $--$ -arenas:

Definition 5.7. Consider A and B $--$ -arenas, and $\sigma : A \vdash B$. Then, for all $m, m' \in |\sigma|$,

$$\begin{aligned} \text{just}(m) &= m' && \text{if } \partial_\sigma(m') \rightarrow_{A \vdash B} \partial_\sigma(m), \\ \text{just}(m) &= \text{init}(m) && \text{if } \partial_\sigma(m) \text{ minimal in } A, \end{aligned}$$

and undefined otherwise.

This leaves the justifier undefined exactly for moves corresponding to minimal moves in B , the *initial moves*. Note that assigning the justifier of m minimal in A to $\text{init}(m)$ ensures that the assignment of justifiers is invariant under currying. It might be helpful to the reader to observe that a causal strategy $\sigma : A \vdash B$ is visible iff for all $\rho \in \text{gcc}(\sigma)$, for all $m \in \rho$, $\text{just}(m) \in \rho$ as well: all gccs are closed under justifiers. We mention in passing this lemma:

Lemma 5.8. Consider A, B $--$ -arenas and $\sigma : A \vdash B$ a causal strategy.

Then, for any non-initial $m \in |\sigma|$, we have $\text{just}(m) <_\sigma m$. Moreover, if $\text{pol}_\sigma(m) = -$, then $\text{just}(m) \rightarrow_\sigma m$ is its (unique) immediate predecessor.

Proof. As a map of event structures, ∂_σ locally reflects causality (Lemma A.2), so $\text{just}(m) <_\sigma m$ if the first clause of Definition 5.7 applies; for the other we clearly have $\text{init}(m) <_\sigma m$.

If $\text{pol}_\sigma(m) = -$, then $\text{just}(e)$ is defined via the first clause since A is negative, and $\partial_\sigma(\text{just}(m)) \rightarrow_{A \vdash B} \partial_\sigma(m)$. Now, m has a predecessor $m' \rightarrow_\sigma m$, by courtesy $\partial_\sigma(m') \rightarrow_{A \vdash B} \partial_\sigma(m)$, so $\partial_\sigma(m') = \partial_\sigma(\text{just}(m))$ as A is forestial, and $m' = \text{just}(m)$ by local injectivity. \square

5.3.2. Justifiers in interactions. We extend justifiers to *interactions* – consider A, B and C three $--$ -arenas, and $\sigma : A \vdash B$ and $\tau : B \vdash C$ causal strategies.

Definition 5.9. We define the partial function $\text{just} : |\tau \circledast \sigma| \rightarrow |\tau \circledast \sigma|$ as $\text{just}(m) = m'$ if:

- (1) $\partial_{\tau \circledast \sigma}(m') \rightarrow_{A \parallel B \parallel C} \partial_{\tau \circledast \sigma}(m)$, or
- (2) $\partial_{\tau \circledast \sigma}(m)$ is minimal in A and $m'_\sigma = \text{init}(m_\sigma)$, or
- (3) $\partial_{\tau \circledast \sigma}(m)$ is minimal in B and $m'_\tau = \text{init}(m_\tau)$,

and undefined otherwise. We say that m' is the **justifier** of m in $\tau \circledast \sigma$.

This leaves $\text{just}(m)$ undefined exactly if it corresponds to a minimal move in C . Clearly the two notions of justifier are compatible, in the sense that for all $m \in |\tau \circledast \sigma|$, if m_σ is defined then $\text{just}(m)_\sigma$ is defined and equal to $\text{just}(m_\sigma)$, and likewise for τ .

5.3.3. *Views of gccs.* We introduce the main technical device on visible causal interactions.

We use polarities in interactions as in Section 4.2.6, and annotate events accordingly. We also write *e.g.* $a^{-,r}$ to indicate that a has polarity $-$ or r . If $\rho \in \text{gcc}(\tau \circledast \sigma)$ with last event m , we say that ρ **ends in σ** if m_σ is defined, and likewise for τ . We now define *views* of gccs, used to project a gcc of the interaction to gccs for both strategies.

Definition 5.10. Consider $\sigma : A \vdash B$ and $\tau : B \vdash C$ with A, B and C $--$ arenas.

If $\rho \in \text{gcc}(\tau \circledast \sigma)$ ends in σ , we (partially) define $\ulcorner \rho \urcorner^\sigma \in \text{gcc}(\sigma)$ by:

$$\begin{aligned} \ulcorner \rho_0 \rightarrow \dots \rightarrow \rho_n \rightarrow \rho_{n+1}^l \urcorner^\sigma &= \ulcorner \rho_0 \rightarrow \dots \rightarrow \rho_n \urcorner^\sigma \cup \{\rho_{n+1}\}, \\ \ulcorner \rho_0 \rightarrow \dots \rightarrow \rho_i \rightarrow \dots \rightarrow \rho_{n+1}^{-,r} \urcorner^\sigma &= \ulcorner \rho_0 \rightarrow \dots \rightarrow \rho_i \urcorner^\sigma \cup \{\rho_{n+1}\} && \text{if } \text{just}(\rho_{n+1}) = \rho_i \text{ in } A \text{ or } B, \\ \ulcorner \rho_0 \rightarrow \dots \rightarrow \rho_i \rightarrow \dots \rightarrow \rho_{n+1}^r \urcorner^\sigma &= \{\rho_{n+1}\} && \text{if } \rho_{n+1} \text{ minimal in } B, \end{aligned}$$

undefined otherwise. For $\rho \in \text{gcc}(\tau \circledast \sigma)$ ending in τ , we (partially) define $\ulcorner \rho \urcorner^\tau \in \text{gcc}(\tau)$:

$$\begin{aligned} \ulcorner \rho_0 \rightarrow \dots \rightarrow \rho_n \rightarrow \rho_{n+1}^r \urcorner^\tau &= \ulcorner \rho_0 \rightarrow \dots \rightarrow \rho_n \urcorner^\tau \cup \{\rho_{n+1}\}, \\ \ulcorner \rho_0 \rightarrow \dots \rightarrow \rho_i \rightarrow \dots \rightarrow \rho_{n+1}^{-,l} \urcorner^\tau &= \ulcorner \rho_0 \rightarrow \dots \rightarrow \rho_i \urcorner^\tau \cup \{\rho_{n+1}\} && \text{if } \text{just}(\rho_{n+1}) = \rho_i; \end{aligned}$$

when defined we call $\ulcorner \rho \urcorner^\sigma \in \text{gcc}(\sigma)$ the σ -**view** of ρ and $\ulcorner \rho \urcorner^\tau \in \text{gcc}(\tau)$ the τ -**view** of ρ .

These definitions almost perfectly follow Definition 3.4. The last clause is only needed for $\ulcorner - \urcorner^\sigma$ and not $\ulcorner - \urcorner^\tau$, because an initial event in C must be the first event of ρ anyway.

That this yields gccs of σ and τ rests on Lemma 4.21, and courtesy of σ and τ . The σ -view and the τ -view are in principle only partially defined, because it may be, when attempting to follow the opponent's pointer, that that justifier lies outside the gcc. For instance $\ulcorner \rho \urcorner^\tau$, for ρ in Figure 47, is not well-defined: when attempting to compute $\ulcorner \mathbf{q}^- \mathbf{q}_0^r \check{\mathbf{v}}_1^l \urcorner^\tau$, none of the clauses apply as $\text{just}(\check{\mathbf{v}}_1^l) = \mathbf{q}_1^r$ is outside ρ . The bulk of the proof of stability of visibility under composition, is to show that this cannot happen for visible strategies:

Proposition 5.11. *Let $\sigma : A \vdash B$ and $\tau : B \vdash C$ be visible causal strategies.*

Then, the views of gccs of $\tau \circledast \sigma$ as in Definition 5.10 are always well-defined.

Proof. We prove by induction on ρ that, for all prefixes of ρ ,

- (1) if ρ ends in σ , then $\ulcorner \rho \urcorner^\sigma$ is well-defined,
- (2) if ρ ends in τ , then $\ulcorner \rho \urcorner^\tau$ is well-defined.

Assume ρ finishes in τ . If the last move has polarity $-$, then either it is initial and there is nothing to prove, or by Lemma 4.21 its justifier is its predecessor in ρ , so $\ulcorner \rho \urcorner^\tau \in \text{gcc}(\tau)$ follows immediately by induction hypothesis (in that case ρ does not end in σ).

If the last move has polarity r , write $\rho = \rho' \rightarrow m_1 \rightarrow m_2^r$. By Lemma 4.21, $m_1 \rightarrow_\tau m_2$, so in particular m_1 is in τ . By induction hypothesis, $\kappa = \ulcorner \rho' \rightarrow m_1 \urcorner^\tau \in \text{gcc}(\tau)$, so

$$\kappa \rightarrow m_2 = \ulcorner \rho \urcorner^\tau \in \text{gcc}(\tau)$$

as well. But if ρ finishes in σ and τ (*i.e.* in B), we must further prove that $\ulcorner \rho \urcorner^\sigma \in \text{gcc}(\sigma)$. In that case, we observe that since $\ulcorner \rho \urcorner^\tau \in \text{gcc}(\tau)$ and τ is visible, it follows that $\text{just}(m_2) \in \ulcorner \rho \urcorner^\tau$. But that is a subset of ρ , so $\text{just}(m_2) \in \rho$. Hence the second clause of Definition 5.10 applies, and we conclude by induction hypothesis. If ρ finishes in σ , the reasoning is symmetric. \square

From this, we are now ready to conclude:

Proposition 5.12. *Let $\sigma : A \vdash B$ and $\tau : B \vdash C$ be visible causal strategies.*

Then, $\tau \odot \sigma : A \vdash C$ is also visible.




Figure 49: Failure of preservation of pre-innocence under composition

Proof. We prove by induction on ρ that for all $\rho \in \text{gcc}(\tau \circledast \sigma)$, $\partial_{\tau \circledast \sigma}(\rho) \in \mathcal{C}(A \parallel B \parallel C)$. If ρ is empty it is clear; take $\rho \rightarrow m \in \text{gcc}(\tau \circledast \sigma)$. By induction hypothesis, $\partial_{\tau \circledast \sigma}(\rho) \in \mathcal{C}(A \parallel B \parallel C)$, we only need that the justifier of m is in ρ . We reason by cases on the polarity of m : if it is σ , then by Proposition 5.11 $\ulcorner \rho \rightarrow m^1 \sigma \urcorner \in \text{gcc}(\sigma)$. But since σ is visible, the justifier of m appears in $\ulcorner \rho^1 \sigma \urcorner$; a subset of ρ . The other cases are symmetric or trivial.

Now, take $\rho_{\odot} \in \text{gcc}(\tau \odot \sigma)$. By definition of $\tau \odot \sigma$, there is a (non-necessarily unique) $\rho_{\circledast} \in \text{gcc}(\tau \circledast \sigma)$ such that ρ_{\odot} comprises exactly those events of ρ_{\circledast} occurring in A or C . By the observation above, $\partial_{\tau \circledast \sigma}(\rho_{\circledast}) \in \mathcal{C}(A \parallel B \parallel C)$, hence $\partial_{\tau \odot \sigma}(\rho_{\odot}) \in \mathcal{C}(A \parallel C)$. \square

5.4. Composition of Innocence. We now address composition of pre-innocence.

We start this section by showing “what could go wrong”. In Figure 49, we show a counter-example to the stability under composition of pre-innocence without visibility, with the corresponding interaction appearing as Figure 50. Let us attempt to explain the phenomenon, calling σ the left hand side strategy (parallel composition) and τ the right hand side one – observe that the dotted lines include the justifications relations from Definition 5.7 rather than just those coming from the arena. Imagine that τ wants to perform an illegal causal merge between the two argument calls of its argument of type $\mathbb{U} \multimap \mathbb{U} \multimap \mathbb{U}$. By pre-innocence it cannot do so directly. However, it can outsource the merge to σ by linking (legally with respect to pre-innocence, but illegally with respect to visibility) the arguments of the parallel composition to those that it wants to merge.

We shall prove that this cannot happen in the presence of visibility. Let us fix, until the end of the section, two visible causal strategies $\sigma : A \vdash B$ and $\tau : B \vdash C$.

5.4.1. *The “forking lemma”.* Taking a closer look at Figure 50, we highlight the two illegally merging gccs in the interaction: while σ is responsible for the merge, the point where these gccs forked is external, outside the scope of σ ! The next lemma, dubbed the “forking lemma”, forbids this: it implies that visible strategies cannot unknowingly close an Opponent fork.

If $\rho = \rho_1 \rightarrow \dots \rightarrow \rho_n$ is a gcc and $1 \leq i \leq n$, $\rho_{\leq i}$ is the gcc $\rho_1 \rightarrow \dots \rightarrow \rho_i$. Two gccs ρ, κ are **forking** iff $\rho \cap \kappa \neq \emptyset$, and for all i, j , if $\rho_i = \kappa_j$ then $\rho_{\leq i} = \kappa_{\leq j}$. If ρ, κ are two forking gccs, we write $\text{gce}(\rho, \kappa)$ for their **greatest common event**. Notice that despite the terminology, two forking gccs can be prefix of one another and never truly go separate ways.

Lemma 5.13 (Forking lemma). *Let $\rho, \kappa \in \text{gcc}(\tau \circledast \sigma)$ be forking gccs ending in σ , s.t. $\ulcorner \rho^1 \sigma \urcorner \cap \ulcorner \kappa^1 \sigma \urcorner \neq \emptyset$ and $\text{gce}(\ulcorner \rho^1 \sigma \urcorner, \ulcorner \kappa^1 \sigma \urcorner)$ negative (the least distinct events, if any, are positive).*

$$\mathbb{U} \multimap \mathbb{U} \multimap \mathbb{U} \vdash (\mathbb{U} \multimap \mathbb{U} \multimap \mathbb{U}) \multimap \mathbb{U}$$





Figure 50: Illegal causal merge

Figure 51: Merging paths in G

Then, $\text{gce}(\ulcorner \rho^1 \sigma, \ulcorner \kappa^1 \sigma) = \text{gce}(\rho, \kappa)$. Moreover, the symmetric property holds for τ .

Proof. We only detail the proof for σ , the proof for τ is exactly the same. We build a directed graph G with vertices $\rho \cup \kappa$, and edges the (disjoint) union of the sets:

$$\begin{aligned} \text{O-edges} &= \{(m_1^1, m_2^1) \mid \text{just}(m_1) = m_2\} \\ \text{P-edges} &= \{(m_1^1, m_2^1) \mid m_2 \rightarrow_{\tau \otimes \sigma} m_1\} \end{aligned}$$

where the annotation m_i^1 indicates the polarity. Each vertex is source of at most one edge, and following edges consists exactly in computing the σ -view. If ρ and κ have the same final move, then $\rho = \kappa$. Otherwise, consider the two paths in G starting with these distinct final moves. Since $\ulcorner \rho^1 \sigma \cap \ulcorner \kappa^1 \sigma \neq \emptyset$, these two paths must intersect – Figure 51 represents a typical G with O -edges in blue and P -edges in red with the two typical cases.

These paths meet at a vertex of incoming degree at least 2; but vertices receive only O -edges, or only P -edges. For the former (as in the bottom of Figure 51), then $\text{gce}(\ulcorner \rho^1 \sigma, \ulcorner \kappa^1 \sigma)$ is positive, which contradicts the hypothesis. For the latter (as in the top of Figure 51), we remark that P -edges are immediate causal links in $\tau \otimes \sigma$; and there is at most one event in $\rho \cup \kappa$ causing two distinct events: if it exists, it must be $\text{gce}(\rho, \kappa)$. \square

This provides the core argument for the compositionality of pre-innocence: intuitively, if a pre-innocent strategy merges two threads, by pre-innocence its *views* of these two threads fork positively. But then the forking lemma ensures that this strategy sees the actual forking point for these threads – which therefore cannot be due to the external Opponent.


5.4.2. *Stability of \rightarrow -pre-innocence.* Now, much of the proof consists in restricting the causal shapes in $\tau \otimes \sigma$ corresponding to a causal merge in $\tau \odot \sigma$, so that the forking lemma applies.

Proposition 5.14. *Consider $\sigma : A \vdash B$ and $\tau : C \vdash C$ visible causal strategies.*

If $\sigma : A \vdash B$ and $\tau : B \vdash C$ are pre-innocent, then so is $\tau \odot \sigma$.

Proof. Consider $m \in |\tau \odot \sigma|$ and distinct $\rho^1 \rightarrow m, \rho^2 \rightarrow m \in \text{gcc}(\tau \odot \sigma)$. *W.l.o.g.* assume that whenever $\rho_i^1 = \rho_j^2, \rho_{\leq i}^1 = \rho_{\leq j}^2$ – or we can change m and ρ^i keeping the same least distinct events, but satisfying this property. Likewise, since ρ^1, ρ^2 are distinct, we assume *w.l.o.g.* that their last moves $m_1 \in \rho^1, m_2 \in \rho^2$ are distinct – or we may replace m with an earlier causal


merge. These two causal chains ρ^1 and ρ^2 may be completed to $\kappa^1 \rightarrow m, \kappa^2 \rightarrow m \in \text{gcc}(\tau \otimes \sigma)$ such that ρ^i consists exactly of the events of κ^i occurring in A or C . Necessarily, the greatest visible events of κ^1 and κ^2 are m_1 and m_2 respectively. Call n the least common event of $\kappa^1 \rightarrow m$ and $\kappa^2 \rightarrow m$ above m_1 and m_2 (which might not be m). The situation is:



with m_1, m_2 and m visible, and no one visible in between. We reason on the polarity of n in $\tau \otimes \sigma$. By Lemma 4.21 and since arenas are forestial, it cannot be negative, so its polarity is either l or r . Assume it is l – the other case is symmetric. We may compute the σ -views:

$$\ulcorner \kappa_{\leq n}^1 \urcorner \sigma \rightarrow n \urcorner \sigma, \quad \ulcorner \kappa_{\leq p}^2 \urcorner \sigma \rightarrow n \urcorner \sigma \in \text{gcc}(\sigma),$$

respectively $\ulcorner \kappa_{\leq n}^1 \urcorner \sigma \rightarrow n_\sigma$ and $\ulcorner \kappa_{\leq p}^2 \urcorner \sigma \rightarrow n_\sigma$, with $\ulcorner \kappa_{\leq n}^1 \urcorner \sigma$ and $\ulcorner \kappa_{\leq p}^2 \urcorner \sigma$ distinct as they respectively contain κ_n^1 and κ_p^2 . Since σ is pointed, $\ulcorner \kappa_{\leq n}^1 \urcorner \sigma \cap \ulcorner \kappa_{\leq p}^2 \urcorner \sigma \neq \emptyset$, so $\kappa^1 \cap \kappa^2 \neq \emptyset$ as well. Call m' the greatest common event of κ^1 and κ^2 , necessarily below m_1 and m_2 , then:



assuming *w.l.o.g.* that $\kappa_{\leq i}^1 = \kappa_{\leq i}^2$ (changing the beginning of κ^2 if required). Summing up some properties, $\xi^1 = \kappa_{i+1}^1 \dots \kappa_n^1$ and $\xi^2 = \kappa_{i+1}^2 \dots \kappa_p^2$ are disjoint. This entails the σ -views

$$\ulcorner \kappa_{\leq n}^1 \urcorner \sigma, \quad \ulcorner \kappa_{\leq p}^2 \urcorner \sigma \in \text{gcc}(\sigma)$$

are *forking*: they coincide on a prefix and disjoint afterwards. Indeed, since ξ^1 and ξ^2 are disjoint, any common event appears in $\kappa_1^1 \rightarrow \dots \rightarrow \kappa_i^1$, before which the σ -view coincides.

Now, since σ is pre-innocent, the least distinct moves m'_1 and m'_2 of $\ulcorner \kappa_{\leq n}^1 \urcorner \sigma$ and $\ulcorner \kappa_{\leq p}^2 \urcorner \sigma$ are positive. Thus their common immediate predecessor is negative – but it is also their greatest common event, since $\ulcorner \kappa_{\leq n}^1 \urcorner \sigma$ and $\ulcorner \kappa_{\leq p}^2 \urcorner \sigma$ are forking. So, by Lemma 5.13,

$$\text{gce}(\ulcorner \kappa_{\leq n}^1 \urcorner \sigma, \ulcorner \kappa_{\leq p}^2 \urcorner \sigma) = \text{gce}(\kappa_{\leq n}^1, \kappa_{\leq p}^2) = m',$$

so m' is negative for σ . In $\tau \otimes \sigma$, m' is negative or in B – in both cases the least visible events in ξ^1 and ξ^2 are positive, but those are our least distinct events of ρ_1 and ρ_2 . \square

Proposition 5.15. *There is $\rightarrow\text{-Strat}^{\text{wb,inn}}$, a sub-Seely category of $\rightarrow\text{-Strat}^{\text{wb}}$ having the same objects and morphisms restricted to parallel innocent causal strategies.*

Proof. Propositions 5.12 and 5.14 ensure that parallel innocent strategies compose. Stability under tensor and pairing are immediate. It remains that structural morphisms are innocent.

We detail it for copycat. Consider A a $--$ -arena. We show that any $(i, a) \in |\mathfrak{C}_A|$, (i, a) is minimal or has exactly one predecessor for $\rightarrow_{\mathfrak{C}_A}$. Assume first $\text{pol}(i, a) = -$. If it is not minimal, take $(j, a') \rightarrow_{\mathfrak{C}_A} (i, a)$. Necessarily $i = j$ and $a' \rightarrow_A a$, so uniqueness follows from A forestial. If $\text{pol}(i, a) = +$, its unique immediate predecessor is $(2 - i, a)$. So, \mathfrak{C}_A is forestial.

Now, for visibility, consider $\rho \in \text{gcc}(\mathfrak{C}_A)$. But since \mathfrak{C}_A is forestial, $\rho \in \mathcal{C}(\mathfrak{C}_A)$, so clearly $\hat{\partial}_{\mathfrak{C}_A}(\rho) = \rho \in \mathcal{C}(A \vdash A)$. For parallel innocence, consider $\rho_1 \rightarrow_{\mathfrak{C}_A}(i, a), \rho_2 \rightarrow_{\mathfrak{C}_A}(i, a) \in \text{gcc}(\mathfrak{C}_A)$ distinct. But since \mathfrak{C}_A is forestial, $\rho_1 = \rho_2$, contradicting their distinctness. \square

5.4.3. *Interpretation of PCF_{\parallel} .* We complete the interpretation. For all sequential primitives of PCF_{\parallel} , the corresponding strategy is forestial: as for copycat, parallel innocence follows. Finally, the strategy $\text{plet}_{\mathbb{X}, \mathbb{Y}}$ is shown parallel innocent by direct inspection (see Figure 38). Altogether, this yields an interpretation of PCF_{\parallel} in $\rightarrow\text{-Strat}_!^{\text{wb,inn}}$, and we have:

Corollary 5.16 (Adequacy). *For $\vdash M : \mathbb{U}$ any term of PCF_{\parallel} , $M \Downarrow$ iff $\llbracket M \rrbracket_{\rightarrow\text{-Strat}_!^{\text{wb,inn}}} \Downarrow$.*

Proof. Consequence of Theorem 4.40, as PCF_{\parallel} is a sub-language of IA_{\parallel} . \square

By the end of the paper, we will have established that $\rightarrow\text{-Strat}^{\text{wb,inn}}$ is intensionally fully abstract for PCF_{\parallel} . However the corresponding technical development is left for last.

6. SEQUENTIALITY AND CAUSAL FULL ABSTRACTION FOR IA AND PCF

In this section, we shall define *sequentiality* on $\rightarrow\text{-Strat}^{\text{wb}}$, then prove full abstraction results

$$\begin{aligned} \rightarrow\text{-Strat}_!^{\text{wb}} + \text{sequentiality} & \text{ is fully abstract for } \text{IA}, \\ \rightarrow\text{-Strat}_!^{\text{wb}} + \text{parallel innocence} + \text{sequentiality} & \text{ is fully abstract for } \text{PCF}, \end{aligned}$$

established by linking with Theorems 3.10 and 3.14 for alternating strategies.

6.1. **Sequentiality.** We construct the Seely category of *sequential* causal strategies.

6.1.1. *Definition.* Intuitively, a causal strategy $\sigma : A$ is *sequential* if it unfolds gracefully to a (deterministic) alternating strategy. That does not mean that Player never throws parallel threads, or always acts deterministically: for instance, the strategy in Figure 31 should be sequential and yet has certain of its configurations enabling two parallel or two conflicting Player moves. But: as long as Opponent follows an alternating discipline, so should Player.

A play $s \in \mathcal{O}\text{-Plays}(A)$ is **alternating** if $s \in \downarrow\uparrow\text{-Plays}(A)$. We shall often use $\mathcal{O}\text{-Plays}(-)$ or $\downarrow\uparrow\text{-Plays}(-)$ on causal strategies – recall that $\downarrow\uparrow\text{-Plays}(\sigma)$ and $\mathcal{O}\text{-Plays}(\sigma)$ are sequences of $|\sigma|$, and must not be confused with $\mathcal{O}\text{-Unf}(\sigma)$ that includes a move-by-move projection via the display map. We start by giving the definition of (deterministic) *sequentiality*:

Definition 6.1. Consider A an arena. A causal (pre)strategy $\sigma : A$ is **sequential** if:

- reachable sequentiality:* for all $tn^+ \in \mathcal{O}\text{-Plays}(\sigma)$, if $t \in \downarrow\uparrow\text{-Plays}(\sigma)$ then $tn \in \downarrow\uparrow\text{-Plays}(\sigma)$.
- sequential determinism:* for all $tn_1^+, tn_2^+ \in \downarrow\uparrow\text{-Plays}(\sigma)$, then $n_1 = n_2$;
- sequential visibility:* every alternating $s \in \mathcal{O}\text{-Unf}(\sigma)$ is P-visible.

For *sequential determinism*, more than merely asking that $\mathcal{O}\text{-Unf}(\sigma)$ acts deterministically on alternating plays, this condition imposes that no *internal* non-deterministic choice is alternatingly reachable, even when this choice would yield no observable non-deterministic behaviour (this is required for the forthcoming alternating projection to preserve symmetry).

Sequential visibility is perhaps puzzling, as P-visibility is usually associated not to sequentiality, but to the absence of higher-order state [AHM98]. From a given control point, a P-visible strategy may only call a procedure bound within the branch of the syntax

tree leading to that control point. In contrast, with higher-order state, a program may call a procedure stored in the memory, originating from a remote program phrase outside the current branch. This phenomenon is independent of sequentiality, but in \rightarrow -Strat the causality due to the syntax tree blurs with that due to interference. So strategies arising from the interpretation of \mathbf{IA}_{\parallel} are “morally” P-visible but formalizing this is nontrivial²¹. For us it is not worth the trouble as P-visibility is not required for full abstraction for \mathbf{IA}_{\parallel} ²². Consequently, it suffices to reinstate it once we restrict to sequential strategies.

6.1.2. *Alternating projection.* We start by defining the *alternating projection*.

Definition 6.2. Consider A, B --arenas, and $\sigma : A \vdash B$ a sequential causal (pre)strategy. Then, we define $\downarrow\uparrow\text{-Unf}(\sigma) = \{\partial_{\Lambda(\sigma)}(t) \mid t \in \downarrow\uparrow\text{-Plays}(\sigma)\}$.

We shall prove that $\downarrow\uparrow\text{-Unf}(\sigma)$ is an alternating (pre)strategy on $A \multimap B$. The two subtle points are that it is deterministic, and uniform – which both rest on the observation:

Lemma 6.3. Consider A, B --arenas, and $\sigma : A \vdash B$ a sequential (pre)strategy.

For any $s \in \downarrow\uparrow\text{-Unf}(\sigma)$, there is a unique $t \in \downarrow\uparrow\text{-Plays}(\sigma)$ such that $s = \partial_{\Lambda(\sigma)}(t)$.

Proof. Immediate by induction, using receptivity and *sequential determinism*. \square

From this it follows that $\downarrow\uparrow\text{-Unf}(\sigma)$ satisfies determinism. For uniformity, we prove:

Lemma 6.4. Consider A an arena and $\sigma, \tau : A \vdash B$ sequential causal (pre)strategies.

If $\sigma \approx \tau$, then $\downarrow\uparrow\text{-Unf}(\sigma) \approx \downarrow\uparrow\text{-Unf}(\tau)$.

Proof. If $\sigma \approx \tau$, by Definition 4.13 there is an isomorphism $\varphi : \sigma \cong \tau$ of ess satisfying

$$\partial_{\tau} \circ \varphi \sim^+ \partial_{\sigma}.$$

We prove by induction on $s^{\sigma} \in \downarrow\uparrow\text{-Unf}(\sigma)$ that if $s^{\sigma} \cong_{A+B} s^{\tau}$ with $s^{\tau} \in \downarrow\uparrow\text{-Unf}(\tau)$, then taking $t^{\sigma} \in \downarrow\uparrow\text{-Plays}(\sigma)$ and $t^{\tau} \in \downarrow\uparrow\text{-Plays}(\tau)$ s.t. $s^{\sigma} = \partial_{\Lambda(\sigma)}(t^{\sigma})$ and $s^{\tau} = \partial_{\Lambda(\tau)}(t^{\tau})$ from Lemma 6.3, we have $\varphi(t^{\sigma}) \cong_{\tau} t^{\tau}$ as in (the obvious generalization of) Definition 2.11.

For s^{σ} empty it is clear. For $s^{\sigma} m_1^- \in \downarrow\uparrow\text{-Unf}(\sigma)$ and $s^{\tau} m_2^- \in \downarrow\uparrow\text{-Unf}(\tau)$ with $s^{\sigma} m_1^- \cong_{A+B} s^{\tau} m_2^-$, there is a unique matching $t^{\sigma} n_1^- \in \downarrow\uparrow\text{-Plays}(\sigma)$. So $\varphi(t^{\sigma}) \varphi(n_1) \in \downarrow\uparrow\text{-Plays}(\tau)$, and by induction hypothesis $\varphi(t^{\sigma}) \cong_{\tau} t^{\tau}$. Moreover, from $s^{\sigma} m_1^- \cong_A s^{\tau} m_2^-$ and $\partial_{\Lambda(\tau)} \circ \varphi \sim \partial_{\Lambda(\sigma)}$,

$$\partial_{\Lambda(\tau)}(\varphi(t^{\sigma}) \varphi(n_1)) \cong_{A+B} \partial_{\Lambda(\tau)}(t^{\tau} n_2),$$

so $\varphi(t^{\sigma}) \varphi(n_1) \cong_{\tau} t^{\tau} n_2$ by \sim -receptivity of τ .

Now, using this auxiliary statement we prove the lemma. For $s^{\sigma} m_1^+ \in \downarrow\uparrow\text{-Unf}(\sigma)$ and $s^{\tau} m_2^+ \in \downarrow\uparrow\text{-Unf}(\tau)$ with $s^{\sigma} m_1^+ \cong_{A+B} s^{\tau} m_2^+$, take $t^{\sigma} n_1 \in \downarrow\uparrow\text{-Plays}(\sigma)$ and $t^{\tau} n_2 \in \downarrow\uparrow\text{-Plays}(\tau)$ s.t. $\partial_{\Lambda(\sigma)}(t^{\sigma} n_1) = s^{\sigma} m_1$ and $\partial_{\Lambda(\tau)}(t^{\tau} n_2) = s^{\tau} m_2$. By induction hypothesis, $\varphi(t^{\sigma}) \cong_{\tau} t^{\tau}$. As the former extends with $\varphi(n_1)$, by *extension* for $\mathcal{S}(\tau)$ there is some (positive) n_2' s.t. $\varphi(t^{\sigma}) \varphi(n_1) \cong_{\tau} t^{\tau} n_2'$ but now, by *reachable determinism* of τ , $n_2 = n_2'$. \square

In particular, if $\sigma : A \vdash B$ is sequential, then $\sigma \approx \sigma$ via the identity isomorphism, consequently $\downarrow\uparrow\text{-Unf}(\sigma) \approx \downarrow\uparrow\text{-Unf}(\sigma)$, i.e. $\downarrow\uparrow\text{-Unf}(\sigma)$ is uniform – from here we conclude:

Proposition 6.5. Consider A, B --arenas, and $\sigma : A \vdash B$ a sequential causal strategy.

Then, $\downarrow\uparrow\text{-Unf}(\sigma) : A \multimap B$ is a P-visible alternating strategy.

²¹This was done by Laird in the first interleaving games model [Lai01b], via explicit threading information.

²²Proposition 4.46 shows that non-visible behaviour characteristic of higher-order state can be mimicked by running several threads in parallel and using signaling via *interference* to jump control between them.




Figure 52: State diagram for sequential interactions

Proof. *Prefix-closure* and *receptivity* are clear. *Determinism* is immediate from Lemma 6.3 and *sequential determinism*, and *uniformity* is immediate from Lemma 6.4.

Finally, *P-visibility* is immediate by definition of *sequential visibility*. \square

6.1.3. *Composition.* We now focus on composition of sequential causal strategies.

We introduce some terminology. If A, B and C are $-$ -arenas and $\sigma : A \vdash B$, $\tau : B \vdash C$ are causal (pre)strategies, then we define $\mathcal{U}\text{-Plays}(\tau \circledast \sigma)$ as in Definition 3.18, referring to polarities of events in $\{-, l, r\}$. If $u \in \mathcal{U}\text{-Plays}(\tau \circledast \sigma)$, we write $u_\sigma \in \mathcal{U}\text{-Plays}(\sigma)$, $u_\tau \in \mathcal{U}\text{-Plays}(\tau)$, and $u_\odot \in \mathcal{U}\text{-Plays}(\tau \odot \sigma)$ for the obvious restrictions. If a play is alternating, it is **in state** O if it has even-length, and **in state** P if it has odd length.

Many properties of causal sequentiality follow from the following crucial observation:

Lemma 6.6. *Consider $\sigma : A \vdash B$ and $\tau : B \vdash C$ sequential causal strategies.*

Then, for any $u \in \mathcal{U}\text{-Plays}(\tau \circledast \sigma)$ such that $u_\odot \in \downarrow\uparrow\text{-Plays}(\tau \odot \sigma)$, we have $u_\sigma \in \downarrow\uparrow\text{-Plays}(\sigma)$ and $u_\tau \in \downarrow\uparrow\text{-Plays}(\tau)$, and we are in one of the following three cases:

- (1) $u_\sigma, u_\tau, u_\odot$ are respectively in state O, O, O ,
- (2) $u_\sigma, u_\tau, u_\odot$ are respectively in state O, P, P ,
- (3) $u_\sigma, u_\tau, u_\odot$ are respectively in state P, O, P .

Proof. By induction on u . If u is empty, this is clear. Consider $um \in \mathcal{U}\text{-Plays}(\tau \circledast \sigma)$. By induction hypothesis u is in one of cases (1), (2) and (3). We distinguish cases:

(1) Seeking a contradiction, assume m occurs in B . Then, one of m_σ or m_τ is positive – say *w.l.o.g.* the former. By induction hypothesis, u_σ is alternating in state O , so ends with a Player move. But so, $u_\sigma m_\sigma^+ \in \mathcal{U}\text{-Plays}(\sigma)$ with $u_\sigma \in \downarrow\uparrow\text{-Plays}(\sigma)$, so $u_\sigma m_\sigma \in \downarrow\uparrow\text{-Plays}(\sigma)$ since σ satisfies reachable sequentiality, contradiction. So, m occurs in A or C – assume *w.l.o.g.* in A . Since u_\odot is in state O , m is negative – then, it is direct that um satisfies (3).

(2) First assume that m occurs in A . Since $u_\odot \in \downarrow\uparrow\text{-Plays}(\tau \circledast \sigma)$ is in state P , then m is positive; then m_σ is positive, contradicting reachable sequentiality of σ with the fact that u_σ is in state O . Similarly, if m occurs in B it has polarity r and we transition to (3), and if m occurs in C it has polarity r and we transition to (1). (3) Symmetric to (2). \square

In other words, as long as the external Opponent respects the alternation discipline, interactions follow the familiar *state diagram* of interactions in alternating game semantics, shown in Figure 52. None of the interacting agents can be the first to break alternation, so the interaction ends up fully alternating. It follows that $\tau \odot \sigma$ satisfies reachable sequentiality:

Lemma 6.7. *Consider $\sigma : A \vdash B$, $\tau : B \vdash C$ sequential causal (pre)strategies.*

Then, $\tau \odot \sigma$ satisfies reachable sequentiality.

Proof. Consider $tn^+ \in \mathcal{U}\text{-Plays}(\tau \odot \sigma)$ s.t. $t \in \downarrow\uparrow\text{-Plays}(\tau \odot \sigma)$. We have either n_σ or n_τ defined and positive, say the former *w.l.o.g.*. Let us complete tn to $un \in \mathcal{U}\text{-Plays}(\tau \otimes \sigma)$. We have $u_\odot = t \in \downarrow\uparrow\text{-Plays}(\tau \odot \sigma)$, so by may distinduish along the three cases of Lemma 6.6:

- (1) If $u_\sigma, u_\tau, u_\odot$ are in state O, O, O . By hypothesis, $u_\sigma n_\sigma^+ \in \mathcal{U}\text{-Plays}(\sigma)$. Since σ satisfies *reachable sequentiality*, $u_\sigma n_\sigma \in \downarrow\uparrow\text{-Plays}(\sigma)$ as well, contradicting u_σ in state O .
- (2) If $u_\sigma, u_\tau, u_\odot$ are in state O, P, P , as in (1) this contradicts *reachable sequentiality*.
- (3) If $u_\sigma, u_\tau, u_\odot$ are in state P, O, P . With $u_\odot = t$ in state P , $tn \in \downarrow\uparrow\text{-Plays}(\tau \odot \sigma)$. \square

It also follows that $\tau \odot \sigma$ satisfies *sequential determinism*: Lemma 6.6 expresses that in an alternatingly reachable interaction, only one agent has control at any point. So any alternatingly reachable non-deterministic choice in $\tau \odot \sigma$ can be attributed to σ and τ :

Lemma 6.8. *Consider $\sigma : A \vdash B$ and $\tau : B \vdash C$ sequential causal (pre)strategies.*

Then, $\tau \odot \sigma$ satisfies sequential determinism.

Proof. Consider $tn_1^+, tn_2^+ \in \downarrow\uparrow\text{-Plays}(\tau \odot \sigma)$ completed to $u_1n_1, u_2n_2 \in \mathcal{U}\text{-Plays}(\tau \otimes \sigma)$, and u' the greatest common prefix of u_1n_1 and u_2n_2 , with $u'm_1 \sqsubseteq u_1n_1$ and $u'm_2 \sqsubseteq u_2n_2$. Necessarily, the visible restriction of u' is a prefix $t' \sqsubseteq t$, so in particular $t' \in \downarrow\uparrow\text{-Plays}(\tau \odot \sigma)$.

We distinguish cases on Lemma 6.6 applied to u' . For (1), m_1 and m_2 must both be the next negative event appearing in t , so $m_1 = m_2$, contradiction. For (2), m_1 and m_2 both have polarity r and we have $u'_\tau \in \downarrow\uparrow\text{-Plays}(\tau)$, $u'_\tau(m_1)_\tau^+, u'_\tau(m_2)_\tau^+ \in \downarrow\uparrow\text{-Plays}(\tau)$. Hence $(m_1)_\tau = (m_2)_\tau$ since τ satisfies sequential determinism – so $m_1 = m_2$ by local injectivity of the projection Π_τ ; contradiction. Finally, for (3), then the reasoning is symmetric. \square

To conclude compositionality, we link with composition of the alternating projections:

Lemma 6.9. *Consider $\sigma : A \vdash B$ and $\tau : B \vdash C$ sequential causal (pre)strategies.*

Then, $\downarrow\uparrow\text{-Unf}(\tau \odot \sigma) = \downarrow\uparrow\text{-Unf}(\tau) \odot \downarrow\uparrow\text{-Unf}(\sigma)$.

Proof. \supseteq . If $s \in \downarrow\uparrow\text{-Unf}(\tau) \odot \downarrow\uparrow\text{-Unf}(\sigma)$, then it is in $\mathcal{U}\text{-Unf}(\tau) \odot \mathcal{U}\text{-Unf}(\sigma)$, thus in $\mathcal{U}\text{-Unf}(\tau \odot \sigma)$ by Proposition 4.30. As s is alternating, we also have $s \in \downarrow\uparrow\text{-Unf}(\tau \odot \sigma)$.

\subseteq . If $s \in \downarrow\uparrow\text{-Unf}(\tau \odot \sigma)$, there is $t \in \downarrow\uparrow\text{-Plays}(\tau \odot \sigma)$ s.t. $s = \partial_{\Lambda(\tau \odot \sigma)}(t)$ completed to $v \in \mathcal{U}\text{-Plays}(\tau \otimes \sigma)$. We display v to $u \in \mathcal{U}\text{-Unf}(\tau) \otimes \mathcal{U}\text{-Unf}(\sigma)$ s.t. $u \upharpoonright A, C = s$. But then, by Lemma 6.6, $v_\sigma \in \downarrow\uparrow\text{-Plays}(\sigma)$ and $v_\tau \in \downarrow\uparrow\text{-Plays}(\tau)$, so $u \upharpoonright A, B$ and $u \upharpoonright B, C$ are actually alternating, and $u \in \downarrow\uparrow\text{-Unf}(\tau) \otimes \downarrow\uparrow\text{-Unf}(\sigma)$. Thus, $s \in \downarrow\uparrow\text{-Unf}(\tau) \odot \downarrow\uparrow\text{-Unf}(\sigma)$. \square

Proposition 6.10. *There is a category $\rightarrow\text{-Strat}^{\text{seq}}$ of $--$ arenas and sequential strategies.*

Moreover, there is a functor $\downarrow\uparrow\text{-Unf}(-) : \rightarrow\text{-Strat}^{\text{seq}} \rightarrow \downarrow\uparrow\text{-Strat}^{\text{vis}}$, preserving \approx .

Proof. Category. It is straightforward that copycat is sequential. Consider $\sigma : A \vdash B$ and $\tau : B \vdash C$ sequential causal strategies. By Lemma 6.7, $\tau \odot \sigma$ satisfies *reachable sequentiality*. By Lemma 6.8, it satisfies *sequential determinism*. By Lemma 6.9 and preservation of P -visible alternating strategies under composition, it satisfies *sequential visibility*.

Functorial projection. Preservation of copycat is a direct verification. Preservation of composition is Lemma 6.9. Preservation of symmetry is Lemma 6.4. \square

6.1.4. *Seely category.* We now show that sequential causal strategies form a Seely category.

We use the following notations, for $\sigma : A \vdash B$ and $\tau : C \vdash D$: if $s \in \mathcal{U}\text{-Plays}(\sigma \otimes \tau)$, then $s_\sigma \in \mathcal{U}\text{-Plays}(\sigma)$ and $s_\tau \in \mathcal{U}\text{-Plays}(\tau)$ are the corresponding restrictions, defined in the obvious way. Stability of sequentiality under tensor uses a state analysis as in Lemma 6.6:

Lemma 6.11. *Consider $\sigma : A \vdash B$ and $\tau : C \vdash D$ sequential (pre)strategies. For any $s \in \downarrow\uparrow\text{-Plays}(\sigma \otimes \tau)$ then $s_\sigma \in \downarrow\uparrow\text{-Plays}(\sigma)$ and $s_\tau \in \downarrow\uparrow\text{-Plays}(\tau)$. Moreover, we are in one of:*

- (1) s_σ, s_τ, s are respectively in state O, O, O ,
- (2) s_σ, s_τ, s are respectively in state O, P, P ,
- (3) s_σ, s_τ, s are respectively in state P, O, P .

Proof. Straightforward by induction on s , using reachable sequentiality of σ and τ . \square

Again, this is the familiar state diagram for alternating plays on a tensor of strategies, see e.g. [Har04]. Finally, there is a state diagram for the functorial action of the exponential – as for tensor, if $s \in \mathcal{C}\text{-Plays}(!\sigma)$, we write s_i for its restriction on copy index i .

Lemma 6.12. *Consider $\sigma : A \vdash B$ be a sequential (pre)strategy.*

For any $s \in \downarrow\uparrow\text{-Plays}(!\sigma)$, then for any $i \in \mathbb{N}$, $s_i \in \downarrow\uparrow\text{-Plays}(\sigma)$; and we are in one of:

- (1) s has state O , and for all $i \in \mathbb{N}$, s_i has state O ,
- (2) s has state P , and there exists a unique $i \in \mathbb{N}$ such that s_i has state P .

Proof. Straightforward by induction on s , using reachable sequentiality of σ and τ . \square

As for composition, the preservation of *reachable sequentiality* and *sequential determinism* under composition are immediate applications. We omit the easy verifications that the alternating projection is compatible with tensor and bang; from which – as for composition – it follows that sequential visibility is preserved. All structural strategies involved in the Seely category structure, being variants of copycat, are easily proved sequential. Overall, we have:

Proposition 6.13. *There is a Seely category $\rightarrow\text{-Strat}^{\text{seq}}$ of \rightarrow -arenas and sequential strategies. Moreover, $\downarrow\uparrow\text{-Unf}(-) : \rightarrow\text{-Strat}^{\text{seq}} \rightarrow \downarrow\uparrow\text{-Strat}^{\text{vis}}$ preserves \approx and the Seely structure.*

6.2. Full Abstraction for IA. Next, we fine-tune $\rightarrow\text{-Strat}^{\text{seq}}$ to get full abstraction for IA.

6.2.1. Interpretation of IA. It only remains to prove that the interpretation of the primitives of IA, i.e. all primitives of IA_{\parallel} except for the parallel let, are sequential. We have:

Lemma 6.14. *The strategies **seq**, **succ**, **if**, **iszero**, **let**, **assign**, **deref**, **grab**, **release** and the prestrategies **cell** _{n} and **lock** _{n} are sequential. Moreover, for each of those (pre)strategies σ , $\downarrow\uparrow\text{-Unf}(\sigma)$ is the corresponding alternating strategy from Sections 2.3.5 and 3.4.3.*

Proof. Routine verification. \square

It follows that we have an adequate interpretation of IA as sequential strategies, and:

Proposition 6.15. *For any $\Gamma \vdash M : A$ in IA, we have $\llbracket M \rrbracket_{\downarrow\uparrow\text{-Strat}_1^{\text{vis}}} = \downarrow\uparrow\text{-Plays}(\llbracket M \rrbracket_{\rightarrow\text{-Strat}_1^{\text{seq}}})$.*

Proof. Straightforward by induction on the derivation $\Gamma \vdash M : A$. \square

The two adequate models of IA, $\downarrow\uparrow\text{-Strat}_1^{\text{vis}}$ and $\rightarrow\text{-Strat}_1^{\text{seq}}$, differ in crucial ways: $\rightarrow\text{-Strat}_1^{\text{seq}}$ is much more expressive, and records intensional *causal* information. Secondly, in $\rightarrow\text{-Strat}_1^{\text{seq}}$ one can also read the behaviour of the program under contexts outside IA.

However, $\rightarrow\text{-Strat}_1^{\text{seq}}$ is not fully abstract for IA – we need to deal with well-bracketing.




Figure 53: Definition 3.20 is weaker than Definition 3.2

6.2.2. *Well-bracketing.* Ideally, we would love to have an interpretation-preserving functor

$$\downarrow\uparrow\text{-Plays}(-) : \rightarrow\text{-Strat}_!^{\text{wb,seq}} \rightarrow \downarrow\uparrow\text{-Strat}_!^{\text{wb,vis}},$$

but Definition 3.20 (for non-alternating plays) is weaker than Definition 3.2 (for alternating plays) when applied on alternating plays, as illustrated in Figure 53. While Definition 3.2 closely follows the operational idea that calls and returns are handled by a single stack, Definition 3.20 restricts the hierarchical relationship between calls and returns.

Fortunately, from a distinguishing test in $\rightarrow\text{-Strat}_!^{\text{wb,seq}}$ one can extract a characteristic *complete* (see Section 3.4.4) alternating play, which – as we shall see – is well-bracketed as in Definition 3.2. This is due to the following lemma, a well-known observation:

Lemma 6.16. *Let $s \in \downarrow\uparrow\text{-Plays}(A)$ be P- and O-visible. Assume that s has the form*

$$s = \dots s_i \overset{\dots}{\dots} s_j \dots$$

where no further move points to s_j . Then, no move after s_j can point within $s_i \dots s_j$.

Proof. By P- or O-visibility, s_{j+1} points strictly before s_i . Then no view can ever see $s_{i+1} \dots s_j$ – so no move can point there. Besides, s_i can only be seen by the player responsible for it, so no move can point to s_i . A proof appears in [CH10, Lemma 5]. \square

From this, we may easily deduce the following:

Lemma 6.17. *Any $s \in \downarrow\uparrow\text{-Plays}(A)$ complete is well-bracketed in the sense of Definition 3.2.*

Proof. Consider $s \in \downarrow\uparrow\text{-Plays}(A)$ complete but with a well-bracketing failure, *i.e.* as in:

$$s = \dots \mathbf{q}_1^Q \overset{\dots}{\dots} \mathbf{q}_2^Q \dots a^A \dots$$

with \mathbf{q}_2 unanswered when playing a . By *answer-closing*, no further move can point to a . Thus by Lemma 6.16, no further move can point to \mathbf{q}_2 , contradicting that s is complete. \square

The play of Figure 53 is well-bracketed as in Definition 3.20, but it cannot be extended to a complete play: the two questions covered will not be addressed ever again.

6.2.3. *Full abstraction.* From that, we may finally conclude:

Theorem 6.18. *The model $\rightarrow\text{-Strat}_!^{\text{wb,seq}}$ is intensionally fully abstract for \mathbf{IA} .*

Proof. Let $\vdash M, N : A$ be terms in \mathbf{IA} , and assume that $\llbracket M \rrbracket \uparrow \llbracket N \rrbracket$, *i.e.* there is a test $\alpha \in \rightarrow\text{-Strat}_!^{\text{wb,seq}}(\llbracket A \rrbracket, \llbracket U \rrbracket)$ such that $\alpha \odot_! \llbracket M \rrbracket \neq \alpha \odot_! \llbracket N \rrbracket$ – assume *w.l.o.g.* that $\alpha \odot_! \llbracket M \rrbracket$ converges while $\alpha \odot_! \llbracket N \rrbracket$ diverges. Writing $\alpha' = \downarrow\uparrow\text{-Unf}(\alpha)$, it follows that

$$\alpha' \odot_! \llbracket M \rrbracket \downarrow_{\downarrow\uparrow\text{-Strat}_!^{\text{wb,vis}}} \quad \alpha' \odot_! \llbracket N \rrbracket \uparrow_{\downarrow\uparrow\text{-Strat}_!^{\text{wb,vis}}},$$

but α' may not be well-bracketed as in Definition 3.3. Consider $s \in \alpha'$ involved in $\alpha' \odot \llbracket M \rrbracket \Downarrow$ – until the rest of the proof, $\llbracket M \rrbracket$ is the interpretation in $\Downarrow\text{-Strat}_!^{\text{wb,vis}}$. The initial question of s has an answer, thus as $\llbracket M \rrbracket$ and α' are well-bracketed in the sense of 3.21, all its questions are answered. It is P-visible and O-visible since both $\llbracket M \rrbracket$ and α' are P-visible. Hence, it is complete, and so by Lemma 6.17, it is well-bracketed in the sense of Definition 3.3.

Consider α' restricted to (plays symmetric to) prefixes of s . Now α' is well-bracketed as in Definition 3.2, and it distinguishes $\llbracket M \rrbracket$ and $\llbracket N \rrbracket$, hence $M \dagger N$ by Theorem 3.14. \square

6.3. Sequential Innocence. A causal strategy $\sigma : A$ is **sequential innocent** if it is both *sequential* and *parallel innocent*; those form a Seely category $\rightarrow\text{-Strat}^{\text{wb,seq,inn}}$.

We already know that $\rightarrow\text{-Strat}^{\text{wb,seq,inn}}$ supports an adequate interpretation of PCF. For (intensional) full abstraction, we shall prove that $\Downarrow\text{-Unf}(-)$ sends sequential innocent strategies to innocent alternating strategies as in Definition 3.6, and rely on Theorem 3.10.

6.3.1. Causal analysis of sequential innocence. First, we shall see that sequential innocent causal strategies are really representations of the *P-view forests* of Section 3.2.

Lemma 6.19. *Consider $\sigma : A$ a sequential, parallel innocent causal strategy.*

Then, σ is an O-branching alternating forest.

Proof. First, we prove that for all $m \in |\sigma|$, its set of dependencies $[m]_\sigma$ is a total order.

Seeking a contradiction, take $m' \in |\sigma|$ minimal with $m' \rightarrow_\sigma m_1$ and $m' \rightarrow_\sigma m_2$ distinct, all within $[m]_\sigma$. By minimality, $[m']_\sigma$ is a total order, *i.e.* a gcc. By Lemma A.3, m_1 and m_2 have the same polarity, opposite of m' . Consider $\rho_1 \in \text{gcc}(\sigma)$ a gcc for m passing through $m' \rightarrow_\sigma m_1$, and $\rho_2 \in \text{gcc}(\sigma)$ a gcc for m passing through $m' \rightarrow_\sigma m_2$. Then ρ_1 and ρ_2 have least distinct events m_1 and m_2 ; hence by pre-innocence m_1 and m_2 are positive.

Now, m' is the only immediate dependency of m_1 and m_2 ; indeed if there was $m'' \rightarrow_\sigma m_i$, then considering $\rho' \rightarrow m_i \in \text{gcc}(\sigma)$ passing through m'' , ρ and ρ' would fork at some event smaller than m' , contradicting its minimality. Hence, $[m'] \cup \{m_i\} \in \mathcal{C}(\sigma)$ for $i \in \{1, 2\}$.

Also writing $[m']$ for the play in $\Downarrow\text{-Plays}(\sigma)$ with events in the same order, we have

$$[m'], [m']m_1, [m']m_2 \in \cup\text{-Plays}(\sigma),$$

but by Lemma A.3, $[m']m_1^+$ and $[m']m_2^+$ are alternating. By *sequential determinism* of σ , it follows that $m_1 = m_2$, contradiction. So, for all $m \in |\sigma|$, $[m]_\sigma$ is a total order.

Thus $(|\sigma|, \leq_\sigma)$ is a forest. Likewise, if $m^- \rightarrow m_1^+$ and $m^- \rightarrow m_2^+$ in σ , by *sequential determinism* and the same reasoning as above, $m_1 = m_2$, so σ is O-branching. Finally, as for any causal strategy $\sigma : A$ on A alternating, we have \rightarrow_σ is alternating as well. \square

Let us call a **branch** of sequential innocent $\sigma : A$ a $s = m_1 \dots m_n \in \Downarrow\text{-Plays}(\sigma)$ s.t.

$$m_1 \rightarrow_\sigma \dots \rightarrow_\sigma m_n \in \text{gcc}(\sigma).$$

Then, $\partial_\sigma(s) \in \Downarrow\text{-Plays}(A)$, but there is more: by courtesy of σ , if $m_i^+ \rightarrow_\sigma m_{i+1}^-$ then $\partial_\sigma(m_i) \rightarrow_A \partial_\sigma(m_{i+1})$, *i.e.* $\partial_\sigma(m_{i+1})$ points to $\partial_\sigma(m_i)$ in $\partial_\sigma(s)$. In other words, $\partial_\sigma(s)$ is actually a *P-view*, *i.e.* an alternating play where Opponent always points to the previous move. One can *display* a sequential innocent causal strategy to a forest of P-views as in Section 3.3.1: we have recovered, as the causal structure of sequential innocent causal strategies, the *forest of P-views* “ σ ”, the “causal presentation” of alternating innocent




Figure 54: Innocence witness for Figure 16




Figure 55: An alternative presentation

strategies from Proposition 3.7. This connection confirms that the *causal strategies* of Section 5 are generalizations of the sets of *P-views* of traditional innocence²³.

6.3.2. *P-views are the causal dependency.* This gives two ways to get *plays* from $\sigma : A$ sequential innocent: following Section 3.3.1, by selecting those plays whose P-views appear in the causal representation; and following Section 5, as displays of alternating plays over σ – we must prove them identical. Figure 54 shows the augmentation explored in the play of

$$\llbracket \lambda f^{\mathbb{U} \rightarrow \mathbb{U}}. f \text{ skip}; f \text{ skip}; \text{tt} \rrbracket : \llbracket (\mathbb{U} \rightarrow \mathbb{U}) \rightarrow \mathbb{B} \rrbracket$$

in Figure 16 – the numbers in red correspond to the order in which that augmentation is explored. The reader should take some time to digest this picture, and in particular observe that for each prefix of the play in Figure 16, the P-view is exactly (the display of) the branch leading to the corresponding move in the configuration. Opponent could explore the same configuration in a different order, corresponding to a different play with the same P-views. On the other hand, only Opponent has any degree of freedom in this exploration: Player has ever at most one possible move, that immediately caused by the last Opponent move²⁴.

As an aside, in Figure 55, we give an alternative presentation of the same augmentation, making explicit how an augmentation of a sequential innocent strategy consists is the underlying configuration (here, Figure 11), enriched with immediate causal links. The set of (isomorphism classes of) configurations reached is essentially the information recorded by the relational model; so this presentation shows plainly that innocent game semantics consist in enriching the relational model with explicit causal / temporal information.

Back to the technical development, we must prove that if $\sigma : A$ is sequential innocent, then $\Downarrow\text{-Unf}(\sigma)$ is innocent as in Definition 3.6. This relies on a link between P-views and the causal structure of $\sigma : A$, which should be expected in the light of Figure 54.

²³This shows that the causal reasoning permitted by traditional innocence, one of the main tools of traditional game semantics, is *not* inherently restricted to innocence. This is a powerful observation, and much of the subsequent line of work in concurrent games has consisted in exploring its implications.

²⁴Different explorations of the same augmentation may be related by permuting contiguous OP pairs of moves. Deterministic innocent strategies may be defined as those stable under the permutations of OP pairs permitted by the arena: this is the idea behind Melliès' presentation of innocence [Mel04].

Lemma 6.20. *Consider A, B --arenas and $\sigma : A \vdash B$ sequential innocent.*

Then, for any $tm \in \downarrow\uparrow\text{-Plays}(\sigma)$, we have $\ulcorner \partial_{\Lambda(\sigma)}(tm) \urcorner = \partial_{\Lambda(\sigma)}([m]_{\sigma})$.

Proof. In the statement above, we treat $[m]_{\sigma}$ as the sequence induced by its total ordering.

The crucial observation is that if $tm^{-}n^{+} \in \downarrow\uparrow\text{-Plays}(\sigma)$, then necessarily $m^{-} \rightarrow_{\sigma} n^{+}$. To prove that, we prove by induction on t that for any $t \in \downarrow\uparrow\text{-Plays}(\sigma)$: (1) if t has even length, then all maximal events of $|t| \in \mathcal{C}(\sigma)$ are positive; and (2) if t has odd length, then $|t| \in \mathcal{C}(\sigma)$ has *exactly one* maximal negative event. Indeed, for $tm^{-} \in \downarrow\uparrow\text{-Plays}(\sigma)$, then t has even length, so $|t|$ has all its maximal events positive. But then $|tm^{-}|$ has exactly one maximal negative event, namely m^{-} . Likewise, for $tm^{+} \in \downarrow\uparrow\text{-Plays}(\sigma)$, then $|t|$ has exactly one maximal negative event. Now, the immediate predecessor of m must be negative. But if it is not maximal in $|t|$, this contradicts Lemma 6.19, and in particular the fact that σ is O -branching. Therefore, the predecessor of m must be the unique maximal negative event of $|t|$, and $|tm|$ has all maximal events positive as required. Now, if $tm^{-}n^{+}$, then $|tm^{-}|$ has exactly one maximal negative event (namely m^{-}); while the maximal events of $|tm^{-}n^{+}|$ are all positive (and comprise n^{+}). Hence, $m^{-} \rightarrow_{\sigma} n^{+}$ as required.

Likewise, if $t_1m^{+}t_2n^{-} \in \downarrow\uparrow\text{-Plays}(\sigma)$ s.t. $\partial_{\Lambda(\sigma)}(m) \rightarrow_A \partial_{\Lambda(\sigma)}(n)$ – so $\ulcorner \partial_{\Lambda(\sigma)}(t_1mt_2n) \urcorner = \ulcorner \partial_{\Lambda(\sigma)}(t_1) \urcorner \partial_{\Lambda(\sigma)}(m) \partial_{\Lambda(\sigma)}(n)$ then we must have $m \rightarrow_{\sigma} n$ by Lemma A.4. From these two facts, the lemma is a direct verification by induction on t . \square

6.3.3. *Innocent alternating unfolding.* From the above, we may now deduce:

Proposition 6.21. *Consider A, B --arenas, and $\sigma : A \vdash B$ a sequential causal strategy.*

If σ is parallel innocent, then $\downarrow\uparrow\text{-Unf}(\sigma)$ is innocent as in Definition 3.6.

Proof. First, we must show that $\downarrow\uparrow\text{-Unf}(\sigma)$ is P-visible. In other words, we must prove that for all $s \in \downarrow\uparrow\text{-Unf}(\sigma)$, $\ulcorner s \urcorner \in \downarrow\uparrow\text{-Plays}(A \multimap B)$. For the empty play there is nothing to prove; so consider $sa \in \downarrow\uparrow\text{-Unf}(\sigma)$ and $tm \in \downarrow\uparrow\text{-Plays}(\sigma)$ such that $sa = \partial_{\Lambda(\sigma)}(tm)$. Now, by Lemma 6.20, we have $\ulcorner sa \urcorner = \partial_{\Lambda(\sigma)}([m]_{\sigma})$, as plays – therefore $\ulcorner sa \urcorner \in \downarrow\uparrow\text{-Plays}(A \multimap B)$ as required.

Now, we prove innocence. Let $sa^{+}, s' \in \downarrow\uparrow\text{-Unf}(\sigma)$ such that $\ulcorner s \urcorner = \ulcorner s' \urcorner$. By definition, there is $tm^{+} \in \downarrow\uparrow\text{-Plays}(\sigma)$ and $t' \in \downarrow\uparrow\text{-Plays}(\sigma)$ such that $sa = \partial_{\Lambda(\sigma)}(tm)$ and $s' = \partial_{\Lambda(\sigma)}(t')$.

Now, by Lemma 6.20, $\ulcorner sa \urcorner = \partial_{\Lambda(\sigma)}([m]_{\sigma})$ as plays, hence also $\ulcorner s \urcorner = \partial_{\Lambda(\sigma)}([n]_{\sigma})$ for $n \rightarrow_{\sigma} m$. Again by Lemma 6.20, $\ulcorner s' \urcorner = \partial_{\Lambda(\sigma)}([n']_{\sigma})$ for n' the last move of t' . Since $\ulcorner s \urcorner = \ulcorner s' \urcorner$, by Lemma 6.3, $[n]_{\sigma} = [n']_{\sigma}$. So, $|t'| \cup \{m\}$ is down-closed. Finally, $|t'| \cup \{m\}$ is negatively compatible since $|t'| \in \mathcal{C}(\sigma)$ and m is positive, hence $|t'| \cup \{m\}$ is compatible as σ satisfies *causal determinism*. Therefore, $t'm \in \downarrow\uparrow\text{-Plays}(\sigma)$, and $s'm = \partial_{\Lambda(\sigma)}(t'm) \in \downarrow\uparrow\text{-Unf}(\sigma)$. \square

Altogether, we have proved the following proposition:

Proposition 6.22. *There is a Seely category $\rightarrow\text{-Strat}^{\text{seq,inn}}$ of --arenas and sequential strategies. Moreover, the alternating unfolding preserves \approx and the Seely category structure:*

$$\downarrow\uparrow\text{-Unf}(-) : \rightarrow\text{-Strat}^{\text{seq,inn}} \rightarrow \downarrow\uparrow\text{-Strat}^{\text{inn}}.$$

Proof. It suffices to prove that the functor of Proposition 6.13 sends sequential innocent causal strategies to innocent alternating strategies, which we know by Proposition 6.21. \square

We are now equipped to show full abstraction for PCF.

6.3.4. *Full abstraction for PCF.* We show that $\rightarrow\text{-Strat}^{\text{wb,seq,inn}}$ is fully abstract for PCF.

Theorem 6.23. *The model $\rightarrow\text{-Strat}^{\text{wb,seq,inn}}$ is intensionally fully abstract for PCF.*

Proof. Let $\vdash M, N : A$ be terms in PCF s.t. $\llbracket M \rrbracket \uparrow \llbracket N \rrbracket$, i.e. there is $\alpha : \llbracket A \rrbracket \vdash \mathbb{U}$ sequential innocent and well-bracketed such that, w.l.o.g., $\alpha \odot \llbracket M \rrbracket \downarrow$ and $\alpha \odot \llbracket N \rrbracket \uparrow$. Then, it follows

$$\downarrow\uparrow\text{-Unf}(\alpha) \odot \llbracket M \rrbracket_{\downarrow\uparrow\text{-Strat}_1} \downarrow, \quad \downarrow\uparrow\text{-Unf}(\alpha) \odot \llbracket N \rrbracket_{\downarrow\uparrow\text{-Strat}_1} \uparrow,$$

with $\downarrow\uparrow\text{-Unf}(\alpha)$ well-bracketed by Proposition 3.15. Hence, $M \uparrow N$ by Theorem 3.10. \square

7. FINITE DEFINABILITY AND FULL ABSTRACTION FOR PCF_{//}

We have now established the following intensional full abstraction results:

$$\begin{array}{ll} \rightarrow\text{-Strat}_1^{\text{wb}} & \text{is fully abstract for } \text{IA}_{//}, \\ \rightarrow\text{-Strat}_1^{\text{wb}} + \textit{sequentiality} & \text{is fully abstract for } \text{IA}, \\ \rightarrow\text{-Strat}_1^{\text{wb}} + \textit{parallel innocence} + \textit{sequentiality} & \text{is fully abstract for } \text{PCF}, \end{array}$$

and we are left with the one outstanding objective:

$$\rightarrow\text{-Strat}_1^{\text{wb}} + \textit{parallel innocence} \quad \text{is fully abstract for } \text{PCF}_{//}.$$

Unfortunately, this is also the most challenging of our full abstraction results: whereas for the others we could leverage earlier work, we must prove finite definability from scratch.

Proving finite definability for parallel innocent strategies involves many steps. In Section 7.1, we introduce a more convenient equivalence between parallel innocent strategies, *positional equivalence*. In Section 7.2 we show it suffices to consider tests that satisfy a stronger, *causal*, form of well-bracketing useful for definability. In Section 7.3 we introduce a notion of finiteness, and show that finite tests suffice. In Section 7.4, we show a factorization result, reducing finite definability to that for first-order strategies. In Section 7.5, we conclude the proof of finite definability. Finally, in Section 7.6, we prove intensional full abstraction for PCF_{//}, concluding the technical contents of the paper.

7.1. The Positional Collapse. Definability will hold only with respect to *positional equivalence*, a congruence amounting to an equal projection in the relational model.

7.1.1. *Positions of arenas.* We will observe strategies only on certain *positions* of arenas.

Definition 7.1. Let A be an arena, and $x \in \mathcal{C}(A)$.

We say that x is **complete** iff every question in x has an answer in x .

Complete configurations mirror the *complete plays* of Section 3.4.4 and onwards. In both cases, all function calls have returned. The difference, however, is that complete plays are sequential whereas complete configurations are not: they are a “static” snapshot presenting all calls and returns and their hierarchical relationships, with no temporal information.

Besides, complete configurations also comprise the ad-hoc choice of copy indices for all replicable moves. So we quotient them out via the following variation.

Definition 7.2. Let A be an arena. The set of **positions** on A , ranged over by x, y, \dots , is:

$$fA = \{x \in \mathcal{C}(A) \mid x \text{ complete}\} / \cong_A.$$

If $x \in \mathcal{C}(A)$ is complete, we write $[x]_{\cong} \in fA$ for its symmetry class.




Figure 56: Quotienting out evaluation order

For instance, the configuration of Figure 11 is complete. The matching position is described intuitively by removing the grey subscripts, *i.e.* the copy indices.

7.1.2. Positions of strategies. The *positional collapse* of a strategy is an explicit *desequentialization*, obtained by forgetting the chronological ordering of complete plays.

Definition 7.3. Consider A an arena, and $\sigma : A$ a causal strategy. The **positions** of σ are

$$f\sigma = \{x \in fA \mid \exists x \in \mathcal{C}(\sigma), x = [\partial_\sigma(x)]_\cong\}.$$

For $\sigma : A \vdash B$, positions are symmetry classes of parallel compositions $x_A \parallel x_B$, also written $x_A \parallel x_B$ for $x_A = [x_A]_\cong, x_B = [x_B]_\cong$. They correspond to pairs $(x_A, x_B) \in (fA) \times (fB)$ – so $f\sigma$ gives a *relation* from fA to fB ; accordingly we also write $(x_A, x_B) \in f\sigma$ for $x_A \parallel x_B \in f\sigma$.

We illustrate the construction in Figure 56: the example illustrates how, by only keeping complete positions, the collapse forgets the evaluation order. We define:

Definition 7.4. Two causal strategies $\sigma, \tau : A$ are **positionally equivalent** iff $f\sigma = f\tau$.

We write $\sigma \equiv \tau$ to denote the fact that σ and τ are positionally equivalent.

This is a drastic quotient, identifying sequential and parallel evaluation. It will help tremendously in our definability procedure, which will not respect the evaluation order. Of course, the more drastic the quotient, the more challenging the corresponding proof obligation that it is preserved by operations on strategies – and in particular by composition.

7.1.3. Deadlocks. Stability of positional equivalence by composition boils down to

$$f(-) : \multimap\text{-Strat} \rightarrow \text{Rel}$$

being functorial, where Rel is the usual category of *sets and relations*.

For morphisms $\sigma \in \multimap\text{-Strat}(A, B)$, we have defined $f\sigma \in \text{Rel}(fA, fB)$. But for now, this operation has no reason to preserve composition! In fact neither inclusion holds: firstly, $f(\tau \odot \sigma) \subseteq (f\tau) \odot (f\sigma)$ may fail as a complete configuration may arise through an interaction involving a non-complete configuration on B . We shall see later on that this can be salvaged by restricting to *well-bracketed* causal strategies, ensuring that an interaction producing a complete configuration on $A \vdash C$ will only involve a complete configuration on B .

However the other direction also fails, and the diagnosis is more serious. To construct $(x_A, x_C) \in (f\tau) \odot (f\sigma)$, one provides $x_B \in fB$ mediating for relational composition, so

$$(x_A, x_B) \in f\sigma, \quad (x_B, x_C) \in f\tau,$$




Figure 57: Deadlocking composition of causal strategies

i.e. with $x^\sigma \in \mathcal{C}(\sigma)$ and $x^\tau \in \mathcal{C}(\tau)$ s.t. writing $\partial_\sigma x^\sigma = x_A^\sigma \parallel x_B^\tau$ and $\partial_\tau x^\tau = x_B^\tau \parallel x_C^\tau$, $x_A^\sigma \in \times_A$, $x_C^\tau \in \times_C$, and $x_B^\sigma, x_B^\tau \in \times_B$, so $x_B^\sigma \cong_B x_B^\tau$ match up to symmetry. In other words, we must provide a *witness position* that both strategies agree (up to symmetry) is *reachable*.

On the other hand, composition of strategies is more rigid: not only should the projections of x^σ and x^τ on B match, they should also arrive at this position *in the same chronological ordering*. This is not always possible: these two notions of composition differ when interaction triggers a *causal deadlock*, *i.e.* pairs of configurations that are *matching* but not *secured* as in Definition 4.5. Figure 57 displays an example: the strategy obtained by composition has no response to the initial Opponent move, while relational composition authorizes 0^+ .

This strikes at the heart of the difference between game and relational semantics: the former is dynamic hence sensitive to deadlocks, while the latter is static. This of course is what lets game semantics model languages with non-commutative effects, but for us, very concretely, it means that positional equivalence is in general not a congruence.

7.1.4. *The deadlock-free lemma.* Our *deus ex machina* is visibility. A powerful – and at first unexpected – consequence of visibility is that any interaction between visible strategies is always deadlock-free. The consequence of visibility that our proof will exploit repeatedly is:

Lemma 7.5. *Consider A, B –arenas, $\sigma : A \vdash B$ visible, and $m, m' \in |\sigma|$ s.t. $m <_\sigma m'$.*

Then, $\text{just}(m')$ is comparable with m with respect to \leq_σ .

Proof. Since $m <_\sigma m'$, there is $\rho \rightarrow m' \in \text{gcc}(\sigma)$ s.t. $m \in \rho$. If $\partial_\sigma(m')$ is minimal in A , $\partial_\sigma(\text{just}(m'))$ is minimal in B , so $\text{just}(m')$ is minimal for \leq_σ by courtesy. But since σ is pointed, $\text{just}(m')$ is the initial move of ρ , obviously comparable with m as ρ is totally ordered.

Else, by visibility $\text{just}(m') \in \rho$. But ρ is totally ordered, so $m, \text{just}(m')$ comparable. \square

We shall prove that the composition of visible causal strategies is deadlock-free. But first, we recall the basic mechanisms of interactions between causal strategies. Consider $\sigma : A \vdash B$ and $\tau : B \vdash C$, and configurations $x^\sigma \in \mathcal{C}(\sigma)$, and $x^\tau \in \mathcal{C}(\tau)$ such that, writing $\partial_\sigma x^\sigma = x_A^\sigma \parallel x_B^\tau$ and $\partial_\tau x^\tau = x_B^\tau \parallel x_C^\tau$, we have $x_B^\sigma = x_B^\tau = x_B$, *i.e.* x^σ and x^τ are *matching*. Then, recall from Definition 4.5 the bijection arising from their synchronization:

$$\varphi : x^\sigma \parallel x_C^\tau \xrightarrow{\partial_\sigma \parallel x_C^\tau} x_A^\sigma \parallel x_B \parallel x_C^\tau \xrightarrow{x_A^\sigma \parallel \partial_\tau^{-1}} x_A^\sigma \parallel x^\tau,$$

whose graph is equipped with a relation importing the causal dependencies from σ and τ :

$$(l, r) \triangleleft (l', r') \iff l <_{\sigma \parallel C} l' \vee r <_{A \parallel \tau} r'.$$

We saw in Definition 4.5 and Proposition 4.7 that (x^σ, x^τ) corresponds to a configuration of the interaction $\tau \circledast \sigma$ exactly when this bijection is *secured*, *i.e.* \triangleleft is acyclic.

If $\sigma : A \vdash B$ and $\tau : B \vdash C$ are visible, we claim that this is always the case. We reason by contradiction: starting with a putative deadlock, we repeatedly push it down the causal dependency of the arena, until it reaches a minimal event – but those cannot appear in a cycle. Before giving the formal proof, we showcase the reasoning on a simplified case.

Consider a **simple deadlock** in φ , given by $p_1 = (l_1, r_1)$ and $p_2 = (l_2, r_2) \in \varphi$ such that

$$l_1 <_{\sigma \parallel C} l_2, \quad r_2 <_{A \parallel \tau} r_1,$$

an immediate causal incompatibility between p_1 and p_2 . In other words we have $p_1 \triangleleft p_2$ and $p_2 \triangleleft p_1$, and we use $p_1 \triangleleft_{\sigma} p_2$ and $p_2 \triangleleft_{\tau} p_1$ to indicate the origin of the causal constraint. Finally, we apply the same conventions for polarity of elements of φ as in Section 4.2.6.

The first observations (skipped here) is that *w.l.o.g.*, the polarities are as in

$$p_1^r \begin{array}{c} \curvearrowright_{\sigma} \\ \curvearrowleft_{\tau} \end{array} p_2^l,$$

where both occur in B but not minimal in B – so we may take $\text{just}(p_i) = (\text{just}(l_i), \text{just}(r_i))$. By Lemma 7.5, l_1 and $\text{just}(l_2)$ are comparable for σ ; while r_2 and $\text{just}(r_1)$ are comparable for τ . If $p_1 \triangleleft_{\sigma} \text{just}(p_2)$ or $p_2 \triangleleft_{\tau} \text{just}(p_1)$, then we respectively have one of the cycles:

$$p_1^r \begin{array}{c} \curvearrowright_{\sigma} \text{just}(p_2)^r \\ \curvearrowleft_{\tau} \end{array} p_2^l \quad \text{or} \quad p_1^r \begin{array}{c} \curvearrowright_{\sigma} \\ \curvearrowleft_{\tau} \text{just}(p_1)^l \end{array} p_2^l$$

so simple deadlocks between p_1 and $\text{just}(p_2)$; or between $\text{just}(p_1)$ and p_2 . The cumulative *depth* in B has decreased. The case $p_1 = \text{just}(p_2)$ or $p_2 = \text{just}(p_1)$ is easily discarded.

The last case has $\text{just}(p_2) \triangleleft_{\sigma} p_1$ and $\text{just}(p_1) \triangleleft_{\tau} p_2$. But p_1 has polarity r , so by Lemma 5.8 the only immediate dependency in σ of l_1^- is $\text{just}(l_1)$. So $\text{just}(p_2) \triangleleft_{\sigma} p_1$ factors as $\text{just}(p_2) \triangleleft_{\sigma} \text{just}(p_1) \triangleleft_{\sigma} p_1$. Symmetrically $\text{just}(p_1) \triangleleft_{\tau} \text{just}(p_2)$, so we have:

$$\text{just}(p_1) \begin{array}{c} \curvearrowright_{\tau} \\ \curvearrowleft_{\sigma} \end{array} \text{just}(p_2),$$

closer to the root of the arena. Repeating this we eventually hit an impossible simple deadlock with a minimal event in B , finally exposing the contradiction. So visibility *structures* the interaction around the dependency of the arena, giving us an effective reasoning principle.

The proof of the deadlock-free lemma is the same in essence, but challenging in form. Firstly, cycles in \triangleleft in Definition 4.5 may have arbitrary length. Secondly, in relational composition strategies synchronize on *symmetry classes* of configurations rather than concrete configurations; so we must account for synchronization through symmetry.

Lemma 7.6. *Consider A, B, C --arenas, $\sigma : A \vdash B$ and $\tau : B \vdash C$ visible causal strategies, $x^{\sigma} \in \mathcal{C}(\sigma)$ and $x^{\tau} \in \mathcal{C}(\tau)$ with a symmetry $\theta : x_B^{\sigma} \cong_B x_B^{\tau}$. Then, the composite bijection*

$$\varphi : x^{\sigma} \parallel x_C^{\tau} \xrightarrow{\partial_{\sigma} \parallel x_C^{\tau}} x_A^{\sigma} \parallel x_B^{\sigma} \parallel x_C^{\tau} \xrightarrow{x_A^{\sigma} \parallel \theta \parallel x_C^{\tau}} x_A^{\sigma} \parallel x_B^{\tau} \parallel x_C^{\tau} \xrightarrow{x_A^{\sigma} \parallel \partial_{\tau}^{-1}} x_A^{\sigma} \parallel x^{\tau},$$

is secured, in the sense that the relation \triangleleft , defined on the graph of φ with

$$(l, r) \triangleleft (l', r')$$

whenever $l (<_{\sigma} \parallel <_C) l'$ or $r (<_A \parallel <_{\tau}) r'$, is acyclic²⁵.

²⁵For θ an identity, this exactly means that x^{σ} and x^{τ} satisfy the *secured* condition of Definition 4.5.

Proof. We use *polarities* \uparrow, \downarrow or $-$ for elements of φ (i.e. pairs (l, r)) as in Section 4.2.6. We say (l, r) *occurs in* A, B or C in the obvious sense. We use a notion of *justifier* of a pair (l, r) non-minimal in B : as θ is an order-isomorphism, $\partial_\sigma(l)$ is minimal in B iff $\partial_\tau(r)$ is. If not, then $\text{just}(l)$ and $\text{just}(r)$ also match up to θ and $(\text{just}(l), \text{just}(r))$ must be in φ as well – we write it $\text{just}(l, r)$. Suppose now \triangleleft is *not* secured, i.e. there is $((l_1, r_1), \dots, (l_n, r_n))$ with

$$(l_1, r_1) \triangleleft (l_2, r_2) \triangleleft \dots \triangleleft (l_n, r_n) \triangleleft (l_1, r_1),$$

written $p_1 \triangleleft \dots \triangleleft p_n \triangleleft p_1$ – the **length** of this cycle is n . First, *w.l.o.g.* the cycle occurs entirely in B . Assume it has minimal length. If it occurs entirely in A or C , then $(l_i)_{1 \leq i \leq n}$ (resp. $(r_i)_{1 \leq i \leq n}$) is a cycle in σ (resp. τ), absurd. So, it passes through B . Next, if *e.g.*

$$p_i^{(B)} \triangleleft p_{i+1}^{(C)} \triangleleft \dots \triangleleft p_{j-1}^{(C)} \triangleleft p_j^{(B)},$$

then it is easy to prove that $r_i <_\tau r_{i+1} <_\tau \dots <_\tau r_{j-1} <_\tau r_j$, so that $p_i^{(B)} \triangleleft p_j^{(B)}$ and the cycle can be shortened, contradicting its minimality – the same argument holds for A .

We restrict to cycles in B . The **depth** of (l, r) is the length of the chain of justifiers to (l_0, r_0) minimal in B – the depth of (l_0, r_0) minimal in B is 0. The *depth* of the cycle is

$$d = \sum_{1 \leq i \leq n} \text{depth}(l_i, r_i),$$

and we assume *w.l.o.g.* the cycle minimal for the product order on pairs (n, d) . In this proof, all arithmetic computations on indices are done modulo n (the length of the cycle).

Next, let us write $p_i \triangleleft_\sigma p_j$ if $l_i (<_\sigma \| <_C) l_j$ and $p_i \triangleleft_\tau p_j$ symmetrically. We notice that \triangleleft_σ and \triangleleft_τ alternate – if not we shorten the cycle by transitivity, contradicting minimality. We assume *w.l.o.g.* that $p_{2k} \triangleleft_\sigma p_{2k+1}$ and $p_{2k+1} \triangleleft_\tau p_{2k+2}$ for all k . But then, $\text{pol}(p_{2k}) = \downarrow$ and $\text{pol}(p_{2k+1}) = \uparrow$ so that polarity in the cycle is alternating as well. Indeed, assume *e.g.*

$$p_{2k+1} \triangleleft p_{2k+2}^{\uparrow} \triangleleft p_{2k+3}$$

with $p_{2k+1} \triangleleft_\tau p_{2k+2}$ and $p_{2k+2} \triangleleft_\sigma p_{2k+3}$. Then, $r_{2k+1} <_{A \parallel \tau} r_{2k+2}^-$. From its polarity, r_{2k+2}^- cannot be minimal in B . By Lemma 5.8, it has a unique predecessor $\text{just}(r_{2k+2}) \rightarrow_{A \parallel \tau}$, so $r_{2k+1} <_{A \parallel \tau} r_{2k+2}^-$ factors as $r_{2k+1} <_{A \parallel \tau} \text{just}(r_{2k+2}) \rightarrow_{A \parallel \tau} r_{2k+2}^-$. Accordingly, $p_{2k+1} \triangleleft_\tau \text{just}(p_{2k+2}) \triangleleft_\tau p_{2k+2}$ – but dependencies in the game are respected by both strategies, so $\text{just}(p_{2k+2}) \triangleleft_\sigma p_{2k+2}$. So $\text{just}(p_{2k+2}) \triangleleft_\sigma p_{2k+3}$, and we can replace the cycle fragment with

$$p_{2k+1} \triangleleft \text{just}(p_{2k+2}) \triangleleft p_{2k+3}$$

which is still in B , has the same length but strictly smaller depth, contradiction. The symmetric argument applies for σ , so any p_{2k+1} has polarity \uparrow and any p_{2k+2} has polarity \downarrow .

Now we show the cycle cannot have an event minimal in B . Seeking a contradiction, if

$$p_{2k+1}^{\uparrow} \triangleleft_\tau p_{2k+2}^{\downarrow} \triangleleft_\sigma p_{2k+3}^{\uparrow}$$

with p_{2k+2} minimal in B , then $l_{2k+2} <_{\sigma \parallel C} l_{2k+3}$ with l_{2k+2} minimal in B , but then $\partial_{\sigma \parallel C}(l_{2k+2}) <_{A \parallel B \parallel C} \partial_{\sigma \parallel C}(l_{2k+3})$. Indeed, if $\partial_{\sigma \parallel C}(l_{2k+2})$ is minimal in B , l_{2k+2} is (by courtesy) minimal in $\sigma \parallel C$. Likewise, since l_{2k+3} occurs in B , $\partial_{\sigma \parallel C}(l_{2k+3})$ depends (for $\leq_{A \parallel B \parallel C}$) on a unique $\partial_{\sigma \parallel C}(l)$ minimal in B , where l must also be minimal in σ . But since σ is pointed, l_{2k+3} has a unique minimal dependency, hence $l = l_{2k+2}$ and $\partial_{\sigma \parallel C}(l_{2k+2}) <_{A \parallel B \parallel C} \partial_{\sigma \parallel C}(l_{2k+3})$ as claimed. But then, $r_{2k+2} <_{A \parallel \tau} r_{2k+3}$, so $p_{2k+1}^{\uparrow} \triangleleft_\tau p_{2k+2}^{\downarrow} \triangleleft_\tau p_{2k+3}^{\uparrow}$ and again the cycle can be shortened by transitivity, contradicting its minimality.

Now, we have proved a minimal cycle has a canonical form where the strategies alternate, polarity alternates, all events are in B and non-minimal. Since $p_{2k}^r \triangleleft_{\sigma} p_{2k+1}^l$, writing $p = (l, r) = \text{just}(p_{2k+1})$, we have that $l = \text{just}(l_{2k+1})$ as well. From Lemma 7.5, we know that $l = \text{just}(l_{2k+1})$ is comparable with l_{2k} in $\sigma \parallel C$ (by visibility of σ). If $\text{just}(l_{2k+1}) = l_{2k}$, then $r_{2k} \triangleleft_{\tau} r_{2k+1}$ as well. This gives $p_{2k-1} \triangleleft_{\tau} p_{2k+2}$, contradicting minimality of the cycle. So $\text{just}(l_{2k+1}) \neq l_{2k}$. Similarly, $\text{just}(r_{2k+2})$ is comparable with r_{2k+1} in $A \parallel \tau$, but distinct.

Assume that we have $p_{2k} \triangleleft_{\sigma} \text{just}(p_{2k+1})$ for some k . Since $\text{just}(p_{2k+1}) \triangleleft_{\tau} p_{2k+1} \triangleleft_{\tau} p_{2k+2}$ we can replace the cycle fragment $p_{2k} \triangleleft p_{2k+1} \triangleleft p_{2k+2}$ with the cycle fragment

$$p_{2k} \triangleleft \text{just}(p_{2k+1}) \triangleleft p_{2k+2}$$

which has the same length but smaller depth, absurd. So, $\text{just}(p_{2k+1}) \triangleleft_{\sigma} p_{2k}$ for all k (symmetrically, $\text{just}(p_{2k+2}) \triangleleft_{\tau} p_{2k+1}$ for all k). In particular, $\text{just}(l_{2k+1}) \triangleleft_{\sigma \parallel C} l_{2k}^-$ but by Lemma 5.8, l_{2k}^- has a unique immediate predecessor $\text{just}(l_{2k})$. So, $\text{just}(p_{2k+1}) \triangleleft_{\sigma} \text{just}(p_{2k})$ for all k ; and likewise $\text{just}(p_{2k+2}) \triangleleft_{\tau} \text{just}(p_{2k+1})$ for all k . So we can replace the full cycle with

$$\text{just}(p_n) \triangleleft \text{just}(p_{n-1}) \triangleleft \dots \triangleleft \text{just}(p_1) \triangleleft \text{just}(p_n)$$

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

Despite its relatively discreet role in the development, we regard the deadlock-free lemma as one of our main contributions. It is a powerful observation with far-reaching consequences in linking game semantics and relational models. It also gives a lot of weight to the notion of visibility, as a simple, well-behaved and fairly general under-approximation of innocence.

7.1.5. Preservation of composition. For preservation of the positional collapse by composition, we need one further lemma: that any complete position is reachable by a *well-bracketed* play.

Lemma 7.7. *Take $\sigma : A$ visible well-bracketed on $--$ -arena A , $x^{\sigma} \in \mathcal{C}(\sigma)$ with $\partial_{\sigma}(x)$ complete. Then, there is $t \in \mathcal{U}\text{-Plays}(\sigma)$ such that $|t| = x^{\sigma}$ and $\partial_{\sigma}(t)$ is well-bracketed.*

Proof. The idea is simple: since σ is well-bracketed, it suffices to show that $\partial_{\sigma}(x)$ is reachable by a well-bracketed Opponent. But we can set up causal constraints forcing Opponent to be well-bracketed, formulated as a visible causal strategy, and apply the deadlock-free lemma.

For $x \in \mathcal{C}(\sigma)$ s.t. $\partial_{\sigma}(x)$ complete, consider x_A^{σ} as an arena with trivial symmetry. We build $\tau : x_A^{\sigma} \vdash \mathbb{U}$ as $|\tau| = |x_A^{\sigma} \vdash \mathbb{U}|$, and \leq_{τ} as the order of the arena enriched with:

$$\begin{array}{lll} (2, \mathbf{q})^- & \rightarrow_{\tau} & (1, q)^{+, \mathcal{Q}} & \text{if } q \text{ is an initial question in } A, \\ (1, a)^{-, \mathcal{A}} & \rightarrow_{\tau} & (2, \checkmark)^+ & \text{if } a \text{ answers an initial question in } A, \\ (1, a_1)^{-, \mathcal{A}} & \rightarrow_{\tau} & (1, a_2)^{+, \mathcal{A}} & \text{if } \text{just}(a_2) \rightarrow_A \text{just}(a_1), \end{array}$$

resulting in τ visible well-bracketed. By Lemma 7.6, there is a linear ordering of $|x_A^{\sigma} \vdash \mathbb{U}|$ compatible with the constraints of both σ and τ . As both are well-bracketed, its projection on the left gives $s \in \mathcal{U}\text{-Unf}(\sigma)$ such that $|s| = x_A^{\sigma}$ and s well-bracketed as required. \square

Proposition 7.8. *Consider $\sigma : A \vdash B$ and $\tau : B \vdash C$ causal strategies.*

If σ and τ are well-bracketed and visible, then $f(\tau \odot \sigma) = (f\tau) \odot (f\sigma)$.

Proof. \subseteq . Consider $(x_A, x_C) \in f(\tau \odot \sigma)$. By definition, $x_A \parallel x_C = f(\partial_{\tau \odot \sigma}(x^{\tau} \odot x^{\sigma}))$ for $x^{\tau} \odot x^{\sigma} \in \mathcal{C}(\tau \odot \sigma)$ with $\partial_{\tau \odot \sigma}(x^{\tau} \odot x^{\sigma}) = x_A \parallel x_C$ complete, and $x_A = fx_A, x_C = fx_C$. By Lemma 7.7, there is $s \in \mathcal{U}\text{-Unf}(\tau \odot \sigma)$ well-bracketed s.t. $|s| = \partial_{\Lambda(\tau \odot \sigma)}(x^{\tau} \odot x^{\sigma})$. By Proposition 4.30, there is $u \in \mathcal{U}\text{-Unf}(\tau) \otimes \mathcal{U}\text{-Unf}(\sigma)$ s.t. $u \upharpoonright A, C = s$. Now, since σ and τ are well-bracketed and s is well-bracketed, it is direct by induction that u is well-bracketed.

Since $x_A \parallel x_C = \partial_{\tau \odot \sigma}(x^\tau \odot x^\sigma)$ is complete, in particular the initial question has an answer – but as u is well-bracketed, all questions in u are answered. Writing $\partial_\sigma(x^\sigma) = x_A \parallel x_B$ and $\partial_\tau(x^\tau) = x_B \parallel x_C$, x_B is complete as well. But then, $[x_B]_{\cong} = x_B \in \mathcal{J}B$, so x^σ witnesses $(x_A, x_B) \in \mathcal{J}\sigma$ and x^τ witnesses $(x_B, x_C) \in \mathcal{J}\tau$; so $(x_A, x_C) \in (\mathcal{J}\tau) \odot (\mathcal{J}\sigma)$.

⊇. Assume we have symmetry classes of complete configurations x_A, x_B and x_C s.t.

$$(x_A, x_B) \in \mathcal{J}\sigma \quad (x_B, x_C) \in \mathcal{J}\tau,$$

so there are $x^\sigma \in \mathcal{C}(\sigma)$ and $x^\tau \in \mathcal{C}(\tau)$ with $\partial_\sigma(x^\sigma) = x_A^\sigma \parallel x_B^\sigma$, $\partial_\tau(x^\tau) = x_B^\tau \parallel x_C^\tau$, with $x_A^\sigma \in x_A$, $x_B^\sigma, x_B^\tau \in x_B$, and $x_C^\tau \in x_C$. In particular, there is $\theta : x_B^\sigma \cong_B x_B^\tau$. Now,

$$\varphi : x^\sigma \parallel x_C^\tau \xrightarrow{\partial_\sigma \parallel x_C^\tau} x_A^\sigma \parallel x_B^\sigma \parallel x_C^\tau \xrightarrow{x_A^\sigma \parallel \theta \parallel x_C^\tau} x_A^\sigma \parallel x_B^\tau \parallel x_C^\tau \xrightarrow{x_A^\sigma \parallel \partial_\tau^{-1}} x_A^\sigma \parallel x^\tau,$$

is secured by Lemma 7.6. They only match up to symmetry, but by Proposition 4.15, there are $y^\sigma \cong_\sigma x^\sigma$ and $y^\tau \cong_\tau x^\tau$ such that $\partial_\sigma(y^\sigma) = y_A^\sigma \parallel y_B^\sigma$, $\partial_\tau(y^\tau) = y_B^\tau \parallel y_C^\tau$ with $y_B^\sigma = y_B^\tau$ – so $y^\tau \odot y^\sigma \in \mathcal{C}(\tau \odot \sigma)$. But since $y^\sigma \cong_\sigma x^\sigma$ and $y^\tau \cong_\tau x^\tau$, we also know that $y_A^\sigma \in x_A$ and $y_C^\tau \in x_C$ still, so $y^\tau \odot y^\sigma$ witnesses $(x_A, x_C) \in \mathcal{J}(\tau \odot \sigma)$ as required. \square

So positional equivalence of well-bracketed innocent causal strategies is preserved under composition. The other constructions pose no challenge. Altogether, we get:

Corollary 7.9. *There is a Seely category $\rightarrow\text{-Strat}^{\text{wb,vis}}/ \equiv$ of $--$ arenas and positional equivalence classes of visible well-bracketed strategies. Finally, the canonical functor*

$$\rightarrow\text{-Strat}^{\text{wb,vis}} \rightarrow \rightarrow\text{-Strat}^{\text{wb,vis}}/ \equiv$$

preserves the interpretation of PCF $\llbracket \cdot \rrbracket$.

7.1.6. *On the relational model.* As claimed, this does indeed yield a functor to Rel:

Proposition 7.10. *The positional collapse defines a functor $f(-) : \rightarrow\text{-Strat}^{\text{wb,vis}} \rightarrow \text{Rel}$.*

Proof. Composition is Proposition 7.8, while for identity it is a direct verification. \square

Unfortunately, this functor is not compatible with the Seely category structure. For negative arenas A, B we do have $f(A \otimes B) \cong (fA) \times (fB)$; but for instance $f(A \& B) \not\cong (fA) + (fB)$ because $f(A \& B)$ includes the empty position, which we do not have enough information to send to the left or to the right. Likewise, $f(!A) \not\cong \mathcal{M}_f(fA)$. Considering only non-empty configurations does not solve the issue, as we lose $f(A \otimes B) \cong (fA) \times (fB)$.

This mismatch can be mitigated by focusing on the cartesian closed Kleisli categories. Say that an $--$ arena A is **strict** if all its minimal (necessarily negative) events are in pairwise conflict – all types and contexts of PCF are interpreted as strict $--$ arenas. If A is an arena, we write $\mathcal{f}A$ for the set of non-empty complete positions of A . Then:

Lemma 7.11. *For A, B, C, D $--$ arenas with C strict there are bijections:*

$$\begin{aligned} f(A \otimes B) &\cong (fA) \otimes (fB) & \mathcal{f}(A \& B) &\cong \mathcal{f}A + \mathcal{f}B \\ \mathcal{f}(A \multimap C) &\cong (fA) \times (\mathcal{f}C) & f(!C) &\cong \mathcal{M}_f(\mathcal{f}C). \end{aligned}$$

For B, C strict it follows that $\mathcal{f}(!B \multimap C) \cong \mathcal{M}_f(\mathcal{f}B) \times (\mathcal{f}C)$, matching relational semantics. For any PCF type A we obtain $\mathcal{f}\llbracket A \rrbracket \cong \llbracket A \rrbracket_{\text{Rel}}$ – such a bijection is easily established for ground types – so for the interpretation of PCF types, points of the web in the relational model exactly correspond to non-empty complete symmetry classes in the




Figure 58: Non-locality of argument sub-strategies

game semantics. This can be extended to an interpretation-preserving functor from games to relations, but we leave the details to a forthcoming paper on the quantitative collapse.

7.2. Well-Bracketed Pruning. We need a sufficiently constrained description of the causal shape of strategies so that we might replicate it syntactically. While innocence is a causal notion, well-bracketing is not, and it leaves the causal shape too liberal. We shall now see how via a well-bracketed pruning a more causal form of well-bracketing can be enforced.

7.2.1. Causal well-bracketing. Recall the idea of Theorem 3.9: each Player question \mathbf{q}^+ corresponds to a call to a variable x . Opponent questions pointing to \mathbf{q}^+ correspond to x calling an argument; and Opponent answers to \mathbf{q}^+ correspond to x evaluating to a return value. The crux of the definability argument is that these subsequent possibilities of Opponent moves pointing to \mathbf{q}^+ *split* (the causal representation of) the innocent strategy under scrutiny into sub-strategies independent from each other – as they must be, if they are to correspond to distinct branches of the desired syntax tree.

Parallel innocent strategies enjoy the same “splitting” property, to an extent:

Lemma 7.12. *Consider A an arena, $\sigma : A$ an innocent causal strategy, and $q^+ \in |\sigma|$.*

If $q^+ \rightarrow_\sigma m_1^-, q^+ \rightarrow_\sigma m_2^-$ distinct, $\{m \in |\sigma| \mid m \geq_\sigma m_1\}$, $\{m \in |\sigma| \mid m \geq_\sigma m_2\}$ disjoint.

Proof. Obvious by pre-innocence. □

Unfortunately, this is too weak. Indeed q^+ in the statement above should correspond to a syntactic variable occurrence x , so that the overall term has form $C[x M_1 \dots M_n]$ where x has arity n , with the M_i written in PCF_\parallel . As standalone pieces of syntax in a language without interference they are indeed independent from each other (as guaranteed semantically by Lemma 7.12), but are also independent from the context C – which Lemma 7.12 does not capture. In fact, the current conditions on strategies do not ensure this.

In Figure 58, we show a counter-example: this satisfies all the conditions for an augmentation of a well-bracketed parallel innocent strategy. Here, \mathbf{q}_1^+ should correspond to a variable call, and \mathbf{q}^- should initiate the exploration of its argument, and independent sub-term. But the subsequent Player move \mathbf{q}_2^+ , the “head variable occurrence” of the argument sub-term, also depends on a parallel thread. This behaviour is not realizable in PCF_\parallel and should be banned before any definability attempt.


We now define which causal behaviour is deemed acceptable for PCF_\parallel .

Definition 7.13. Consider A an arena, and $\sigma : A$ parallel innocent.

We say that σ is **causally well-bracketed** iff it satisfies the two conditions below:

wb-threads: for every $\rho \in \text{gcc}(\sigma)$, ρ is well-bracketed in the sense of Definition 3.2,


globular: for any diagram in σ with $X = \{m_1, \dots, m_n\}$ and $Y = \{n_1, \dots, n_p\}$ disjoint:



then every question in X (resp. Y) is answered in X (resp. Y).

The Question/Answer labeling is implicitly imported from $|A|$ to $|\sigma|$ via ∂_σ – likewise, for $m^\mathcal{Q}, n^A \in |\sigma|$, we say that n **answers** m if $\partial_\sigma(m) \rightarrow_A \partial_\sigma(n)$, *i.e.* $\partial_\sigma(n)$ answers $\partial_\sigma(m)$.

We call such a diagram a *globule*. This bans directly behaviours as in Figure 58. Only Player can merge parallel threads, and only if he is responsible for the fork (by parallel innocence). So, polarities in the definition of globules are not restrictive. One can further observe that globules always have Question/Answer assignments as in



Indeed, if m_1 was an answer, it would be maximal in σ (by Lemma A.7, as answers are maximal in A) and the merge would be impossible. So m_1 is a question, and by *globular* it has an answer in $\{m_1, \dots, m_{2n}\}$. By courtesy this answer depends immediately on m_1 in σ and so must be m_2 . Repeating this we get the description above. Hence in causally well-bracketed strategies, parallel threads that might merge follow a strict call/return discipline.

A variant of the proof of Proposition 5.14, relying on Lemma 5.13, shows that causal well-bracketing composes. This was done in [CCW15] (see also [Cas17]), but it is not our route here: we must prove full abstraction with respect to the *same* well-bracketing condition used before. We shall see that the situation is similar as in Section 6.2.2: in *complete* configurations, which suffice for tests, the weaker well-bracketing implies the stronger.

7.2.2. Strengthening well-bracketing. We will show that if $\sigma : A$ is innocent and well-bracketed (as in Definition 4.31) and $x^\sigma \in \mathcal{C}(\sigma)$ is complete (*i.e.* all its questions have an answer within x^σ), then the corresponding augmentation is also *causally* well-bracketed – so restricting an innocent well-bracketed $\sigma : A$ to its completable part yields a positionally equivalent, causally well-bracketed strategy; a procedure we call “well-bracketed pruning”.

On arenas arising from types, gccs of innocent strategies are *already* well-bracketed – this holds even without well-bracketing. The proof is morally as for Proposition 3.15.

Lemma 7.14. Consider A a type, $\sigma : \llbracket A \rrbracket$ an innocent causal strategy, and $\rho \in \text{gcc}(\sigma)$.

Then, ρ is well-bracketed in the sense of Definition 3.2.

Proof. Consider $\rho = \rho_1 \rightarrow \dots \rightarrow \rho_n \in \text{gcc}(\sigma)$ and assume, seeking a contradiction, that ρ_n is an answer not pointing to the pending question. It follows that ρ has the form:

$$\rho_1 \rightarrow \dots \rightarrow \rho_i^{-,Q} \rightarrow \dots \rightarrow \rho_j^{+,Q} \rightarrow \rho_{j+1}^{-,Q} \rightarrow \dots \rightarrow \rho_n^{+,A}$$

where $\rho_{j+1}^{-,\mathcal{Q}}$ was pending. As $\llbracket A \rrbracket$ is the interpretation of a type, $\rho_{j+1}^{-,\mathcal{Q}}$ is sibling to countably many symmetric copies, *i.e.* we may write it as b_k with some copy index $k \in \mathbb{N}$, and consider a copy $b_{k+1} \in |\sigma|$ with both b_k and b_{k+1} pointing to ρ_j . Write also $m = \rho_n$.


Consider $x^\sigma \in \mathcal{C}(\sigma)$ with ρ in x^σ , *w.l.o.g.* we can assume that (the augmentation) x^σ has top element m . Consider any $t \in \mathcal{C}\text{-Plays}(x^\sigma)$ s.t. $|t| = x^\sigma$. Then t may be written as $t_1 \cdot b_k \cdot t_2 \cdot m$ where without loss of generality, we may assume that all events after b_k depend on it. Because ρ is in x^σ , m points in t_1 . Now, by uniformity of σ , there is also some

$$t' = t_1 \cdot b_{k+1} \cdot t'_2 \cdot m' \in \mathcal{C}\text{-Plays}(\sigma),$$

such that $t \cong_\sigma t'$, where (therefore) m' points to the same move as m in t_1 . There is some $y^\sigma \in \mathcal{C}(\sigma)$ such that $t' \in \mathcal{C}\text{-Plays}(y^\sigma)$, *w.l.o.g.* assume $|t'| = y^\sigma$. By innocence, any causal branching in x^σ and y^σ is due to Player. Next, we observe that x^σ and y^σ are negatively compatible. Indeed if $n_1^- \in |x^\sigma|$, $n_2^- \in |y^\sigma|$ are in minimal conflict in σ , then by Lemma A.8, $\partial_\sigma(n_1)$ and $\partial_\sigma(n_2)$ are in minimal conflict in $\llbracket A \rrbracket$. But by property of arenas originating from types, this implies that n_1 and n_2 have the same justifier n . As moves in $t_2 \cdot m$ depend on b_k and moves in $t'_2 \cdot m'$ depend on b_{k+1} , by pre-innocence (Lemma 7.12) these must be disjoint, so n appears in t_1 . But again by pre-innocence, since x^σ has top element m and y^σ has top element m' , only one Opponent move can point to n in x^σ ; and likewise for y^σ . As $t \cong_\sigma t'$ preserves pointers, that means that n_1 and n_2 must arrive at the same (chronological) index in t, t' , and so their display in $\llbracket A \rrbracket$ are related by a symmetry in $\llbracket A \rrbracket$. But in arenas arising from PCF types, no conflicting events can be related by a symmetry, contradiction. \square

So condition *wb-threads* of Definition 7.13 is automatic, even without assuming well-bracketing. However, condition *globular* is *not* automatic as illustrated by Figure 58.


We shall now prove *globular* on *complete* augmentations of well-bracketed innocent strategies. Intuitively, the reason is simple. Consider a globule



in x^σ . As x^σ is complete, every question is eventually answered. But after the merge, it is too late for questions in the a_i s and b_j s: if some b such that $a \geq b$ answers a_i then by visibility a_i must appear in all its gccs; but a gcc to b may go through the b_j s and avoids a_i entirely. Of course, turning this idea into a proof takes some work. First we prove:

Lemma 7.15. *Consider a --arena A , $\sigma : A$ visible well-bracketed, and $x^\sigma \in \mathcal{C}(\sigma)$ complete.*

For all $q_1^{-,\mathcal{Q}} \in x^\sigma$, for all $q_2^{+,\mathcal{Q}} \in x^\sigma$ such that $\partial_\sigma(q_1) \rightarrow_A \partial_\sigma(q_2)$, we have:



with a_1 and a_2 the (unique) answers to q_1 and q_2 .

Proof. First, a_1 and a_2 exist as x^σ is complete, and are unique by *answer-linear*.

Seeking a contradiction, assume $a_2 \not\leq_\sigma a_1$, *i.e.* $a_2 \notin [a_1]_\sigma$. As $a_2 \notin [q_2]_\sigma$, writing $y = [a_1]_\sigma \vee [q_2]_\sigma$, $y \in \mathcal{C}(\sigma)$ with $a_1, q_2 \in y$ and $a_2 \notin y$. By Lemma 7.7, there is $t \in \mathcal{C}\text{-Plays}(y)$

well-bracketed s.t. $|t| = y$. So by *wait* of Definition 3.20, q_2 is answered in $|t| = y$. But since $a_2 \notin y$, this means q_2 is answered twice in x^σ , which contradicts *answer-linear*. \square

Using that, we may now formalize and complete the intuitive argument above.

Lemma 7.16. *For $\sigma : \llbracket A \rrbracket$ well-bracketed innocent and $x^\sigma \in \mathcal{C}(\sigma)$ complete, x^σ is globular.*

Proof. Consider a *globule* in x^σ , i.e. a diagram

$$\begin{array}{ccccccc} m_0^- & \rightarrow & m_1^+ & \rightarrow & \dots & \rightarrow & m_n^- & \rightarrow & m^+ \\ & & \searrow & & & & \searrow & & \\ & & n_1^+ & \rightarrow & \dots & \rightarrow & n_p^- & \rightarrow & \end{array}$$

with $X = \{m_1, \dots, m_n\}$ and $Y = \{n_1, \dots, n_p\}$ disjoint. Seeking a contradiction, consider $m_i^{\mathcal{Q}} \in X$ unanswered in X . First, we consider $q^- = \text{just}(m^+)$. Since σ is visible, q^- appears in any gcc of m , so $q^- \leq_\sigma m_0^-$. But hence m is a question, or a gcc like

$$\dots \rightarrow \dots \rightarrow q^-, \mathcal{Q} \rightarrow \dots \rightarrow m_0^- \rightarrow m_1^+ \rightarrow \dots \rightarrow m_n^- \rightarrow m^+, \mathcal{A}$$

fails well-bracketing as m_i is unanswered, forbidden by Lemma 7.14. So m^+ has an answer a^- in x^σ . But q^- also has an answer b^+ in x^σ , and by Lemma 7.15, $a^- \leq_\sigma b^+$. So altogether

$$\dots \rightarrow \dots \rightarrow q^-, \mathcal{Q} \rightarrow \dots \rightarrow m_0^- \rightarrow m_1^+ \rightarrow \dots \rightarrow m_n^- \rightarrow m^+, \mathcal{Q} \rightarrow a^-, \mathcal{A} \rightarrow \dots \rightarrow b^+, \mathcal{A}$$

well-bracketed by Lemma 7.14. So in $m^+ \rightarrow a^- \rightarrow \dots \rightarrow b^+$, some move must answer m_i ; and in particular point to m_i . But m_i does not appear in the gcc $\dots \rightarrow m_0 \rightarrow n_1 \rightarrow \dots \rightarrow n_p \rightarrow m \rightarrow a \rightarrow \dots \rightarrow b$, contradicting visibility. Therefore, X is complete. \square

7.2.3. Well-bracketed pruning. For $x \in \mathcal{C}(\sigma)$, write $x^+ \in \mathcal{C}(\sigma)$ for the greatest $+$ -covered configuration s.t. $x^+ \subseteq x$, obtained by removing trailing Opponent moves.

Proposition 7.17. *For A a $--$ -arena and $\sigma : A$ a well-bracketed innocent strategy, set*

$$|\text{comp}(\sigma)| = \cup \{x \in \mathcal{C}(\sigma) \mid x^+ \subseteq y \in \mathcal{C}(\sigma) \text{ complete}\},$$

with all other components directly inherited from σ .

Then $\text{comp}(\sigma) \equiv \sigma : A$ is innocent, well bracketed, causally well-bracketed²⁶.

Proof. Most conditions are immediate consequences of those from σ . The only non-trivial property is that $\mathcal{S}(\sigma)$ restricted to $|\text{comp}(\sigma)|$ is still an isomorphism family.

First, we prove that for any $\theta : x_1 \cong_\sigma x_2$, $x_1 \in \mathcal{C}(\text{comp}(\sigma))$ iff $x_2 \in \mathcal{C}(\text{comp}(\sigma))$. Indeed, assume $x_1^+ \subseteq y_1 \in \mathcal{C}(\sigma)$ complete. Then, since θ is an order-iso that preserves polarities, by *restriction* it restricts to $\theta' : x_1^+ \cong_\sigma x_2^+$. Now, by *extension*, θ' extends to $\theta'' : y_1 \cong_\sigma y_2$ for some $x_2^+ \subseteq y_2$. But since θ'' preserves the Question/Answer labeling, $y_2 \in \mathcal{C}(\sigma)$ is complete; hence $x_2 \in \mathcal{C}(\text{comp}(\sigma))$ as required. From that, it is straightforward that $\mathcal{S}(\text{comp}(\sigma))$ comprising symmetries between configurations of $\text{comp}(\sigma)$ is an isomorphism family.

For causal well-bracketing, $\text{comp}(\sigma)$ satisfies *wb-threads* by Lemma 7.14. For *globular*, taking a diagram as in Definition 7.13, m^+ appears in a $+$ -covered configuration of $\text{comp}(\sigma)$; hence in a complete configuration of σ . Thus, the condition follows by Lemma 7.16.

²⁶A strategy may well be causally well-bracketed without being well-bracketed: an example of that is a strategy $\sigma : \mathbb{U} \vdash \mathbb{U}$ that simultaneously calls its argument (but does nothing with the result) and returns.

Finally, we must show $\sigma \equiv \text{comp}(\sigma)$ – in fact, we show both have the same complete configurations. For that, any complete $x \in \mathcal{C}(\sigma)$ must also be $+$ -covered: take $x \in \mathcal{C}(\sigma)$ complete. If x has a maximal negative event m^- , since x is complete, m is an answer. But $\text{just}(\text{just}(m))$ is a Question answered by some $a^{+,A}$ in x – but then $m \leq_{\sigma} a$ by Lemma 7.15, contradicting maximality of m . Using this we conclude: clearly, any complete configuration of $\text{comp}(\sigma)$ is a complete configuration of σ . Reciprocally, a complete $x \in \mathcal{C}(\sigma)$ is $+$ -covered, and thus also a (complete) configuration of $\text{comp}(\sigma)$. So, $f\sigma = f\text{comp}(\sigma)$. \square

By construction, we also have that $\text{comp}(\sigma)$ is **complete**, in the following sense:

Definition 7.18. Consider A an arena, and $\sigma : A$ a causal strategy.

We say that σ is **complete**, if for any $x \in \mathcal{C}^+(\sigma)$ there is $x \subseteq y \in \mathcal{C}(\sigma)$ complete.

7.3. Meager Form. As for sequential strategies, definability applies for *finite* strategies, defined through a notion of *meager form*. But to define meager forms we will first need to restrict to *concrete arenas* in the sense of Section 2.20, which we must first update.

7.3.1. Updating concrete arenas. We enrich and update Definition 2.20.

Definition 7.19. A **concrete arena** is $(A, A^0, \text{lbl}, \text{ind})$ with A an arena, A^0 meager arena,

$$\text{lbl} : |A| \rightarrow |A^0|, \quad \text{ind} : |A| \rightarrow \mathbb{N}$$

two functions, satisfying, additionally to the conditions of Definition 2.20, the conditions:

- jointly injective:* for $a_1, a_2 \in |A|$, if $\text{lbl}(a_1) = \text{lbl}(a_2)$, $\text{ind}(a_1) = \text{ind}(a_2)$, and $\mathbf{pred}(a_1) = \mathbf{pred}(a_2)$, then $a_1 = a_2$.
- \mathcal{Q} -wide:* for any $q_1^{\mathcal{Q}} \in |A|$ non-minimal, for any $n \in \mathbb{N}$, there is $q_2^{\mathcal{Q}} \in |A|$ such that $\text{lbl}(q_1) = \text{lbl}(q_2)$, $\mathbf{pred}(q_1) = \mathbf{pred}(q_2)$ and $\text{ind}(q_2) = n$.
- A -narrow:* for any $a^- \in |A|$ minimal or $a^A \in |A|$, $\text{ind}(a) = 0$,
- A -conflicting:* if $a_1, a_2 \in |A|$ are distinct, they are in minimal conflict iff they are both minimal with the same polarity and copy index, or they are both answers to the same question.
- $+$ -transparent:* for any $\theta : x \cong_A y$, then $\theta \in \mathcal{S}_+(A)$ iff for all $a^- \in x$, $\text{ind}(\theta(a)) = \text{ind}(a)$.
- $-$ -transparent:* for any $\theta : x \cong_A y$, then $\theta \in \mathcal{S}_-(A)$ iff for all $a^+ \in x$, $\text{ind}(\theta(a)) = \text{ind}(a)$.

We call $\text{ind}(a)$ the **copy index** of a , and $\mathbf{pred}(-)$ is the (unique) *immediate predecessor*.

In arenas for ground types, all moves have copy index 0. For $A \& B$, the copy index function is simply inherited. For $A \rightarrow B$, the copy index of an initial $(1, (i, a))$ in A is simply i – in all other cases it is inherited. It is direct that all requirements are met.

7.3.2. Meager innocent strategies. We now introduce the causal counterpart of the meager alternating innocent strategies of Section 3.3.1. Those are parallel innocent strategies in the sense of Section 5, but on a restricted arena authorizing only Player replications:

Proposition 7.20. Consider A a concrete arena. Then, setting events:

$$|A^+| = \{a' \in |A| \mid \forall a^- \leq_A a', \text{ind}(a) = 0\},$$

with other components inherited from A , yields an arena A^+ .

Proof. All verifications are straightforward. For symmetry, if $\theta : x \cong_A y$ with $x, y \in \mathcal{C}(A^+)$, then by Definition 7.19, $\theta \in \mathcal{S}_+(A)$. But likewise by Definition 7.19, for $\theta : x \cong_A^+ y$, $x \in \mathcal{C}(A^+)$ iff $y \in \mathcal{C}(A^+)$. Together those imply that the restriction of $\mathcal{S}(A)$ to A^+ satisfies *extension* – other axioms are easy. Finally, from Definition 7.19 again the polarized isomorphism families are $\mathcal{S}_+(A^+) = \mathcal{S}(A^+)$, and $\mathcal{S}_-(A^+)$ restricted to identities. \square

So A^+ is A where Opponent has only access to copy index 0. This lets us define:

Definition 7.21. Consider A a concrete arena.

A **meager causal (pre)strategy** on A is a causal (pre)strategy $\sigma : A^+$.

As intended, this eliminates the infinity originating from Opponent repetitions. As a side-effect, the isomorphism family of σ becomes trivial: as the only non-trivial symmetries of A^+ are positive, it follows by condition *thin* (see Lemma 3.28 of [CCW19]) that symmetries of σ are reduced to identities – so a meager strategy is really a plain event structure.

Any causal strategy $\sigma : A$ yields a meager strategy $\text{mf}(\sigma) : A^+$, simply by restriction:

Proposition 7.22. Consider A a concrete arena, and $\sigma : A$ any causal strategy. Setting

$$|\text{mf}(\sigma)| = \{m' \in |\sigma| \mid \forall m^- \leq_\sigma m', \text{ind}(\partial_\sigma(m)) = 0\}$$

with other components inherited from σ , yields a meager causal strategy $\text{mf}(\sigma) : A^+$.

Proof. All conditions are straightforward verifications. \square

We call $\text{mf}(\sigma) : A^+$ the **meager form** of $\sigma : A$. As A^+ is closed under positive symmetry, it is immediate that $\text{mf}(-)$ preserves positive isomorphism: if $\sigma \approx \sigma'$, then $\text{mf}(\sigma) \approx \text{mf}(\sigma')$. As illustration, we show in Figure 59 the meager form of $\mathbf{plet}_{\mathbb{U}, \mathbb{U}}$ as defined in Section 4.4.2. Unlike Figure 38, this is now not merely a symbolic representation, but an exhaustive display of the full event structure of the meager form – we shall see that as for alternating innocent strategies, meager forms of parallel innocent causally well-bracketed strategies provide a way to give complete formal descriptions of the full infinite strategy.

Without parallel innocence, taking the meager form is a lossy operation. In Figure 60 we show a typical augmentation of the causal strategy obtained as the interpretation of

$$\vdash \mathbf{newref} \ x := 0 \ \mathbf{in} \ \lambda f^{\mathbb{U} \rightarrow \mathbb{U}}. f \ (\mathbf{let} \ v = !x \ \mathbf{in} \ x := 1; \ \mathbf{assume} \ (v =_{\mathbb{N}} 1)) : (\mathbb{U} \rightarrow \mathbb{U}) \rightarrow \mathbb{U},$$

displaying behaviour that is lost when taking the meager form: any interference between different copies of the same branch – a behaviour typically banned by parallel innocence.

This lets us define finiteness for parallel innocent strategies as in Section 3.3.1:

Definition 7.23. Consider A a concrete arena, and $\sigma : A$ a parallel innocent strategy.

We say that σ is **finite** if the set $|\text{mf}(\sigma)|_+ = \{m \in |\text{mf}(\sigma)| \mid \text{pol}(m) = +\}$ is finite.


If σ is finite, its **size** is the cardinal of $|\text{mf}(\sigma)|_+$.

7.3.3. Finite tests suffice. The key mechanism behind the reduction to finite tests is to be able to restrict a parallel innocent strategy following a finite subset of its meager form.

Say $x \in \mathcal{C}(\sigma)$ is **normal** iff for all $m^+ \rightarrow_\sigma m_1^-$ and $m^+ \rightarrow_\sigma m_2^-$ in x , $m_1 = m_2$. We show that every normal $x \in \mathcal{C}(\sigma)$ has a unique representative in $\text{mf}(\sigma)$.

Lemma 7.24. Consider A a concrete arena, $\sigma : A$ a causal strategy, and $x \in \mathcal{C}(\sigma)$ normal.

There is a unique $\text{mf}(x) \in \mathcal{C}(\text{mf}(\sigma))$ s.t. $x \cong_\sigma \text{mf}(x)$, and $\theta_x : x \cong_\sigma \text{mf}(x)$ is unique.

$$!((\mathbb{U} \& \mathbb{U}) \& ((\mathbb{U} \& \mathbb{U}) \rightarrow \mathbb{U})) \vdash \mathbb{U}$$
Figure 59: The meager form of $\text{plet}_{\mathbb{U},\mathbb{U}}$


$$(\mathbb{U} \rightarrow \mathbb{U}) \rightarrow \mathbb{U}$$


Figure 60: Lossy meager form

Proof. Existence. By induction on x . Consider $x \xrightarrow{m} \subset y$. If m is positive, by *extension* there is an extension of θ_x with (m, n) . As n is positive its negative dependencies are in $\text{mf}(x)$ so their display have copy index 0, so $n \in |\text{mf}(\sigma)|$. If m is negative, by *extension* on A , there is

$$\partial_\sigma \theta_x \cup \{(\partial_\sigma(m), a)\} : \partial_\sigma(x) \cup \{\partial_\sigma(m)\} \cong_A \partial_\sigma(\text{mf}(x)) \cup \{a\},$$

with a characterised by $\text{ind}(a)$, $\text{lbl}(a)$, and $\text{pred}(a)$. If m is an answer, so is a and by \mathcal{A} -*narrow*, $\text{ind}(a) = 0$. If m is a Question, we may not have $\text{ind}(a) = 0$. But then by \mathcal{Q} -*wide*, there is $a' \in |A|$ s.t. $\text{pred}(a') = \text{pred}(a)$, $\text{lbl}(a') = \text{lbl}(a)$ and $\text{ind}(a') = 0$; but we must prove that a' is not already in $\partial_\sigma(\text{mf}(x))$. If it is, there is $\theta_x(m') \in \text{mf}(x)$ s.t. $\partial_\sigma(\theta_x(m')) = a'$. By courtesy m' and m are negative in y with the same predecessor, contradicting normality of y . So y extends with a' with copy index 0. By *transparent*, $\partial_\sigma(\theta_x) \cup \{(\partial_\sigma(m), a')\} \in \mathcal{S}(A)$. So, by \sim -receptivity of σ , θ_x extends with $(m, \text{mf}(m))$ in σ s.t. $\text{mf}(m) \in |\text{mf}(\sigma)|$ as required.

Uniqueness. For $y_1, y_2 \in \mathcal{C}(\text{mf}(\sigma))$ s.t. $\theta_1 : x \cong_\sigma y_1$ and $\theta_2 : x \cong_\sigma y_2$, then $\theta = \theta_2 \circ \theta_1^{-1} : y_1 \cong_\sigma y_2$. But copy indices of Opponent events in $\partial_\sigma y_1$ and $\partial_\sigma y_2$ are 0, so by $+$ -*transparent*, $\partial_\sigma \theta \in \mathcal{S}_+(A)$. By Lemma 3.28 of [CCW19], $y_1 = y_2$ and $\theta = \text{id}$, so $\theta_1 = \theta_2$. \square

This does not depend on parallel innocence – which comes in when transporting *events*:

Lemma 7.25. *Consider A a concrete arena, $\sigma : A$ parallel innocent, and $m \in |\sigma|$.*

Then, there exists a unique $\text{mf}(m) \in |\text{mf}(\sigma)|$ such that $[m]_\sigma \cong_\sigma [\text{mf}(m)]_\sigma$.

Proof. *Pre-innocence* exactly states that the prime configuration $[m]_\sigma$ is normal. Hence, by Lemma 7.24, there is a unique $y \in \mathcal{C}(\text{mf}(\sigma))$ such that $[m]_\sigma \cong_\sigma y$. But then, as symmetries are order-isomorphisms, y is a prime configuration $y = [\text{mf}(m)]_\sigma$ as required. \square

Moreover, this assignment is preserved under symmetry:

Lemma 7.26. *Consider A a concrete arena, $\sigma : A$ parallel innocent.*

For any $\theta : x \cong_\sigma y$ and $m \in x$, $\text{mf}(m) = \text{mf}(\theta(m))$.

Proof. As θ is an order-iso, it restricts to $[m]_\sigma \cong_\sigma [\theta(m)]_\sigma$. Composition with $[m]_\sigma \cong [\text{mf}(m)]_\sigma$ and $[\theta(m)]_\sigma \cong [\text{mf}(\theta(m))]_\sigma$ yields $\varphi : [\text{mf}(m)]_\sigma \cong_\sigma [\text{mf}(\theta(m))]_\sigma$, and $\partial_\sigma(\varphi) \in \mathcal{S}_+(A)$ by $+$ -*transparent*. Hence, by Lemma 3.28 of [CCW19], φ is an identity. \square

From that, we may deduce the following:

Corollary 7.27. *For A concrete, $\sigma : A$ parallel innocent and $\tau \leq \text{mf}(\sigma)$ finite,*

$$|\sigma \upharpoonright \tau| = \{m \in |\sigma| \mid \text{mf}(m) \in |\tau|\}$$

with all other components inherited from σ yields a finite innocent $\sigma \upharpoonright \tau : A$ s.t. $\sigma \upharpoonright \tau \leq \sigma$. Moreover, if σ is well-bracketed (resp. causally well-bracketed), so is $\sigma \upharpoonright \tau$.

Proof. The only non-trivial condition, *extension* for $\mathcal{S}(\sigma \upharpoonright \tau)$, follows from Lemma 7.26. \square

From that, we may finally prove that finite tests suffice.

Corollary 7.28. *Consider A a concrete $--$ arena, and $\sigma_1, \sigma_2 : A$ parallel innocent.*

If there is $\alpha : !A \vdash \mathbb{U}$ parallel innocent and well-bracketed, such that

$$\alpha \odot_! \sigma_1 \Downarrow, \quad \alpha \odot_! \sigma_2 \Uparrow,$$

then there is $\alpha' \leq \alpha$ parallel innocent, well-bracketed, and finite, s.t. $\alpha' \odot_! \sigma_1$ and $\alpha' \odot_! \sigma_2$.

Proof. Since $\alpha \odot_! \sigma_1 \Downarrow$, there is $x^\alpha \odot x^{\sigma_1} \in \mathcal{C}(\alpha \odot \sigma_1^\dagger)$ s.t. $\partial_\alpha x^\alpha = x_A \parallel \{\mathbf{q}^-, \checkmark^+\}$. In particular, x^α is finite. So the set $X = \{\text{mf}(m) \mid m \in x^\alpha\}$ is finite, so there is $\tau \leq \text{mf}(\alpha)$ a finite meager strategy s.t. $X \subseteq |\tau|$. By Corollary 7.27, $\alpha' = \alpha \upharpoonright \tau : A$ is a finite parallel innocent strategy s.t. $\alpha' \leq \alpha$. Moreover, by construction, $x^\alpha \in \mathcal{C}(\alpha')$ with the same causal ordering as in α , so that $\alpha' \odot_! \sigma_1 \Downarrow$ still. Finally, $\alpha' \odot_! \sigma_2 \Downarrow$ would contradict $\alpha \odot_! \sigma_2 \Uparrow$ since $\alpha' \leq \alpha$. As $\mathcal{C}\text{-Plays}(\alpha') \subseteq \mathcal{C}\text{-Plays}(\alpha)$, α' is still well-bracketed. \square

In fact, meager forms of parallel innocent strategies can be expanded back to the original strategy. This is not used in the technical development, but we include it as Appendix D.5.

7.4. Factorization. We focus on finite tests. Unlike in the sequential argument of Section 3.3.1, parallel innocent strategies have no “first Player move” to reproduce first syntactically. Hence we organize our definability process differently. Its core is a *factorization result* (Corollary 7.43): namely, that every finite test $\alpha : !(\&A_i) \vdash \mathbb{X}$ may be obtained as

$$\alpha \equiv \mathbf{fo}(\alpha) \odot_! \langle x_i \alpha_{k,1} \dots \alpha_{k,p_i} \mid i \in I, k \in K_i \rangle, \quad (7.1)$$

with $\mathbf{fo}(\alpha)$ a strategy on a *first-order type* and $\alpha_{k,j}$ strictly smaller. This reduces finite definability to that for finite *first-order* strategies, dealt with in Section 7.5.1.

We first extract the components mentioned in (7.1): the *first-order substrategy* $\mathbf{fo}(\alpha)$, and the *argument substrategies* $\alpha_{k,j}$. We use as illustration the strategy with typical maximal augmentations in Figure 61. The *first-order sub-strategy*, in red, has events those depending on no Opponent question besides the initial move: it is independent of Opponent’s exploration of the arguments, and is purely first-order. The Player questions in this first-order part play a special role; we call them *primary questions*. Intuitively, they correspond to occurrences of variables not appearing in an argument to a variable call. In Figure 61, the primary questions are \mathbf{q}_0^+ , \mathbf{q}_1^+ and \mathbf{q}_2^+ . Depending on their type, the primary questions admit arguments that Opponent can access by playing questions pointing to them. Parts of the strategy accessed in this way are the *argument sub-strategies* – in Figure 61 there are three, respectively prompted by (*i.e.* causally depending on) $\mathbf{q}_{0,i}^-$, $\mathbf{q}_{2,k}^-$ and $\mathbf{q}_{1,r}^-$ and colored accordingly.

The strategy of Figure 61 is exactly definable; as illustration we show the term in Figure 62, with subterms colored so as to match the four components of the strategy²⁷.

²⁷In the end, our definability process will not quite give the term of Figure 62 for the strategy of Figure 61, but a sequential version as we do not know how to define first-order strategies in general – see Section 7.5.1.




Figure 61: A parallel innocent causally well-bracketed strategy

$$f : \mathbb{B} \rightarrow \mathbb{U}, g : (\mathbb{U} \rightarrow \mathbb{B}) \rightarrow \mathbb{B} \vdash$$

$$\text{let} \left(\begin{array}{l} x = f \mathbf{tt} \\ y = g(\lambda z^{\mathbb{U}}. \text{let} \left(\begin{array}{l} u = z \\ v = g \perp \end{array} \right) \text{in } v) \end{array} \right) \text{in} (\text{if } y \text{ then } (f \mathbf{ff}; \mathbf{tt}) \text{ else } \mathbf{ff}) : \mathbb{B}$$

Figure 62: Factorization and definability for Figure 61

The proof of factorization is organized as follows. In Section 7.4.1, we extract the first-order part, and in Section 7.4.2 we reorganize it so as to be able to re-compose it better. In Section 7.4.3, we extract the argument substrategies. In Section 7.4.4, we conclude.

7.4.1. *Shallow substrategy.* Consider A a type, and $\alpha : \llbracket A \rrbracket$ well-bracketed, parallel innocent, causally well-bracketed, finite, and complete. We call such a strategy a **test strategy**.

Necessarily A has the form $A_1 \rightarrow \dots \rightarrow A_n \rightarrow \mathbb{X}$ where $A_i = A_{i,1} \rightarrow \dots \rightarrow A_{i,p_i} \rightarrow \mathbb{X}_i$. Recall that \mathbb{X}, \mathbb{X}_i range over ground types, *i.e.* \mathbb{U}, \mathbb{B} and \mathbb{N} . Up to currying, we write

$$\alpha : !(\&_{1 \leq i \leq n} A_i) \vdash \mathbb{X},$$

omitting semantic brackets. We often shorten the left hand side part to $!(\&A_i)$, and reuse A for the arena $!(\&A_i) \vdash \mathbb{X}$. Now, we start with the *shallow substrategy*.

Proposition 7.29. *We define the **shallow substrategy** as having set of events*

$$|\text{sh}(\alpha)| = \{m \in |\alpha| \mid [m]_{\alpha} \text{ has at most one Opponent question}\},$$

and other components inherited from α . Then, $\text{sh}(\alpha) : !(\&\mathbb{X}_i) \vdash \mathbb{X}$ is a test strategy.

Proof. Write $\text{sh}(A)$ for the arena $!(\&\mathbb{X}_i) \vdash \mathbb{X}$. First, for each $m \in |\text{sh}(\alpha)|$, $\partial_{\alpha}(m) \in |\text{sh}(A)|$: indeed, the least events in $|A|$ but not in $|\text{sh}(A)|$ are Opponent questions. For *extension*, as symmetries are order-isos preserving polarities and Q/A labeling, they preserve $\text{sh}(\alpha)$. The conditions for a test strategy are direct by restriction from α .

We show $\text{sh}(\alpha)$ complete. Consider $x \in \mathcal{C}^+(\text{sh}(\alpha))$. Since α is complete, there is $x \subseteq y \in \mathcal{C}(\alpha)$ complete. In particular, there is an answer a^+ to the initial q_0^- . By Lemma 7.15, all moves of x are below a^+ for \leq_{σ} . So setting $z = [a]_{\sigma}$, $x \subseteq z$. By Lemma 7.14, all gccs leading to a are well-bracketed, so $[a]_{\sigma}$ is complete. Finally, $[a]_{\sigma} \in \mathcal{C}(\text{sh}(\alpha))$: indeed,




Figure 63: Augmentations of the flow substrategy for Figure 61

if it has a non-initial Opponent question q^- , it has an answer b^+ distinct from a^+ . But Player answers are maximal in gccs, contradicting that any gcc can be extended to a . \square

This captures the red part of Figure 61. All events of $\text{sh}(\alpha)$ are in $\text{mf}(\alpha)$: the negative dependencies of $m \in |\text{sh}(\alpha)|$ are either answers or the initial question, so by \mathcal{A} -narrow their display has copy index 0. Since α is finite, so is the set of positive events of $\text{sh}(\alpha)$.

7.4.2. The flow substrategy. We must reconstruct α using the categorical structure of $\rightarrow\text{-Strat}$. But as distinct Player questions in the same \mathbb{X}_i may receive distinct argument substrategies, we need to relabel $\text{sh}(\alpha)$ to send those to distinct components. First we define:

Definition 7.30. A **primary question** of α is any $q^{\mathcal{Q},+} \in |\text{sh}(\alpha)|$. We write \mathcal{Q} for the set of primary questions, and \mathcal{Q}_i for the primary questions displaying to \mathbb{X}_i .

By construction, $\mathcal{Q} = \uplus_{1 \leq i \leq n} \mathcal{Q}_i$. As $|\text{sh}(\alpha)|$ is finite, so is \mathcal{Q} , which allows us to set:

Definition 7.31. The **flow substrategy** $\text{flow}(\alpha) : \bigotimes_{1 \leq i \leq n} \bigotimes_{q \in \mathcal{Q}_i} \mathbb{X}_i \vdash \mathbb{X}$ is $\text{sh}(\alpha)$ with

$$\begin{aligned} \partial_{\text{flow}(\alpha)}(m^{\mathcal{Q},-}) &= (2, a) & \text{if } \partial_{\text{sh}(\alpha)}(m^{\mathcal{Q},-}) &= (2, a) \\ \partial_{\text{flow}(\alpha)}(m^{\mathcal{A},+}) &= (2, a) & \text{if } \partial_{\text{sh}(\alpha)}(m^{\mathcal{A},+}) &= (2, a) \\ \partial_{\text{flow}(\alpha)}(m^{\mathcal{Q},+}) &= (1, (i, (m, a))) & \text{if } \partial_{\text{sh}(\alpha)}(m^{\mathcal{Q},+}) &= (1, (j, (i, a))) \\ \partial_{\text{flow}(\alpha)}(m^{\mathcal{A},-}) &= (1, (i, (\text{just}(m), a))) & \text{if } \partial_{\text{sh}(\alpha)}(m^{\mathcal{A},-}) &= (1, (j, (i, a))), \end{aligned}$$

i.e. sending each $q \in \mathcal{Q}$ and its answers to the copy of \mathbb{X}_i specified by indices i, q .

It is a test strategy. We show in Figure 63 the maximal augmentations of the flow substrategy for Figure 61, tagging each component by the corresponding primary question.

7.4.3. The argument substrategies. Next, we focus on the *higher-order* structure, aiming to extract the arguments to (the variable calls corresponding to) the primary questions.

Fix a primary question $q \in \mathcal{Q}_i$. It displays to an initial event in A_i , which is:

$$!A_{i,1} \multimap \dots \multimap !A_{i,p_i} \multimap \mathbb{X}_i.$$

Argument substrategies are accessed by Opponent questions pointing to primary questions. Up to symmetry, there are p_i Opponent questions pointing to q , matching the p_i arguments of A_i . From now on, if $q \in \mathcal{Q}_i$ is a primary question and $q \rightarrow_{\sigma} m^{\mathcal{Q},-}$ an Opponent

question, we shall say that m is **in component** j if it displays to an initial move of $!A_{i,j}$. For $q \in \mathcal{Q}_i$ and $1 \leq j \leq p_i$, we shall extract the *argument sub-strategy* $\alpha_{q,j}$ initiated by Opponent questions pointing to q in component j . As the strategy provides the information for an argument of A_i it must live in $A_{i,j}$; but it can still access the context, so we aim for:

$$\alpha_{q,j} : !(\&A_i) \vdash !A_{i,j}.$$

To do this, we assign to all events of α tags, as follows:

Definition 7.32. Let $q \in \mathcal{Q}_i$ a primary question, $1 \leq j \leq p_i$, and $m \in |\alpha|$. We write:

$$\begin{aligned} m \in \text{sh}(\alpha) &\Leftrightarrow [m]_\alpha \text{ comprises exactly one Opponent question,} \\ m \in \alpha_{q,j} &\Leftrightarrow \text{there is } q \rightarrow_\alpha n^{\mathcal{Q},-} \text{ in component } j, \text{ such that } n \leq_\alpha m. \end{aligned}$$


Any $m \in |\alpha|$ receives a tag as either in $\text{sh}(\alpha)$; or in one of the argument sub-strategies. Crucially, each event receives *exactly one tag*. This is where all our structural constraints on strategies strike in, finally banning the pathological phenomenon of Figure 58.

Lemma 7.33. *Every event $m \in |\alpha|$ receives exactly one tag following Definition 7.32.*

Proof. First, each $m \in |\alpha|$ receives *at least* one tag. Any $[m]_\alpha$ contains at least one Opponent question: the initial move. If it contains exactly one Opponent question, $m \in \text{sh}(\alpha)$. Assume there are at least two. Take $n \leq_\alpha m$ minimal s.t. it is a non-initial Opponent question. Then its immediate predecessor is some $q \in \mathcal{Q}_i$; and so there is $1 \leq j \leq p_i$ s.t. $m \in \alpha_{q,j}$.

We prove that m receives *at most* one tag. Clearly if $m \in \alpha_{q,j}$ for some $q \in \mathcal{Q}_i$ and $1 \leq j \leq p_i$, there are at least two Opponent questions in $[m]_\alpha$ so we cannot have $m \in \text{sh}(\alpha)$. Assume $m \in \alpha_{q,j}$ and $m \in \alpha_{q',j'}$ for $q \in \mathcal{Q}_i, q' \in \mathcal{Q}_{i'}, 1 \leq j \leq p_i$ and $1 \leq j' \leq p_{i'}$. We first show that $q = q'$; seeking a contradiction assume they are distinct. But q and q' cannot be comparable: if $q \leq_\alpha q'$, $[q']_\alpha$ has at least two Opponent questions, contradicting $q' \in |\text{sh}(\alpha)|$.

Take $\rho \rightarrow m, \rho' \rightarrow m \in \text{gcc}(\alpha)$, respectively passing through q and q' . Diagrammatically:



and since q, q' are distinct, the diagram may be chosen with $X = \{m_1, \dots, m_k\}$ and $Y = \{n_1, \dots, n_p\}$ disjoint. By parallel innocence m_1 and n_1 are positive. By Lemma A.6 so must be m_{k+1} , so m_k and n_p are negative. We have a *globule* as in Definition 7.13, so X and Y are complete, and in particular q is answered in ρ . Writing $q = m_i$, m_{i+1} must answer q . But since $m \in \alpha_{q,j}$, $q \rightarrow_\alpha n^{\mathcal{Q},-} \leq_\alpha m$ where n is a negative question in component j . Then necessarily, $n = m_{i+1}$, or we get a contradiction with parallel innocence. Thus m_{i+1} is both a question and an answer, contradiction. So, $q = q'$. Finally, from Lemma 7.12, $j = j'$. \square

This shows any $m \in |\alpha|$ can always be attributed to exactly *one* of the sub-strategies we wish to extract. Accordingly, the *argument sub-strategies* will be defined with events

$$|\alpha_{q,j}| = \{m \in |\alpha| \mid m \in \alpha_{q,j}\},$$

completed to ess by inheriting the components from α , as will be made explicit later.

The display map requires a careful reindexing of events ending up on the right hand side, illustrated in Figures 64 and 65. For this we split $|\alpha_{q,j}|$ in two subsets: on the one hand, we have those events that depend *statically*, i.e. with respect to \leq_A (through ∂_α) on the primary question q – in Figure 64, those are $\mathbf{q}_{i,r}^-, \mathbf{q}_{1,r,0}^+, \check{\nu}_{1,r,0}^-$ and $b_{1,r}^+$. On the other hand,



$!(\mathbb{B} \rightarrow \mathbb{U}) \& ((\mathbb{U} \rightarrow \mathbb{B}) \rightarrow \mathbb{B}) \vdash \mathbb{B}$
 $!(\mathbb{B} \rightarrow \mathbb{U}) \& ((\mathbb{U} \rightarrow \mathbb{B}) \rightarrow \mathbb{B}) \vdash !(\mathbb{U} \rightarrow \mathbb{B})$

 Figure 64: Events of $|\alpha_{q_1^+, 1}|$ of Figure 61

 Figure 65: The corresponding aug. of $\alpha_{q_1^+, 1}$

the remaining events must follow from new calls to variables in the context – in Figure 64, those are q_{2r+3}^+ and b_{2r+3}^- . These two subsets are treated differently when defining the new display map: the former are left unchanged, while the latter are reindexed as in Figure 65.

We introduce notations for the canonical embeddings of the set of moves $|A_i|$ and $|A_{i,j}|$ into $|A|$. More precisely, it will be convenient, for each primary question $q \in \mathcal{Q}_i$, to write

$$\text{inj}_q(-) : |A_i| \rightarrow |A|$$

the injection adding the sequence of tags addressing A_i within A , originating from the tagged disjoint unions involved in all arena constructions – in particular, it maps the initial move of A_i to $\partial_\alpha(q)$. Likewise, $\text{inj}_{q,j} : |A_{i,j}| \rightarrow |A|$ addresses the j -th argument of q . Then:

Definition 7.34. We define a display map for $\alpha_{q,j}$ by setting, for $m \in |\alpha_{q,j}|$:

$$\begin{aligned} \partial_{\alpha_{q,j}}(m) &= \text{inj}_r(a) \quad \text{if} \quad \partial_\alpha(m) = \text{inj}_{q,j}(a), \\ \partial_{\alpha_{q,j}}(m) &= \partial_\alpha(m) \quad \text{otherwise,} \end{aligned}$$

where $\text{inj}_l(a) = (1, a)$ and $\text{inj}_r(a) = (2, a)$.

Altogether, this lets us extract $\alpha_{q,j}$ as intended:

Proposition 7.35. Consider $q \in \mathcal{Q}_i$ and $1 \leq j \leq p_i$. The *argument substrategy* for q, j is $(|\alpha_{q,j}|, \leq_{q,j}, \#_{q,j}, \mathcal{S}(\alpha_{q,j}), \partial_{\alpha_{q,j}})$, with components $\leq_{q,j}$ and $\#_{q,j}$ the restrictions of α , $\mathcal{S}(\alpha_{q,j}) = \{\theta \cap |\alpha_{q,j}|^2 \mid \theta \in \mathcal{S}(\alpha)\}$, and $\partial_{\alpha_{q,j}}$ in Definition 7.34.

Then, $\alpha_{q,j} : !(\&A_i) \vdash !A_{i,j}$ is well-bracketed, causally well-bracketed, parallel innocent.

Proof. A routine verification. The key point is that as symmetries of α are order-isomorphisms displayed to symmetries of A , it follows that they preserve the tag as in Definition 7.32. \square

Finally, we get rid of the $!$ on the right hand side, using *dereliction* $\text{der}_A : !A \vdash A$.

Proposition 7.36. Consider $q \in \mathcal{Q}_i$ and $1 \leq j \leq p_i$. Then, $\alpha_{q,j}^\bullet = \text{der}_{A_{i,j}} \odot \alpha_{q,j} : !(\&A_i) \vdash A_{i,j}$ is a test strategy with size strictly lesser than α .

Proof. Using Proposition 4.12, it is easy that $\alpha_{q,j}^\bullet$ is isomorphic, as an ess, to sub-ess of $\alpha_{q,j}$ where Opponent only opens the copy index 0 on the right hand side. There is at least one positive event in $\text{mf}(\alpha)$ but not of $\alpha_{q,j}^\bullet$, namely the primary question q . \square

$\frac{x_A \in \int A \quad x_B \in \int B}{x_A \otimes x_B \in \int(A \otimes B)}$	$\frac{x_A \in \int A \quad x_C \in \int C}{x_A \multimap x_C \in \int(A \multimap C)}$	$\frac{x \in \int A_i \quad (i \in I)}{(i, x) \in \int(\&_{i \in I} A_i)}$	$\frac{(x_C^i \in \int C)_{i \in I}}{[x_C^i \mid i \in I] \in \int C}$
---	---	--	--

Figure 66: Syntax for positions for A, B, C arenas with C strict

Note $\alpha_{q,j}$ is recovered from $\alpha_{q,j}^\bullet$, via the ‘‘Bang lemma’’ [AJM00] (see Appendix D.4):

Lemma 7.37. *For concrete arenas A, B with B pointed and $\sigma \in \multimap\text{-Strat}(!A, !B)$,*

$$(\text{der}_B \odot \sigma)^\dagger \approx \sigma.$$

Summing up, from the original strategy $\alpha : !(\&A_i) \vdash \mathbb{X}$, we have now constructed:

$$\begin{aligned} \text{flow}(\alpha) &: \otimes_{1 \leq i \leq n} \otimes_{q \in \mathcal{Q}_i} \mathbb{X}_i \vdash \mathbb{X} \\ \alpha_{q,j}^\bullet &: !(\&A_i) \vdash A_{i,j} \quad \text{for each } q \in \mathcal{Q}_i \text{ and } 1 \leq j \leq p_i. \end{aligned}$$

For each primary question $q \in \mathcal{Q}_i$, we use the cartesian closed internal language to form

$$x_1 : A_1, \dots, x_n : A_n \vdash x_i \alpha_{q,1}^\bullet \dots \alpha_{q,p_i}^\bullet : \mathbb{X}_i.$$

Let us write $\alpha_q : !(\&A_i) \vdash \mathbb{X}_i$ for the resulting strategy. Then, finally,

$$\text{recomp}(\alpha) = \text{flow}(\alpha) \odot (\otimes_{1 \leq i \leq n} \otimes_{q \in \mathcal{Q}_i} \alpha_q) \odot \delta_{\&A_i} : !(\&A_i) \vdash \mathbb{X}, \quad (7.2)$$

is our candidate to reconstruct α . Here, for B an arena and $n \in \mathbb{N}$, we write $\delta_B : !B \vdash (!B)^{\otimes n}$ for the obvious strategy (leaving n implicit). In the sequel we may only write δ .

7.4.4. *Positions of $\text{recomp}(\alpha)$.* We expect that $\text{recomp}(\alpha) \approx \alpha$, but we shall only prove $\text{recomp}(\alpha) \equiv \alpha$ – this is simpler as positions compose simply relationally.

To help reason on positions, we adopt a syntax presented in Figure 66, following the bijections of Lemma 7.11. We also write $x_A \vdash x_B \in \int(A \vdash B)$ for all $x_A \in \int A$ and $x_B \in \int B$.

We now analyse the positions of the recomposition (7.2). We start with:

Lemma 7.38. *For B strict, the positions $\int \delta_B$ of $\delta_B : !B \vdash (!B)^{\otimes n}$ are exactly those*

$$\left(\sum_{1 \leq i \leq n} x_i \vdash \otimes_{1 \leq i \leq n} x_i \right) \in \int (!B \vdash (!B)^{\otimes n})$$

where $x_i \in \int B$ for all $1 \leq i \leq n$.

By a direct variation of Lemma 4.18. Next we analyse the positions of α_q for $q \in \mathcal{Q}_i$:

Lemma 7.39. *For any $q \in \mathcal{Q}_i$, the non-empty positions $\int \alpha_q$ are exactly those of the form*

$$\left([(i, y_{q,1} \multimap \dots \multimap y_{q,p_i} \multimap v_q)] + \sum_{1 \leq j \leq p_i} x_{q,j} \right) \vdash v_q \in \int (!(\&A_i) \vdash \mathbb{X}_i)$$

where for each $1 \leq j \leq p_i$, $(x_{q,j} \vdash y_{q,j}) \in \int \alpha_{q,j}$, and for $v_q \in \int \mathbb{X}_i$.

Proof. From the laws of Seely categories, α_q is positively isomorphic to the composition

$$!(\&A_i) \xrightarrow{\delta} \otimes_{p_i+1} !(\&A_i) \xrightarrow{(\otimes_{1 \leq j \leq p_i} (\alpha_{q,j}^\bullet)^\dagger) \otimes [x_i]} (\otimes_{1 \leq j \leq p_i} !A_{i,j}) \otimes A_i \xrightarrow{\text{ev}} \mathbb{X}_i$$

in $\multimap\text{-Strat}$. The characterisation follows from Proposition 7.8, Lemma 7.38, Lemma 7.37, and a direct verification analogous to Lemma 7.38 for other copycat-like strategies involved. \square

From all those, we may characterise the positions of $\text{recomp}(\alpha)$ as

Corollary 7.40. *Non-empty positions of $\text{recomp}(\alpha)$ are exactly those of the form*

$$\sum_{1 \leq i \leq n} \sum_{q \in \mathcal{Q}'_i} \left([(i, y_{q,1} \multimap \dots \multimap y_{q,p_i} \multimap v_q)] + \sum_{1 \leq j \leq p_i} x_{q,j} \right) \vdash v \quad \in \quad \int (!(\&A_i) \vdash \mathbb{X})$$

where for all $1 \leq i \leq n$, \mathcal{Q}'_i is a subset of \mathcal{Q}_i , all v_q are non-empty, and:

$$\left(\otimes_{1 \leq i \leq n} \otimes_{q \in \mathcal{Q}'_i} v_q \vdash v \right) \in \mathfrak{flow}(\alpha), \quad ((x_{q,j} \vdash y_{q,j}) \in \mathfrak{f}\alpha_{q,j})_{q \in \mathcal{Q}'_i, 1 \leq j \leq p_i}.$$

Proof. Direct from Lemmas 7.38, 7.39 and Proposition 7.8. \square

7.4.5. *Positions of α .* We write $\mathcal{C}^c(\alpha)$ for the complete, non-empty configurations of α . If $x \in \mathcal{C}^c(\alpha)$, then $\text{sh}(x) = \{m \in x \mid m \in \text{sh}(\alpha)\} \in \mathcal{C}^c(\text{sh}(\alpha))$ and we write $\mathcal{Q}^x = \mathcal{Q} \cap x$, and \mathcal{Q}_i^x likewise. For each $q \in \mathcal{Q}_i^x$ and $1 \leq j \leq p_i$, we also have $x_{q,j} = \{m \in x \mid m \in \alpha_{q,j}\} \in \mathcal{C}(\alpha_{q,j})$.

From Lemma 7.33, this informs $x = \text{sh}(x) \uplus \left(\bigsqcup_{1 \leq i \leq n} \bigsqcup_{q \in \mathcal{Q}_i^x} \bigsqcup_{1 \leq j \leq p_i} x_{q,j} \right)$.

We show that all complete non-empty configurations of α arise in this way:

Lemma 7.41. *This yields a bijection between $\mathcal{C}^c(\alpha)$ and pairs $(x, (x_{q,j})_{1 \leq i \leq n, q \in \mathcal{Q}_i^x, 1 \leq j \leq p_i})$ where $x \in \mathcal{C}^c(\text{sh}(\alpha))$, $x_{q,j} \in \mathcal{C}(\alpha_{q,j})$ complete for all $1 \leq i \leq n$, $q \in \mathcal{Q}_i^x$ and $1 \leq j \leq p_i$.*

Moreover, writing $[x, (x_{q,j})_{i,q,j}] = x \uplus (\uplus_{i,q,j} x_{q,j}) \in \mathcal{C}(\alpha)$ this correspondence, we have

$$\partial_\alpha([x, (x_{q,j})_{i,q,j}]) = \left(\bigsqcup_{1 \leq i \leq n} \bigsqcup_{q \in \mathcal{Q}_i^x} \left([\text{inj}_q(z_{q,1} \multimap \dots \multimap z_{q,p_i} \multimap v_q)] \uplus \left[\bigsqcup_{1 \leq j \leq p_i} \text{inj}_j(y_{q,j}) \right] \right) \right) \uplus \text{inj}_r(v),$$

where we have, for all $1 \leq i \leq n$, $q \in \mathcal{Q}_i^x$ and $1 \leq j \leq p_i$:

$$\partial_{\text{flow}(\alpha)}(x) = \text{inj}_l \left(\otimes_{1 \leq i \leq n} \otimes_{q \in \mathcal{Q}_i^x} v_q \right) \uplus \text{inj}_r(v) \quad \partial_{\alpha_{q,j}}(x_{q,j}) = \text{inj}_l(y_{q,j}) \uplus \text{inj}_r(z_{q,j}).$$

Proof. From $x \in \mathcal{C}^c(\alpha)$, we get $(\text{sh}(x), (x_{q,j})_{i,q,j})$ from the decomposition above. Reciprocally, from $(x, (x_{q,j})_{i,q,j})$ we get a configuration in $\mathcal{C}^c(\alpha)$ as their union; it is down-closed by construction and consistent by determinism of α (as any immediate negative conflict originates from the arena, and hence appears in one of the components). Finally, the characterization of the display map follows from display maps of $\text{flow}(\alpha)$ and $\alpha_{q,j}$. \square

From this we may finally conclude the proof of factorization:

Corollary 7.42. *The strategies α and $\text{recomp}(\alpha)$ are positionally equivalent.*

Proof. Taking symmetry classes from Lemma 7.41, we obtain the same characterization of non-empty complete positions of α as in Corollary 7.40. \square

7.4.6. *Syntactic factorization.* Finally, we must reformulate the result relying on the cartesian closed structure only. The **first-order substrategy** $\mathfrak{fo}(\alpha) \in \multimap\text{-Strat}_!(\&_{1 \leq i \leq n} \&_{q \in \mathcal{Q}_i} \mathbb{X}_i, \mathbb{X})$ is obtained in the obvious way from $\text{flow}(\alpha)$ using the Seely category structure. Using Corollary 7.42, Proposition 7.36, and laws of a Seely category, we conclude:

Corollary 7.43. *Any test strategy $\alpha : !(\&A_i) \vdash \mathbb{X}$ factors as a composition of test strategies*

$$\alpha \equiv \mathfrak{fo}(\alpha) \odot ! \langle x_i \alpha_{q,1}^\bullet \dots \alpha_{q,p_i}^\bullet \mid 1 \leq i \leq n, q \in \mathcal{Q}_i \rangle,$$

where for all $1 \leq i \leq n$, $q \in \mathcal{Q}_i$ and $1 \leq j \leq p_i$, $\alpha_{q,j}^\bullet$ has size strictly lesser than α .




Figure 67: An undefinable first-order strategy

7.5. Finite Definability. Corollary 7.43 allows us to handle the higher-order structure, it only remains to prove definability for first-order test strategies.

7.5.1. First-order definability. Not every first-order test strategy is exactly definable in PCF_{\parallel} . For instance, that in Figure 67 is not series-parallel, while it is fairly easy to prove that all PCF_{\parallel} -definable terms on this type yield a series-parallel causal dependency. In general, we have not yet managed to properly understand which first-order strategies are definable. Luckily, we do not need to. Indeed, given a test α it is sufficient to find M such that $\llbracket M \rrbracket$ is *positionally equivalent* from α . In PCF_{\parallel} , without interference, the order of evaluation is unobservable; and positional equivalence is not sensitive to it. So our definability process will simply sequentialize α , while preserving its positions.

Consider $\Gamma = x_1 : \mathbb{X}_1, \dots, x_n : \mathbb{X}_n$, some ground type \mathbb{X} , and a test strategy:

$$\alpha : !(\&\mathbb{X}_i) \vdash \mathbb{X}.$$

If $\phi\alpha$ is empty, any diverging term M will satisfy $\phi\llbracket M \rrbracket = \emptyset$. Otherwise, there is some $x \in \mathcal{C}^c(\alpha)$. If α has no primary question, then $x = \{q_0^{\mathcal{Q},-}, a^{\mathcal{A},+}\}$, with a answering q_0 – write $\partial_\alpha(a) = v$ some answer in \mathbb{X} . But then by determinism of α , it cannot have any other move and $\alpha \approx \llbracket v \rrbracket$. Otherwise, if α has a primary question, it has one $a \in \mathcal{Q}$ which is **minimal**, *i.e.* it only depends on the initial move. But then q appears in every $x \in \mathcal{C}^c(\alpha)$:

Lemma 7.44. *For any minimal primary question q , for any $y \in \mathcal{C}^c(\alpha)$, we have $q \in y$.*

Proof. By determinism, $y \cup \{q_0, q\} \in \mathcal{C}(\alpha)$. Since y is complete, there is $a^{\mathcal{A},+} \in y$ such that a answers q_0 . But then, by Lemma 7.15 we have $q \leq_\alpha a$. It follows that $q \in y$ as required. \square

Choose $q \in \mathcal{Q}_i$ minimal. As q appears in all non-empty complete configurations, it is safe to first making a call to x_i , then branching on the possible return values. Since α is finite, there is a finite set V of values leading to an observable result. Now, for each $v \in V$, we define $\alpha_{(qv)}$ the *residual* of α after q yields value v ; and then proceed inductively. To define this residual, our first step is to rename α to isolate this first call:

Lemma 7.45. *For $q \in \mathcal{Q}_i$ minimal, there is a test strategy $\alpha_{(q)} : !((\&i\mathbb{X}_i) \ \& \ \mathbb{X}_i) \vdash \mathbb{X}$ s.t.*

- (1) for all $x \vdash w \in \phi\alpha_{(q)}$, then $x = x' + [(n+1, v)]$ such that $x' + [(i, v)] \vdash w \in \phi\alpha$,
- (2) for all $x \vdash w \in \phi\alpha$, then $x = x' + [(i, v)]$ such that $x' + [(n+1, v)] \vdash w \in \phi\alpha_{(q)}$.

Proof. The strategy $\alpha_{(q)}$ has components as for α except the display map, which sends q and its answers to the new component. The characterisation of positions is straightforward. \square

We obtain the residual $\alpha_{(qv)}$ as $\alpha_{(qv)} = \alpha_{(q)} \odot \langle \text{id}_{\&\mathbb{X}_i}, v \rangle : !(\&\mathbb{X}_i) \vdash \mathbb{X}$ writing $v : !(\&\mathbb{X}_i) \vdash \mathbb{X}_i$ the constant strategy. In order to characterize its positions, we note:

Lemma 7.46. *For any $v \in V$, the positions of $\langle \text{id}_{\&\mathbb{X}_i}, v \rangle^\dagger$ are exactly those of the form*

$$x \vdash x + p \cdot [(n+1, v)] \in f(!(\&\mathbb{X}_i) \vdash !(\&\mathbb{X}_i) \& \mathbb{X}_i),$$

where $p \cdot [(n+1, v)]$ denotes the p -fold sum, and for any $x \in f(!(\&\mathbb{X}_i))$ and $p \in \mathbb{N}$.

Proof. As $f(-)$ preserves the identity, $f(\text{id}_{\&\mathbb{X}_i}^\dagger)$ comprises exactly positions of the form $x \vdash x$, and likewise, $f v^\dagger$ comprises exactly positions of the form $p \cdot v$ for some $p \in \mathbb{N}$. The lemma then follows from Proposition 7.8 and by applying the laws of Seely categories for \rightarrow -Strat. \square

Using this lemma, we link the positions of α and $\alpha_{(qv)}$.

Lemma 7.47. *We have the following properties:*

- (1) *for any $x \vdash w \in \mathcal{F}\alpha$, there is $x = x' + [(i, v)]$ such that $x' \vdash w \in \mathcal{F}\alpha_{(qv)}$,*
- (2) *for any $x \vdash w \in \mathcal{F}\alpha_{(qv)}$, we have $x + [(i, v)] \vdash w \in \mathcal{F}\alpha$.*

Proof. By definition, we have $\alpha_{(qv)} = \alpha_{(q)} \odot \langle \text{id}_{\&\mathbb{X}_i}, v \rangle^\dagger$, so by Proposition 7.8,

$$f\alpha_{(qv)} = f(\alpha_{(q)}) \odot f(\langle \text{id}_{\&\mathbb{X}_i}, v \rangle^\dagger).$$

The lemma directly follows by Lemmas 7.45 and 7.46. \square

We have $\alpha_{(qv)}$ a test strategy with size strictly smaller than that of α . By IH there is

$$x_1 : \mathbb{X}_1, \dots, x_n : \mathbb{X}_n \vdash N_{(qv)} : \mathbb{X}_i,$$

for each $v \in V$, such that $\llbracket N_{(qv)} \rrbracket \equiv \alpha_{(qv)}$. Finally, we define $x_1 : \mathbb{X}_1, \dots, x_n : \mathbb{X}_n \vdash M : \mathbb{X}$ as

$$\begin{array}{ll} \text{case } x_i \text{ of} & \text{let } x = x_i \text{ in} \\ v_1 \mapsto N_{(qv_1)} & \text{if } x =_{\mathbb{X}} v_1 \text{ then } N_{(qv_1)} \\ v_2 \mapsto N_{(qv_2)} & \text{else if } x =_{\mathbb{X}} v_2 \text{ then } N_{(qv_2)} \\ \dots & \dots \\ v_p \mapsto N_{(qv_p)} & \text{else if } x =_{\mathbb{X}} v_p \text{ then } N_{(qv_p)} \\ & \text{else } \perp \end{array} \quad \stackrel{\text{def}}{=} \quad$$

where $V = \{v_1, \dots, v_p\}$, using the syntactic sugar introduced in Section 1.3. Write

$$x_1 : \mathbb{X}_1, \dots, x_n : \mathbb{X}_n, x : \mathbb{X}_i \vdash M' : \mathbb{X}$$

for the iterated if statement, *i.e.* M is **let** $x = x_i$ **in** M' . It remains to analyze the positions of $\llbracket M \rrbracket$ and $\llbracket M' \rrbracket$ to show that $\llbracket M \rrbracket \equiv \alpha$ as required. We start with the positions of $\llbracket M' \rrbracket$.

Lemma 7.48. *We have the following properties:*

- (1) *for any $x + p \cdot [(n+1, v)] \vdash w \in \mathcal{F}\llbracket M' \rrbracket$ where $x \in f!(\&_{1 \leq i \leq n} \mathbb{X}_i)$, $x \vdash w \in \mathcal{F}\llbracket N_{(qv)} \rrbracket$,*
- (2) *for any $x \vdash w \in \mathcal{F}\llbracket N_{(qv)} \rrbracket$, then there is $p \in \mathbb{N}$ s.t. $x + p \cdot [(n+1, v)] \vdash w \in \mathcal{F}\llbracket M' \rrbracket$.*

Proof. It is a direct verification, amounting to the correctness of our definition for equality test and the usual laws for conditionals, that for any $v \in V$ we have $\llbracket M' \rrbracket \odot \langle \text{id}, v \rangle = \llbracket N_{(qv)} \rrbracket$. The claim then follows by Proposition 7.8 and Lemma 7.46. \square

It remains to take the interpretation of the **let** construction into account. Recall that

$$\llbracket M \rrbracket = \text{let}_{\mathbb{X}_i, \mathbb{X}} \odot! \langle \pi_i, \Lambda^1(\llbracket M' \rrbracket) \rangle : !(\&\mathbb{X}_i) \vdash \mathbb{X}$$

where $\text{let}_{\mathbb{X}_i, \mathbb{X}} : !(\mathbb{X}_i \& (!\mathbb{X}_i \multimap \mathbb{X})) \vdash \mathbb{X}$. We characterize the positions of $\text{let}_{\mathbb{X}_i, \mathbb{X}}$ as follows.

Lemma 7.49. *The non-empty positions of $\text{let}_{\mathbb{X}_i, \mathbb{X}}$ are exactly those of the form*

$$[(1, \mathbf{v})] + [(2, ((p \cdot [\mathbf{v}]) \multimap \mathbf{w}))] \vdash \mathbf{w} \in f(!(\mathbb{X}_i \& (!\mathbb{X}_i \multimap \mathbb{X})) \vdash \mathbb{X})$$

for $\mathbf{v} \in \mathcal{F}\mathbb{X}_i$, $\mathbf{w} \in \mathcal{F}\mathbb{X}$, and $p \in \mathbb{N}$.

Proof. A direct analysis of positions reached by complete configurations of $\text{let}_{\mathbb{X}_i, \mathbb{X}}$. \square

We can now wrap up, showing that $\llbracket M \rrbracket$ has the same non-empty positions as α .

Lemma 7.50. *We have the following two properties:*

- (1) *for any $\mathbf{x} \vdash \mathbf{w} \in \mathcal{F}\llbracket M \rrbracket$, there is $\mathbf{x} = \mathbf{x}' + [(i, \mathbf{v})]$ such that $\mathbf{x}' \vdash \mathbf{w} \in \mathcal{F}\llbracket N_{(qv)} \rrbracket$,*
- (2) *for any $\mathbf{x} \vdash \mathbf{w} \in \mathcal{F}\llbracket N_{(qv)} \rrbracket$, we have $\mathbf{x} + [(i, \mathbf{v})] \vdash \mathbf{w} \in \mathcal{F}\llbracket M \rrbracket$.*

Proof. By Lemmas 7.48 and 7.49. \square

So we have $\llbracket M \rrbracket \equiv \alpha$ as desired. Summing up, we have proved:

Proposition 7.51. *For $\alpha : !(\&\mathbb{X}_i) \vdash \mathbb{X}$ any test strategy, there is*

$$x_1 : \mathbb{X}_1, \dots, x_n : \mathbb{X}_n \vdash M : \mathbb{X}$$

a term of PCF (not using parallel evaluation) such that $\llbracket M \rrbracket \equiv \alpha$.

7.5.2. *Finite definability.* We may now conclude the proof of finite definability.

Corollary 7.52. *Let $\Gamma \vdash A$ be a PCF typing judgment, and $\alpha : \llbracket \Gamma \rrbracket \vdash \llbracket A \rrbracket$ a test strategy.*

Then, there is $\Gamma \vdash M : A$ such that $\llbracket M \rrbracket \equiv \alpha$.

Proof. Up to currying, we write $\alpha : !(\&_{1 \leq i \leq n} A_i) \vdash \mathbb{X}$, writing $A_i = A_{i,1} \rightarrow \dots \rightarrow A_{i,p_i} \rightarrow \mathbb{X}_i$ for $1 \leq i \leq n$. We reason by induction on the size of α . By Corollary 7.43, α factors as

$$\alpha \equiv \mathbf{fo}(\alpha) \odot! \langle x_i \alpha_{q,1}^\bullet \dots \alpha_{q,p_i}^\bullet \mid 1 \leq i \leq n, q \in \mathcal{Q}_i \rangle,$$

with each \mathcal{Q}_i finite, and for $1 \leq i \leq n$, $q \in \mathcal{Q}_i$ and $1 \leq j \leq p_i$, $\alpha_{q,j}^\bullet : !(\&A_i) \vdash A_{i,j}$ a test strategy of size strictly smaller than α . By induction hypothesis, there is

$$x_1 : A_1, \dots, x_n : A_n \vdash N_{q,j} : A_{i,j}$$

such that $\llbracket N_{q,j} \rrbracket \equiv \alpha_{q,j}$. Let us write $\mathcal{Q}_i = \{q_{i,1}, \dots, q_{i,k_i}\}$. By Proposition 7.51 there is also

$$x_{q_{i,1}} : \mathbb{X}_1, \dots, x_{q_{i,k_i}} : \mathbb{X}_1, \dots, x_{q_{n,1}} : \mathbb{X}_n, \dots, x_{q_{n,k_n}} : \mathbb{X}_n \vdash M_{\mathbf{fo}} : \mathbb{X}$$

such that $\llbracket M_{\mathbf{fo}} \rrbracket \equiv \mathbf{fo}(\alpha)$. Then, we define the term $x_1 : A_1, \dots, x_n : A_n \vdash M : \mathbb{X}$ as

$$x_1 : A_1, \dots, x_n : A_n \vdash M_{\mathbf{fo}}[x_i N_{q_{i,l},1} \dots N_{q_{i,l},p_i}/x_{q_{i,l}}] : \mathbb{X}.$$

Then we may finally compute

$$\begin{aligned} \llbracket M \rrbracket &= \llbracket M_{\mathbf{fo}} \rrbracket \odot! \langle [x_i N_{q_{i,l},1} \dots N_{q_{i,l},p_i}] \mid 1 \leq i \leq n, q_{i,l} \in \mathcal{Q}_i \rangle \\ &\equiv \llbracket M_{\mathbf{fo}} \rrbracket \odot! \langle x_i \llbracket N_{q,1} \rrbracket \dots \llbracket N_{q,n} \rrbracket \mid 1 \leq i \leq n, q \in \mathcal{Q}_i \rangle \\ &\equiv \mathbf{fo}(\alpha) \odot! \langle x_i \alpha_{q,1}^\bullet \dots \alpha_{q,p_i}^\bullet \mid 1 \leq i \leq n, q \in \mathcal{Q}_i \rangle \end{aligned}$$

using the substitution lemma for cartesian closed categories, compatibility of interpretation with the internal language, the properties of $M_{\mathbf{fo}}$ and $N_{q,j}$ and that \equiv is a congruence. \square

7.6. **Full Abstraction for PCF_{\parallel} .** We may now prove our final full abstraction result.

Theorem 7.53. *The model $\rightarrow\text{-Strat}_!^{\text{wb,inn}}$ is intensionally fully abstract for PCF_{\parallel} .*

Proof. Consider $\vdash M, N : A$ such that $M \sim N$, and assume that $\llbracket M \rrbracket \uparrow \llbracket N \rrbracket$, i.e. there is a test $\alpha \in \rightarrow\text{-Strat}_!^{\text{wb,inn}}(\llbracket A \rrbracket, \llbracket \mathbb{U} \rrbracket)$ such that, w.l.o.g., $\alpha \odot! \llbracket M \rrbracket \Downarrow$ and $\alpha \odot! \llbracket N \rrbracket \Uparrow$.

By Corollary 7.28, we assume α is additionally finite. We consider $\text{comp}(\alpha)$ as defined in Proposition 7.17. By construction it is well-bracketed, parallel innocent, and finite. By Proposition 7.17 it is causally well-bracketed, so a *test strategy*. Proposition 7.17 also states that $\text{comp}(\alpha) \equiv \alpha$ which is a congruence by Corollary 7.9, so

$$\text{comp}(\alpha) \odot! \llbracket M \rrbracket \Downarrow, \quad \text{comp}(\alpha) \odot! \llbracket N \rrbracket \Uparrow.$$

By Corollary 7.52, there is a term $x : A \vdash T : \mathbb{U}$ such that $\llbracket T \rrbracket \equiv \text{comp}(\alpha)$. Defining the context $C[-] = (\lambda x^A. T) [-]$, it follows from the laws of cartesian closed categories that

$$\llbracket C[M] \rrbracket = \llbracket (\lambda x^A. T) M \rrbracket \approx \llbracket T[M/x] \rrbracket \approx \llbracket T \rrbracket \odot! \llbracket M \rrbracket \equiv \text{comp}(\alpha) \odot! \llbracket M \rrbracket \Downarrow,$$

and likewise, $\llbracket C[N] \rrbracket \equiv \text{comp}(\alpha) \odot! \llbracket N \rrbracket$. By Theorem 4.40, $C[M] \Downarrow$. By hypothesis $M \sim N$, so $C[N] \Downarrow$. By Theorem 4.40 again, $\llbracket C[N] \rrbracket \Downarrow$, hence $\text{comp}(\alpha) \odot! \llbracket N \rrbracket \Downarrow$, contradiction. \square

We may finally answer our main question positively, with the following theorem.

Theorem 7.54. *The model $\rightarrow\text{-Strat}_!^{\text{wb}}$ supports parallel innocence and sequentiality, s.t.*

$$\begin{aligned} \rightarrow\text{-Strat}_!^{\text{wb}} & \text{ is fully abstract for } \text{IA}_{\parallel}, \\ \rightarrow\text{-Strat}_!^{\text{wb,inn}} & \text{ is fully abstract for } \text{PCF}_{\parallel}, \\ \rightarrow\text{-Strat}_!^{\text{wb,seq}} & \text{ is fully abstract for } \text{IA}, \\ \rightarrow\text{-Strat}_!^{\text{wb,seq,inn}} & \text{ is fully abstract for } \text{PCF}, \end{aligned}$$

Thus parallel innocence exactly bans interference, and sequentiality exactly bans parallelism. Through this theorem, we have successfully disentangled parallelism and interference.

8. CONCLUSION

It is puzzling that disentangling parallelism and interference requires such an intricate machinery whereas the original *semantic cube* arose almost “by accident” from minor variations of the Hyland-Ong model of PCF.

Our interpretation is that computational effects may be organized in distinct categories. Some effects, such as interference and control, bring more freedom as to how execution and its control flow explores a piece of code. Once a sufficiently general mathematical description of the control flow is given (such as the original Hyland-Ong setting for sequential deterministic computation), this additional freedom may be studied and characterized. In contrast, other effects such as non-determinism and parallelism, affect the inherent structure of execution itself: non-determinism quite explicitly so by generating non-deterministic branching, and parallelism in a more DAG-like fashion – we refer to both under the umbrella name “branching effects”. What our paper demonstrates – along with earlier papers on non-deterministic innocence [CCW14, TO15] – is that if we aim to realize fully the “unified semantic landscape” of Abramsky’s programme, we must first deal with branching effects. Non-branching effects should follow by characterizing the causal patterns they allow²⁸.

²⁸The line is not always so clear between branching and non-branching effects: for instance, in a sequential setting interference is non-branching, but in a parallel setting it generates non-deterministic choice.

What other branching effects are around? One currently at the focus of the semantics community is probabilistic choice. By itself, the probabilistic branching structure is not much harder than non-deterministic choice [TO14, CCPW18]. However its interaction with non-deterministic and parallel branching is a significant challenge, and one of the remaining scientific and technical barriers for a truly unified game semantics landscape.

Acknowledgments. First we would like to thank Glynn Winskel. Though he has not participated in the developments in this paper per se, we rely deeply on his pioneering work - and our subsequent collaborations - on concurrent games. We owe a lot to his insight.

This work was supported by the ANR project DyVerSe (ANR-19-CE48-0010-01); and by Labex MiLyon (ANR-10-LABX-0070) of Université de Lyon, within the program “Investissements d’Avenir” (ANR-11-IDEX-0007), operated by the French ANR.

REFERENCES

- [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, 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. *Electron. Notes Theor. Comput. Sci.*, 3:2–14, 1996.
- [AM99a] Samson Abramsky and Guy McCusker. Game semantics. In *Computational logic*, 1999.
- [AM99b] Samson Abramsky and Paul-André Mellès. Concurrent games and full completeness. In *14th Annual IEEE Symposium on Logic in Computer Science, Trento, Italy, July 2-5, 1999*, pages 431–442. IEEE Computer Society, 1999.
- [Cas17] Simon Castellan. *Concurrent structures in game semantics. (Structures concurrentes en sémantique des jeux)*. PhD thesis, University of Lyon, France, 2017.
- [CC16] Simon Castellan and Pierre Clairambault. Causality vs. interleavings in concurrent game semantics. In José Desharnais and Radha Jagadeesan, editors, *27th International Conference on Concurrency Theory, CONCUR 2016, August 23-26, 2016, Québec City, Canada*, volume 59 of *LIPICs*, pages 32:1–32:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016.
- [CCHW18] Simon Castellan, Pierre Clairambault, Jonathan Hayman, and Glynn Winskel. Non-angelic concurrent game semantics. In Christel Baier and Ugo Dal Lago, editors, *Foundations of Software Science and Computation Structures - 21st International Conference, FOSSACS 2018, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2018, Thessaloniki, Greece, April 14-20, 2018, Proceedings*, volume 10803 of *Lecture Notes in Computer Science*, pages 3–19. Springer, 2018.
- [CCPW18] Simon Castellan, Pierre Clairambault, Hugo Paquet, and Glynn Winskel. The concurrent game semantics of probabilistic PCF. In *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2018, Oxford, UK, July 09-12, 2018*, pages 215–224, 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 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*, pages 28:1–28:10. ACM, 2014.
- [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*, pages 232–243. IEEE Computer Society, 2015.
- [CCW19] Simon Castellan, Pierre Clairambault, and Glynn Winskel. Thin games with symmetry and concurrent hyland-ong games. *Log. Methods Comput. Sci.*, 15(1), 2019.
- [CD15] Pierre Clairambault and Peter Dybjer. Game semantics and normalization by evaluation. 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 56–70. Springer, 2015.
- [CdV20] Pierre Clairambault and Marc de Visme. Full abstraction for the quantum lambda-calculus. *Proc. ACM Program. Lang.*, 4(POPL):63:1–63:28, 2020.
- [CH10] Pierre Clairambault and Russ Harmer. Totality in arena games. *Ann. Pure Appl. Logic*, 161(5):673–689, 2010.
- [Cur98] Pierre-Louis Curien. Abstract Böhm trees. *Mathematical Structures in Computer Science*, 8(6):559–591, 1998.
- [Cur06] Pierre-Louis Curien. Notes on game semantics. 2006. Unpublished.
- [CY19] Simon Castellan and Nobuko Yoshida. Two sides of the same coin: session types and game semantics: a synchronous side and an asynchronous side. *Proc. ACM Program. Lang.*, 3(POPL):27:1–27:29, 2019.
- [DH00] Vincent Danos and Russ Harmer. Probabilistic game semantics. In *15th Annual IEEE Symposium on Logic in Computer Science, Santa Barbara, California, USA, June 26-29, 2000*, pages 204–213. IEEE Computer Society, 2000.
- [DHR96] Vincent Danos, Hugo Herbelin, and Laurent Regnier. Game semantics & abstract machines. In *Proceedings, 11th Annual IEEE Symposium on Logic in Computer Science, New Brunswick, New Jersey, USA, July 27-30, 1996*, pages 394–405. IEEE Computer Society, 1996.
- [dV20] Marc de Visme. *Quantum Game Semantics*. PhD thesis, University of Lyon, France, 2020.
- [FP09] Claudia Faggian and Mauro Piccolo. Partial orders, event structures and linear strategies. In Pierre-Louis Curien, editor, *Typed Lambda Calculi and Applications, 9th International Conference, TLCA 2009, Brasilia, Brazil, July 1-3, 2009. Proceedings*, volume 5608 of *Lecture Notes in Computer Science*, pages 95–111. Springer, 2009.
- [GG12] Murdoch Gabbay and Dan R. Ghica. Game semantics in the nominal model. In Ulrich Berger and Michael W. Mislove, editors, *Proceedings of the 28th Conference on the Mathematical Foundations of Programming Semantics, MFPS 2012, Bath, UK, June 6-9, 2012*, volume 286 of *Electronic Notes in Theoretical Computer Science*, pages 173–189. Elsevier, 2012.
- [Gir87] Jean-Yves Girard. Linear logic. *Theor. Comput. Sci.*, 50:1–102, 1987.
- [GM03] Dan R. Ghica and Guy McCusker. The regular-language semantics of second-order idealized ALGOL. *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. Log.*, 151(2-3):89–114, 2008.
- [GMO06] Dan R. Ghica, Andrzej S. Murawski, and C.-H. Luke Ong. Syntactic control of concurrency. *Theor. Comput. Sci.*, 350(2-3):234–251, 2006.
- [Gou17] Jean Goubault-Larrecq. Isomorphism theorems between models of mixed choice. *Math. Struct. Comput. Sci.*, 27(6):1032–1067, 2017.
- [Gri90] Timothy Griffin. A formulae-as-types notion of control. In Frances E. Allen, editor, *Conference Record of the Seventeenth Annual ACM Symposium on Principles of Programming Languages, San Francisco, California, USA, January 1990*, pages 47–58. ACM Press, 1990.
- [GT12] Dan R. Ghica and Nikos Tzevelekos. A system-level game semantics. In Ulrich Berger and Michael W. Mislove, editors, *Proceedings of the 28th Conference on the Mathematical Foundations of Programming Semantics, MFPS 2012, Bath, UK, June 6-9, 2012*, volume 286 of *Electronic Notes in Theoretical Computer Science*, pages 191–211. Elsevier, 2012.
- [Har04] Russ Harmer. Innocent game semantics. *Lecture notes*, 2007, 2004.
- [HHM07] Russ 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, 2007.

- [HM99] Russ Harmer and Guy McCusker. A fully abstract game semantics for finite nondeterminism. In *14th Annual IEEE Symposium on Logic in Computer Science, Trento, Italy, July 2-5, 1999*, pages 422–430. IEEE Computer Society, 1999.
- [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.
- [Jab15] Guilhem Jaber. Operational nominal game semantics. 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 264–278. Springer, 2015.
- [Lai97] James Laird. Full abstraction for functional languages with control. In *LICS*, pages 58–67, 1997.
- [Lai01a] James Laird. A fully abstract game semantics of local exceptions. In *LICS*, 2001.
- [Lai01b] James Laird. A game semantics of idealized CSP. *Electr. Notes Theor. Comput. Sci.*, 45:232–257, 2001.
- [Lai04] James Laird. A calculus of coroutines. In Josep Díaz, Juhani Karhumäki, Arto Lepistö, and Donald Sannella, editors, *Automata, Languages and Programming: 31st International Colloquium, ICALP 2004, Turku, Finland, July 12-16, 2004. Proceedings*, volume 3142 of *Lecture Notes in Computer Science*, pages 882–893. Springer, 2004.
- [Ler09] Xavier Leroy. Formal verification of a realistic compiler. *Commun. ACM*, 52(7):107–115, 2009.
- [Lev14] Paul Blain Levy. Transition systems over games, 2014. Talk at the *Chocola* seminar.
- [Loa01] Ralph Loader. Finitary PCF is not decidable. *Theor. Comput. Sci.*, 266(1-2):341–364, 2001.
- [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 *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*, pages 64:1–64:10, 2014.
- [McC98] Guy McCusker. *Games and full abstraction for a functional metalanguage with recursive types*. CPH/BCS distinguished dissertations. Springer, 1998.
- [Mel03] Paul-André Melliès. Asynchronous games 1: A group-theoretic formulation of uniformity. *Manuscript, Available online*, 2003.
- [Mel04] Paul-André Melliès. Asynchronous games 2: The true concurrency of innocence. In Philippa Gardner and Nobuko Yoshida, editors, *CONCUR 2004 - Concurrency Theory, 15th International Conference, London, UK, August 31 - September 3, 2004, Proceedings*, volume 3170 of *Lecture Notes in Computer Science*, pages 448–465. Springer, 2004.
- [Mel05] Paul-André Melliès. Asynchronous games 4: A fully complete model of propositional linear logic. In *20th IEEE Symposium on Logic in Computer Science (LICS 2005), 26-29 June 2005, Chicago, IL, USA, Proceedings*, pages 386–395. IEEE Computer Society, 2005.
- [Mel09] Paul-André Melliès. Categorical semantics of linear logic. *Panoramas et Synthèses*, 2009.
- [Mel19] Paul-André Melliès. Categorical combinatorics of scheduling and synchronization in game semantics. *Proc. ACM Program. Lang.*, 3(POPL):23:1–23:30, 2019.
- [MM07] Paul-André Melliès and Samuel Mimram. Asynchronous games: Innocence without alternation. In Luís Caires and Vasco Thudichum Vasconcelos, editors, *CONCUR 2007 - Concurrency Theory, 18th International Conference, CONCUR 2007, Lisbon, Portugal, September 3-8, 2007, Proceedings*, volume 4703 of *Lecture Notes in Computer Science*, pages 395–411. Springer, 2007.
- [MT10] Paul-André Melliès and Nicolas Tabareau. Resource modalities in tensor logic. *Ann. Pure Appl. Log.*, 161(5):632–653, 2010.
- [MT13] Andrzej S. Murawski and Nikos Tzevelekos. Deconstructing general references via game semantics. In Frank Pfenning, editor, *Foundations of Software Science and Computation Structures - 16th International Conference, FOSSACS 2013, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2013, Rome, Italy, March 16-24, 2013. Proceedings*, volume 7794 of *Lecture Notes in Computer Science*, pages 241–256. Springer, 2013.
- [Mur03] Andrzej S. Murawski. On program equivalence in languages with ground-type references. In *18th IEEE Symposium on Logic in Computer Science (LICS 2003), 22-25 June 2003, Ottawa, Canada, Proceedings*, page 108. IEEE Computer Society, 2003.

- [Mur07] Andrzej S. Murawski. Bad variables under control. 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 558–572. Springer, 2007.
- [Mur10] Andrzej S. Murawski. Full abstraction without synchronization primitives. *Electr. Notes Theor. Comput. Sci.*, 265:423–436, 2010.
- [Ong02] C.-H. Luke Ong. Observational equivalence of 3rd-order idealized algol is decidable. In *17th IEEE Symposium on Logic in Computer Science (LICS 2002), 22-25 July 2002, Copenhagen, Denmark, Proceedings*, pages 245–256. IEEE Computer Society, 2002.
- [Paq20] Hugo Paquet. *Probabilistic concurrent game semantics*. PhD thesis, PhD thesis, University of Cambridge, 2020.
- [Plo77] Gordon D. Plotkin. LCF considered as a programming language. *Theor. Comput. Sci.*, 5(3):223–255, 1977.
- [Plo81] Gordon D. Plotkin. Post-graduate lecture notes in advanced domain theory (incorporating the "Pisa Notes"). *Dept. of Computer Science, Univ. of Edinburgh*, 1981.
- [RW11] Silvain Rideau and Glynn Winskel. Concurrent strategies. In *Proceedings of the 26th Annual IEEE Symposium on Logic in Computer Science, LICS 2011, June 21-24, 2011, Toronto, Ontario, Canada*, pages 409–418. IEEE Computer Society, 2011.
- [SBCA15] Gordon Stewart, Lennart Berlinger, Santiago Cuellar, and Andrew W. Appel. Compositional compcert. In Sriram K. Rajamani and David Walker, editors, *Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2015, Mumbai, India, January 15-17, 2015*, pages 275–287. ACM, 2015.
- [SYW⁺16] Sam Staton, Hongseok Yang, Frank D. Wood, Chris Heunen, and Ohad Kammar. Semantics for probabilistic programming: higher-order functions, continuous distributions, and soft constraints. In Martin Grohe, Eric Koskinen, and Natarajan Shankar, editors, *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science, LICS '16, New York, NY, USA, July 5-8, 2016*, pages 525–534. ACM, 2016.
- [TO14] Takeshi Tsukada and C.-H. Luke Ong. Innocent strategies are sheaves over plays - deterministic, non-deterministic and probabilistic innocence. *CoRR*, abs/1409.2764, 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*, pages 220–231. IEEE Computer Society, 2015.
- [VKS19] Matthijs Vákár, Ohad Kammar, and Sam Staton. A domain theory for statistical probabilistic programming. *Proc. ACM Program. Lang.*, 3(POPL):36:1–36:29, 2019.
- [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 *FoSSaCS*, pages 418–433, 2013.

APPENDIX A. CONCURRENT GAMES TOOLBOX

A.1. Event Structures and Maps. We start with recalling *maps of event structures*.

Definition A.1. Consider E, F two es. A **map of event structures** from E to F is a function $f : |E| \rightarrow |F|$ satisfying the following two conditions:

- valid:* for all $x \in \mathcal{C}(E)$, we have $f x \in \mathcal{C}(F)$,
local injectivity: for all $e_1, e_2 \in x \in \mathcal{C}(E)$, if $f e_1 = f e_2$ then $e_1 = e_2$.

If E and F are esps, a **map of esps** is additionally required to preserve polarities.

Lemma A.2. Consider $f : E \rightarrow F$ a map of event structures, and $e_1, e_2 \in x \in \mathcal{C}(E)$.

If $f(e_1) \leq_F f(e_2)$, then $e_1 \leq_E e_2$.

Proof. Seeking a contradiction, assume we do not have $e_1 \leq_E e_2$. This means that $e_1 \notin [e_2]_E$. But $[e_2]_E \in \mathcal{C}(E)$, so $f[e_2]_E$ is down-closed. Moreover, $f e_2 \in f[e_2]_E$, so $f e_1 \in f[e_2]_E$. So there is $e'_1 \in [e_2]_E$ such that $f e'_1 = f e_1$. Finally, $[e_2]_E \subseteq x$, so $e_1, e'_1 \in x \in \mathcal{C}(E)$. Hence, by local injectivity, $e_1 = e'_1$. We deduce that $e_1 \in [e_2]_E$ after all, contradiction. \square

A.2. Basic Properties of Strategies. We gather some basic properties.

Lemma A.3. Consider A an arena, $\sigma : A$ a causal strategy, and $m, n \in |\sigma|$.

If $m \rightarrow_\sigma n$, then $\text{pol}(m) \neq \text{pol}(n)$.

Proof. If $\text{pol}(m) = \text{pol}(n)$, then by *courteous* we have $\partial_\sigma(m) \rightarrow_A \partial_\sigma(n)$ as well. But since arenas are alternating, this is absurd. \square

Lemma A.4. Consider A an arena, $\sigma : A$ a causal strategy, and $m, n^- \in |\sigma|$ compatible.

Then, $m \rightarrow_\sigma n$ iff $\partial_\sigma(m) \rightarrow_A \partial_\sigma(n)$.

Proof. *If.* Assume $\partial_\sigma(m) \rightarrow_A \partial_\sigma(n)$. Since $[n]_\sigma \in \mathcal{C}(\sigma)$ and ∂_σ is a map of event structures, $\partial_\sigma[n]_\sigma \in \mathcal{C}(A)$, so it is down-closed. Thus, there is $m' \in [n]_\sigma$ such that $\partial_\sigma(m') = \partial_\sigma(m)$. In particular, n cannot be minimal, so there is $m'' \rightarrow_\sigma n$. By *courteous*, since n is negative we have $\partial_\sigma(m'') \rightarrow_A \partial_\sigma(n)$. But A is *forestial*, so $\partial_\sigma(m'') = \partial_\sigma(m)$. Now, since m, n are compatible they appear in a configuration $x \in \mathcal{C}(\sigma)$, and in particular $m, m'' \in x$. Thus, $m = m''$ by local injectivity. *Only if.* If $m \rightarrow_\sigma n^-$, then $\partial_\sigma(m) \rightarrow_A \partial_\sigma(n)$ by *courteous*. \square

Lemma A.5. Consider A an arena, $\sigma : A$ a causal strategy, and $m^- \in |\sigma|$ s.t. $\partial_\sigma m$ non-minimal. Then, there is a unique $n \rightarrow_\sigma m$.

Proof. *Existence.* Write $a \rightarrow \partial_\sigma(m)$, which is unique since A is *forestial*. Since $[m]_\sigma \in \mathcal{C}(\sigma)$, we have $\partial_\sigma[m]_\sigma \in \mathcal{C}(A)$, therefore it is down-closed and must contain a . Therefore, there is $n \leq_\sigma m$ such that $\partial_\sigma(n) = a$. By Lemma A.4, we then have $n \rightarrow_\sigma m$ as required.

Uniqueness. If $n_1 \rightarrow_\sigma m$ and $n_2 \rightarrow_\sigma m$, by *courteous* $\partial_\sigma n_1 \rightarrow_A \partial_\sigma m$ and $\partial_\sigma n_2 \rightarrow_A \partial_\sigma m$. But then, $\partial_\sigma n_1 = \partial_\sigma n_2$ as A is *forestial*. So $n_1 = n_2$ as ∂_σ is *locally injective*. \square

Lemma A.6. Consider A an arena, $\sigma : A$ a causal strategy, and $m, n_1, n_2 \in |\sigma|$.

If $n_1 \rightarrow_\sigma m$ and $n_2 \rightarrow_\sigma m$ with n_1, n_2 distinct, then $\text{pol}(m) = +$.

Proof. Seeking a contradiction, assume $\text{pol}(m) = -$. Then, by *courteous*, we have $\partial_\sigma(n_1) \rightarrow_A \partial_\sigma(m)$ and $\partial_\sigma(n_2) \rightarrow_A \partial_\sigma(m)$. As A is *forestial*, $\partial_\sigma(n_1) = \partial_\sigma(n_2)$. As $n_1, n_2 \in [m]_\sigma \in \mathcal{C}(\sigma)$, we have $n_1 = n_2$ by local injectivity, contradiction. \square

Lemma A.7. Consider A an arena, $\sigma : A$ a causal strategy, and $m^+ \in x \in \mathcal{C}(\sigma)$.

Then, m is maximal in x iff $\partial_\sigma m$ is maximal in $\partial_\sigma x$.

Proof. If. Assume $\partial_\sigma m$ is maximal in $\partial_\sigma x$. Seeking a contradiction, assume that $m \rightarrow_\sigma n$. By Lemma A.3, $\text{pol}(n) = -$. Therefore, by *courteous*, $\partial_\sigma(m) \rightarrow_A \partial_\sigma(n)$, contradiction.

Only if. Straightforward by Lemma A.2. \square

Lemma A.8. Consider A an arena, $\sigma : A$ a causal strategy, and $m_1^-, m_2^- \in |\sigma|$. If m_1 and m_2 are in minimal conflict in σ , then $\partial_\sigma(m_1)$ and $\partial_\sigma(m_2)$ are in minimal conflict in A .

Proof. We first prove that $\partial_\sigma(m_1)$ and $\partial_\sigma(m_2)$ are in conflict. Seeking a contradiction, assume that it is not the case. Then, as m_1 and m_2 are in *minimal* conflict, we have $\partial_\sigma([m_1]_\sigma \cup [m_2]_\sigma) \in \mathcal{C}(A)$. Hence, by *receptive*, there is a unique $m'_2 \in |\sigma|$ such that

$$[m_1]_\sigma \cup [m_2]_\sigma \vdash_\sigma m'_2$$

with $\partial_\sigma(m'_2) = \partial_\sigma(m_2)$ – where $[m_2]_\sigma = \{n \in |\sigma| \mid n <_\sigma m_2\}$. But then, by Lemma A.4, m_2 is minimal in σ iff m'_2 is minimal in σ , and so $m_2 = m'_2$ by *receptive*, contradicting that $m_1 \#_\sigma m_2$. Otherwise, consider $n \rightarrow_\sigma m_2$ and $n' \rightarrow_\sigma m'_2$. Then, by *courteous*,

$$\partial_\sigma(n) \rightarrow_A \partial_\sigma(m_2), \quad \partial_\sigma(n') \rightarrow_A \partial_\sigma(m'_2),$$

but then $\partial_\sigma(n) = \partial_\sigma(n')$ since A is *forestial*. But $n, n' \in [m_1]_\sigma \cup [m_2]_\sigma \in \mathcal{C}(\sigma)$, so we must have $n = n'$ by locally injectivity. By Lemma A.5, n is the unique predecessor of m_2 and m'_2 . So, $[n]_\sigma \vdash_\sigma m_2$ and $[n]_\sigma \vdash_\sigma m'_2$ with the same image. So $m_2 = m'_2$, contradiction.

Finally, minimality of the conflict is obvious from that of m_1 and m_2 . \square

APPENDIX B. CONSTRUCTION OF ALTERNATING STRATEGIES

In this second section of the appendix, we give more details on the construction of $\downarrow\uparrow\text{-Strat}$.

First, a warning: quite a few superficial complications come from the general construction $A \multimap B$ for B non-pointed, with morphisms from A to B being strategies on $A \multimap B$. An alternative is to only consider $A \multimap B$ for B strict. Then we do not have a symmetric monoidal closed category, only an exponential ideal. This would be sufficient for the languages considered in this paper, however, we opted to link with traditional categorical models.

B.1. Basic Categorical Structure. We start by establishing the categorical structure.

B.1.1. Arrow arena. First, we give postponed proofs on the construction of the arrow arena.

Lemma B.1. Consider A and B two $--$ arenas.

Then, there is a unique $\#_{A \multimap B}$ making $A \multimap B$ a $--$ arena such that for all down-closed finite $x \subseteq |A \multimap B|$, $x \in \mathcal{C}(A \multimap B)$ iff $\chi_{A,B} x \in \mathcal{C}(A^\perp \parallel B)$ with $\chi_{A,B}$ injective on x .

Proof. Existence. We set $\#_{A \multimap B}$ as the following relation:

$$\begin{aligned} (2, b_1) \#_{A \multimap B} (2, b_2) &\Leftrightarrow b_1 \#_B b_2 \\ (1, (b, a_1)) \#_{A \multimap B} (1, (b, a_2)) &\Leftrightarrow a_1 \#_A a_2 \\ (2, b_1) \#_{A \multimap B} (1, (b_2, a)) &\Leftrightarrow b_1 \#_B b_2 \\ (1, (b_1, a)) \#_{A \multimap B} (2, b_2) &\Leftrightarrow b_1 \#_B b_2 \\ (1, (b_1, a_1)) \#_{A \multimap B} (1, (b_2, a_2)) &\Leftrightarrow (b_1 \#_B b_2) \vee (\min(a_1) = \min(a_2)) \vee (a_1 \#_A a_2), \end{aligned}$$

$$\begin{array}{lll}
\varepsilon \upharpoonright A, B = \varepsilon & \varepsilon \upharpoonright B, C = \varepsilon & \varepsilon \upharpoonright A, C = \varepsilon \\
u[r, c] \upharpoonright A, B = u \upharpoonright A, B & u[r, c] \upharpoonright B, C = (u \upharpoonright B, C)[r, c] & u[r, c] \upharpoonright A, C = (u \upharpoonright A, C)[r, c] \\
u[m, b]_c \upharpoonright A, B = (u \upharpoonright A, B)[r, b] & u[m, b]_c \upharpoonright B, C = (u \upharpoonright B, C)[l, b]_c & u[m, b]_c \upharpoonright A, C = u \upharpoonright A, C \\
u[l, a]_{b,c} \upharpoonright A, B = (u \upharpoonright A, B)[l, a]_b & u[l, a]_{b,c} \upharpoonright B, C = u \upharpoonright B, C & u[r, a]_{b,c} \upharpoonright A, C = (u \upharpoonright A, C)[l, a]_c
\end{array}$$

Figure 68: Restrictions of pre-interactions

where b_1, b_2 are assumed distinct, all other pairs left consistent, and with $\min(a)$ the unique minimal antecedent of a . It is routine that this conflict makes $A \multimap B$ a $--$ -arena.

Now, we check the additional condition. Consider $x \subseteq |A \multimap B|$ down-closed, written as

$$x = (\parallel_{b \in I} x_{A,b}) \parallel x_B$$

where I is a subset of the minimal events of B . Then, we show that $x \in \mathcal{C}(A \multimap B)$ iff

$$\chi_{A,B} x = (\cup_{b \in I} x_{A,b}) \parallel x_B \in \mathcal{C}(A \vdash B)$$

and $\chi_{A,B}$ is injective on x . *Only if* is a direct verification. For *if*, if $\chi_{A,B} x \in \mathcal{C}(A \multimap B)$ then the only possible conflict in x is of the form, with $b_1 \#_B b_2$ or $\min(a_1) = \min(a_2)$:

$$(1, (b_1, a_1)) \#_{A \multimap B} (1, (b_2, a_2))$$

In the former case, by down-closure, $(2, b_1), (2, b_2) \in x$, contradicting $\chi_{A,B} x \in \mathcal{C}(A \vdash B)$. In the latter case, by down-closure, $(1, (b_1, \min(a_1))), (1, (b_2, \min(a_2))) \in x$ – but they have the same image through $\chi_{A,B}$, contradicting that it should be injective on x . Uniqueness follows as with fixed events, an event structure is determined by configurations. \square

B.1.2. Composition. Consider fixed A, B and C three $--$ -arenas, and $\sigma : A \multimap B, \tau : B \multimap C$ alternating prestrategies. Recall from the main text:

Definition B.2. A **pre-interaction** on A, B, C is $u \in |(A \multimap B) \multimap C|^*$ satisfying:

$$\text{valid: } \forall 1 \leq i \leq n, \{s_1, \dots, s_i\} \in \mathcal{C}((A \multimap B) \multimap C)$$

Remember that in $A \multimap B$, events are either $(2, b)$ for $b \in |B|$, or $(1, (b, a))$ for $b \in \min(B)$ and $a \in |A|$. By convention, in this section, we write $[r, b]$ for $(2, b)$ and $[l, a]_b$ for $(1, (b, a))$. In that case, $[r, b]$ is the unique immediate predecessor of $[l, a]_b$, *i.e.* its *justifier*. Similarly, in $(A \multimap B) \multimap C$, events can be $(2, c)$ written $[r, c]$; $(1, (c, (2, b)))$ written $[m, b]_c$; and $(1, (c, (1, (b, a))))$ written $[l, a]_{b,c}$; we respectively say that they are **in C**, **in B**, or **in A**.

Using this, we define in Figure 68 three **restrictions** of a pre-interaction u , namely $u \upharpoonright A, B \in |A \multimap B|^*$, $u \upharpoonright B, C \in |B \multimap C|^*$, and $u \upharpoonright A, C \in |A \multimap C|^*$. Now we set:

Definition B.3. An **interaction** $u \in \tau \otimes \sigma$ between σ and τ is a pre-interaction u s.t.

$$u \upharpoonright A, B \in \sigma, \quad u \upharpoonright B, C \in \tau, \quad u \upharpoonright A, C \in \downarrow\uparrow\text{-Plays}(A \multimap C).$$

The **composition** of σ and τ is $\tau \odot \sigma = \{u \upharpoonright A, C \mid u \in \tau \otimes \sigma\}$.

The first step is to ensure that $\tau \odot \sigma$ is a prestrategy, and that if σ and τ are strategies, then so is $\tau \odot \sigma$. We start with the conditions of Definition 2.8, postponing uniformity. *Non-empty* and *prefix-closed* follow from those on σ and τ . For *deterministic*, we need more tools. The main tool to study the interaction of alternating strategies is the *state analysis* of plays and interactions. Recall from Section 6.1.3 that s alternating is **in state** O if it has even length, and **in state** P otherwise. Then, we have the following property:

Lemma B.4. *If $u \in \tau \otimes \sigma$, then we are in one of the following three cases:*

- (1) $u \upharpoonright A, B, u \upharpoonright B, C$ and $u \upharpoonright A, C$ are respectively in state O, O, O .
- (2) $u \upharpoonright A, B, u \upharpoonright B, C$ and $u \upharpoonright A, C$ are respectively in state O, P, P .
- (3) $u \upharpoonright A, B, u \upharpoonright B, C$ and $u \upharpoonright A, C$ are respectively in state P, O, P .

Proof. Standard argument, direct by induction on u . □

Next, we can prove the key property of the composition of alternating strategies.

Lemma B.5. *Consider $s \in \tau \odot \sigma$ of even length.*

*Then, there is a unique **witness** $u \in \tau \otimes \sigma$ such that $u \upharpoonright A, C = s$.*

Proof. *Existence* is obvious by definition.

Uniqueness. Seeking a contradiction, consider $u_1, u_2 \in \tau \otimes \sigma$ distinct such that $u_1 \upharpoonright A, C = u_2 \upharpoonright A, C = s$. First, since s has even length, $u_i \upharpoonright A, C$ is in state O , so u_1 and u_2 must be in state (1) of Lemma B.4. It follows that their immediate prefix cannot be in state (1), from which it follows last move is visible (*i.e.* in A or C). So, u_1 and u_2 cannot be comparable for the prefix order. Therefore, there is u' maximal such that $u' \sqsubseteq u_1$ and $u' \sqsubseteq u_2$, say we have $u'm_1 \sqsubseteq u_1$ and $u'm_2 \sqsubseteq u_2$ for m_1, m_2 distinct. By Lemma B.4, u' is in one of the states (1), (2) or (3). If it is in state (1), then the next moves m_1 and m_2 are in $u \upharpoonright A, C = s$, so they cannot be distinct. Say u' is in state (3) – the case (2) is similar but simpler. Necessarily, m_1 and m_2 are in A or B , so that

$$u'm_1 \upharpoonright A, B = (u' \upharpoonright A, B)n_1 \quad u'm_2 \upharpoonright A, B = (u' \upharpoonright A, B)n_2$$

are two plays of σ with even length – so that $n_1 = n_2$, by *determinism*. From the definition of restriction, the only case where $m_1 \neq m_2$ with $n_1 = n_2$ is if $m_1 = [l, a]_{b,c}$, $m_2 = [l, a]_{b,c'}$ with $c \neq c'$, so $n_1 = n_2 = [l, a]_b$. Then, as u' satisfies condition *valid*, this entails that their immediate dependencies $[m, b]_c, [m, b]_{c'} \in |u'|$ as well – impossible since $[m, b]_c \# [m, b]_{c'}$. □

It is then straightforward to prove that $\tau \odot \sigma$ satisfies *deterministic*.

Proposition B.6. *We have that $\tau \odot \sigma : A \multimap C$ is a prestrategy.*

Moreover, if σ and τ are strategies, then so is $\tau \odot \sigma$.

Proof. To obtain a prestrategy, it remains that $\tau \odot \sigma$ satisfies *deterministic*. Consider $sn_1^+, sn_2^+ \in \tau \odot \sigma$. Consider $u_1m_1, u_2m_2 \in \tau \otimes \sigma$ their unique witness as given by Lemma B.5. We reason as in Lemma B.5: if there is a diverging point between u_1m_1 and u_2m_2 , by Lemma B.4 the divergence can be attributed to either σ or τ , contradicting *determinism*.

Now, assume that σ and τ satisfy condition *receptive*. Consider $s \in \tau \odot \sigma$ such that $sm^- \in \downarrow\uparrow\text{-Plays}(A \multimap C)$. More precisely, assume that $m = [(l, a)]_c$ as the case in C is simpler. Consider now $u \in \tau \otimes \sigma$ such that $u \upharpoonright A, C = s$; necessarily u is in state (1) of Lemma B.4. Now, as m is negative its immediate predecessor of m in $A \multimap C$ is some $[(l, a')]_c$ in s . Since $s = u \upharpoonright A, C$, it corresponds to some $[(l, a')]_{b,c}$ in u , for some $b \in \min(B)$. But then, it is a direct verification that $(u \upharpoonright A, B)[(l, a)]_b \in \downarrow\uparrow\text{-Plays}(A \multimap B)$, so $(u \upharpoonright A, B)[(l, a)]_b \in \sigma$ by *receptive*. Therefore, $u[(l, a)]_{b,c} \in \tau \otimes \sigma$ witnessing that $sm \in \tau \odot \sigma$ as required. □

Note that this argument and the prior state analysis of Lemma B.4 are also found in the proof of composition of sequentiality for causal strategies in Section 6.1.3.

B.1.3. Associativity. We now sketch associativity, which follows standard lines, see *e.g.* [Har04]. Let us fix $\sigma : A \multimap B, \tau : B \multimap C$ and $\delta : C \multimap D$ three alternating prestrategies.

The first step is to define a notion of interaction between three strategies. First:

Definition B.7. A **3-pre-interaction** on A, B, C, D is $w \in |((A \multimap B) \multimap C) \multimap D|^*$ s.t.

$$\text{valid: } \forall 1 \leq i \leq n, \{w_1, \dots, w_i\} \in \mathcal{C}((A \multimap B) \multimap C),$$

where $w = w_1 \dots w_n$.

For a 3-pre-interaction w , we define $w \upharpoonright A, B \in |A \multimap B|^*$, $w \upharpoonright B, C \in |B \multimap C|^*$, $w \upharpoonright C, D \in |C \multimap D|^*$ and $w \upharpoonright A, D \in |A \multimap D|^*$ with the obvious adaptation of Figure 68. An **interaction** of σ, τ and δ is a 3-pre-interaction w s.t.

$$w \upharpoonright A, B \in \sigma, \quad w \upharpoonright B, C \in \tau, \quad w \upharpoonright C, D \in \delta, \quad w \upharpoonright A, D \in \downarrow\uparrow\text{-Plays}(A \multimap D),$$

written $w \in \delta \otimes \tau \otimes \sigma$. We have four additional restrictions $w \upharpoonright A, B, C \in \tau \otimes \sigma$, $w \upharpoonright B, C, D \in \delta \otimes \tau$, $w \upharpoonright A, C, D \in \delta \circledast (\tau \odot \sigma)$ and $w \upharpoonright A, B, D \in (\delta \odot \tau) \otimes \sigma$, defined in the obvious way.

Then, the key argument of associativity is the so-called ‘‘zipping lemma’’:

Lemma B.8 (Zipping). Consider $u \in \delta \circledast (\tau \odot \sigma)$ and $v \in \tau \otimes \sigma$ such that $u \upharpoonright A, C = v \upharpoonright A, C$.

Then, there is a unique $w \in \delta \otimes \tau \otimes \sigma$ such that $w \upharpoonright A, C, D = u$ and $w \upharpoonright A, B, C = v$.

Proof. By induction on u – by Lemma B.4 the moves in B from v can be interleaved with those in v in a unique way; likewise there is a unique way to set their dependency. \square

We also have the mirror image, zipping $u \in (\delta \odot \tau) \otimes \sigma$ with $v \in \delta \otimes \tau$. Altogether,

Proposition B.9. We have $(\delta \odot \tau) \odot \sigma = \delta \odot (\tau \odot \sigma)$.

Proof. Consider $s \in \delta \odot (\tau \odot \sigma)$. It has a (unique) witness $u \in \delta \circledast (\tau \odot \sigma)$. Then, $u \upharpoonright A, C \in \tau \odot \sigma$, thus there is again $v \in \tau \otimes \sigma$ s.t. $u \upharpoonright A, C = v \upharpoonright A, C$. By Lemma B.8, there is $w \in \delta \otimes \tau \otimes \sigma$ s.t. $w \upharpoonright A, C, D = u$ and $w \upharpoonright A, B, C = v$. But then we may restrict w to $w \upharpoonright B, C, D \in \delta \otimes \tau$, so $w \upharpoonright B, D \in \delta \odot \tau$. Moreover $w \upharpoonright A, B \in \sigma$ so $w \upharpoonright A, B, D \in (\delta \odot \tau) \otimes \sigma$, from which $w \upharpoonright A, D \in (\delta \odot \tau) \odot \sigma$. The other direction is symmetric. \square

Note that associativity holds for prestrategies and does not depend on *receptive*.

B.1.4. Identities. Fix some \multimap -arena A . For $s \in \downarrow\uparrow\text{-Plays}(A \multimap A)$, we define its restrictions

$$\begin{array}{ll} \varepsilon \upharpoonright l = \varepsilon & \varepsilon \upharpoonright r = \varepsilon \\ s[(l, a)]_{a'} \upharpoonright l = (s \upharpoonright l) a & s[(l, a)]_{a'} \upharpoonright r = s \upharpoonright r \\ s[(r, a)] \upharpoonright r = s \upharpoonright r & s[(r, a)] \upharpoonright l = (s \upharpoonright l) a \end{array}$$

using which we may define the identity for alternating strategies:

Definition B.10. The **copycat** $\mathfrak{c}_A : A \multimap A$ comprises all $s \in \downarrow\uparrow\text{-Plays}(A \multimap A)$ s.t.

- (1) for all $s' \sqsubseteq s$ of even length, $s' \upharpoonright l = s' \upharpoonright r$,
- (2) for all $[(l, a)]_{a'} \in |s|$, with $a \in \min(A)$, $a = a'$.

Condition (2) means that when playing a minimal event on the left hand side, copycat justifies it with the same move on the right. Such a condition is also required in Hyland-Ong games (though sometimes mistakenly omitted). Copycat strategies provide identities:

Proposition B.11. The \multimap -arenas and alternating strategies form a category $\downarrow\uparrow\text{-Strat}$.

Proof. It remains that for any $\sigma : A \multimap B$, $\sigma = \mathfrak{c}_B \odot \sigma \odot \mathfrak{c}_A$, which is elementary. \square

$$\begin{array}{ll}
\varepsilon \upharpoonright A_1, B_1 = \varepsilon & \varepsilon \upharpoonright A_2, B_2 = \varepsilon \\
s[(l, (1, a))]_{(1,b)} \upharpoonright A_1, B_1 = (s \upharpoonright A_1, B_1)[(l, a)]_b & s[(l, (1, a))]_{(i,b)} \upharpoonright A_1, B_1 = s \upharpoonright A_2, B_2 \\
s[(l, (2, a))]_{(i,b)} \upharpoonright A_1, B_1 = s \upharpoonright A_1, B_1 & s[(l, (2, a))]_{(2,b)} \upharpoonright A_2, B_2 = (s \upharpoonright A_2, B_2)[(l, a)]_b \\
s[(r, (1, b))] \upharpoonright A_1, B_1 = (s \upharpoonright A_1, B_1)[(r, b)] & s[(r, (1, b))] \upharpoonright A_2, B_2 = s \upharpoonright A_2, B_2 \\
s[(r, (2, b))] \upharpoonright A_1, B_1 = s \upharpoonright A_1, B_1 & s[(r, (2, b))] \upharpoonright A_2, B_2 = (s \upharpoonright A_2, B_2)[(r, b)]
\end{array}$$

Figure 69: Partial restrictions for the tensor

B.2. Monoidal Closed Structure. We now describe the monoidal structure.

B.2.1. Tensor product. On $--$ arenas, we have defined $A \otimes B$ simply as $A \parallel B$.

For strategies, the critical step is a suitable notion of restriction. More precisely, for A_1, A_2, B_1, B_2 $--$ arenas and $s \in \downarrow\uparrow\text{-Plays}(A_1 \otimes A_2 \multimap B_1 \otimes B_2)$, we give a partial definition

$$s \upharpoonright A_1, B_1 \in |A_1 \multimap B_1|^* \quad s \upharpoonright A_2, B_2 \in |A_2 \multimap B_2|^*$$

in Figure 69 – partial, because *e.g.* $[(l, (1, a))]_{(2,b)} \upharpoonright A_1, B_1$ is left undefined. We then set:

Definition B.12. Consider $\sigma_1 : A_1 \multimap B_1$ and $\sigma_2 : A_2 \multimap B_2$ alternating strategies. Then:

$$\sigma_1 \otimes \sigma_2 = \{s \in \downarrow\uparrow\text{-Plays}(A_1 \otimes A_2 \multimap B_1 \otimes B_2) \mid \forall i \in \{1, 2\}, s \upharpoonright A_i, B_i \in \sigma_i\},$$

implying in particular that for each $s \in \sigma_1 \otimes \sigma_2$ and $i \in \{1, 2\}$, $s \upharpoonright A_i, B_i$ is defined.

By definition, $\sigma_1 \otimes \sigma_2$ satisfies *non-empty* and *prefix-closed*. As for composition, *determinism* involves performing a state analysis expressing that at each point, only one of σ_1 or σ_2 has control. We skip the details. See *e.g.* [Har04] for an analogous proof, also reflected in the proof in Section 6.1.4 that sequential causal strategies are stable under tensor.

Proposition B.13. Consider $\sigma_1 : A_1 \multimap B_1$ and $\sigma_2 : A_2 \multimap B_2$ alternating strategies.

Then, $\sigma_1 \otimes \sigma_2 : A_1 \otimes A_2 \multimap B_1 \otimes B_2$ is an alternating strategy.

Fix $\sigma_1 : A_1 \multimap B_1, \sigma_2 : A_2 \multimap B_2, \tau_1 : B_1 \multimap C_1$ and $\tau_2 : B_2 \multimap C_2$. For $w \in (\tau_1 \otimes \tau_2) \otimes (\sigma_1 \otimes \sigma_2)$, we first define partially $w \upharpoonright A_1, B_1, C_1$ and $w \upharpoonright A_2, B_2, C_2$ analogously to Figure 69 – it is direct to prove that $w \upharpoonright A_1, B_1, C_1 \in \tau_1 \otimes \sigma_1$ and $w \upharpoonright A_2, B_2, C_2 \in \tau_2 \otimes \sigma_2$.

Functoriality is analogous to associativity in that it relies on a zipping lemma:

Lemma B.14. Consider $u_1 \in \tau_1 \otimes \sigma_1, u_2 \in \tau_2 \otimes \sigma_2$, and $s \in \downarrow\uparrow\text{-Plays}(A_1 \otimes A_2 \multimap C_1 \otimes C_2)$ such that $s \upharpoonright A_1, C_1 = u_1 \upharpoonright A_1, C_1$ and $s \upharpoonright A_2, C_2 = u_2 \upharpoonright A_2, C_2$.

Then, there is a unique $w \in (\tau_1 \otimes \tau_2) \otimes (\sigma_1 \otimes \sigma_2)$ such that

$$s = w \upharpoonright A_1 \otimes A_2, C_1 \otimes C_2, \quad s \upharpoonright A_1, B_1, C_1 = u_1, \quad s \upharpoonright A_2, B_2, C_2 = u_2.$$

Proposition B.15. The construction \otimes extends to a bifunctor

$$\otimes : \downarrow\uparrow\text{-Strat} \times \downarrow\uparrow\text{-Strat} \rightarrow \downarrow\uparrow\text{-Strat}.$$

Proof. Preservation of identities is direct. Functoriality follows by Lemma B.14 \square

Next, we complete the symmetric monoidal structure by providing the structural natural isomorphisms. We first introduce a few tools useful in giving a clean definition of such structural isomorphisms, which are variants of the copycat strategy. We shall make use of certain morphisms, called *renamings*, to act on strategies.

Definition B.16. A **renaming** from arena A to B is a function $f : |A| \rightarrow |B|$ satisfying:

$$\begin{aligned}
\text{validity:} & \quad \forall x \in \mathcal{C}(A), fx \in \mathcal{C}(B) \\
\text{local injectivity:} & \quad \forall a_1, a_2 \in x \in \mathcal{C}(A), fa_1 = fa_2 \implies a_1 = a_2 \\
\text{polarity-preserving:} & \quad \forall a \in |A|, \text{pol}_B(fa) = \text{pol}_A(a) \\
\text{symmetry-preserving:} & \quad \forall \theta \in \mathcal{S}(A) \text{ (resp. } \mathcal{S}_+(A), \mathcal{S}_-(A)), f\theta \in \mathcal{S}(B), \text{ (resp. } \mathcal{S}_+(B), \mathcal{S}_-(B)) \\
\text{strong-receptivity:} & \quad \text{for all } \theta \in \mathcal{S}(A), \text{ for all } f\theta \subseteq^- \varphi \in \mathcal{S}(B), \exists! \theta' \subseteq^- \theta' \in \mathcal{S}(A) f\theta' = \varphi \\
\text{courtesy:} & \quad \forall a_1 \rightarrow_A a_2, (\text{pol}_A(a_1) = + \vee \text{pol}_A(a_2) = -) \implies fa_1 \rightarrow_B fa_2.
\end{aligned}$$

We write $f : A \rightarrow B$ to mean that f is a renaming from A to B .

This construction is imported from [CCW19]. Then we have:

Definition B.17. Consider $\sigma : A \multimap B$ and renamings $g : B \rightarrow B', f : A^\perp \rightarrow A'^\perp$. We set

$$g \cdot \sigma \cdot f = \{g \cdot s \cdot f \mid s \in \sigma\} : A' \multimap B'$$

where $g \cdot s \cdot f$ acts on s event-wise, sending $[(l, a)]_b$ to $[(l, f(a))]_{g(b)}$ and $[(r, b)]$ to $[(r, g(b))]$.

It is direct that $g \cdot \sigma \cdot f$ is a strategy. The structural isomorphisms that we aim to define are obtained by *lifting* renamings. Indeed if $f : A \rightarrow B$ is a renaming, we may define

$$\overrightarrow{f} = f \cdot \mathfrak{c}_A : A \multimap B.$$

a renaming from copycat. Likewise, from $f : B^\perp \rightarrow A^\perp$ we set $\overleftarrow{f} = \mathfrak{c}_B \cdot f : A \multimap B$.

The main property satisfied by these constructions is the following *lifting lemma*.

Lemma B.18. Consider $\sigma : A \multimap B$ a strategy, and $f : A^\perp \rightarrow A'^\perp, g : B \rightarrow B'$ two renamings. Then, we have $\overrightarrow{g} \odot \sigma \odot \overleftarrow{f} = g \cdot \sigma \cdot f$.

Proof. A direct adaptation of the neutrality of copycat under composition. \square

Before we put these to use to construct the symmetric monoidal structure, we deduce a few properties of lifted strategies. In the statement below, for any renaming $f : A \rightarrow B$ which is additionally an isomorphism, then we write $f^\perp : B^\perp \rightarrow A^\perp$ for its inverse with polarities reversed – it is immediate that it still satisfies the conditions of a renaming.

Proposition B.19. *Lifting is a functor from the category of renamings to $\downarrow\uparrow\text{-Strat}$.*

Moreover, if $f : A \rightarrow B$ is an iso, \overrightarrow{f} is an iso; and we have $\overrightarrow{f}^{-1} = \overleftarrow{f^\perp}$ and $\overleftarrow{f} = \overrightarrow{f^\perp}$.

Proof. By Lemma B.18 and direct verifications. \square

With this, we may now define the structural isomorphisms for the symmetric monoidal structure. We notice that for all arenas A, B, C , there are invertible renamings:

$$\begin{aligned}
\rho_A & : & A \otimes 1 & \rightarrow A \\
\lambda_A & : & 1 \otimes A & \rightarrow A \\
\alpha_{A,B,C} & : & (A \otimes B) \otimes C & \rightarrow A \otimes (B \otimes C) \\
s_{A,B} & : & A \otimes B & \rightarrow B \otimes A
\end{aligned}$$

where 1 is the empty arena, satisfying the coherence laws of a symmetric monoidal category. The required structural isomorphisms are simply obtained by lifting those as $\overrightarrow{\rho_A}, \overrightarrow{\lambda_A}, \overrightarrow{\alpha_{A,B,C}}$ and $\overrightarrow{s_{A,B}}$. By Proposition B.19, the required coherence diagrams are still satisfied.

Proposition B.20. *The category $\downarrow\uparrow\text{-Strat}$ is a symmetric monoidal category.*

Proof. It remains to prove naturality, easy from Proposition B.19 and direct verifications. \square

B.2.2. *Cartesian products.* $\downarrow\uparrow\text{-Strat}$ has cartesian products, given by $\&$ on objects.

To perform the **pairing** of $\sigma_1 : A \multimap B_1$ and $\sigma_2 : A \multimap B_2$, we first build prestrategies

$$\text{inj}_1(\sigma_1) : A \multimap B_1 \& B_2, \quad \text{inj}_2(\sigma_2) : A \multimap B_1 \& B_2$$

defined by applying event-wise $\text{inj}_i([(r, b)]) = [(r, (i, b))]$, $\text{inj}_i([(l, a)]_b) = [(l, a)]_{(i, b)}$. We obtain $\langle \sigma_1, \sigma_2 \rangle = \text{inj}_1(\sigma_1) \uplus \text{inj}_2(\sigma_2)$. It is direct that this yields a bijection:

$$\langle -, - \rangle : \downarrow\uparrow\text{-Strat}(A, B_1) \times \downarrow\uparrow\text{-Strat}(A, B_2) \rightarrow \downarrow\uparrow\text{-Strat}(A, B_1 \& B_2)$$

The **projections** are $(\pi_1, \pi_2) = \langle -, - \rangle^{-1}(\text{id}_{B_1 \& B_2})$, and we can verify

$$\pi_1 \odot \langle \sigma_1, \sigma_2 \rangle = \sigma_1, \quad \pi_2 \odot \langle \sigma_1, \sigma_2 \rangle = \sigma_2$$

which is enough to complete the cartesian structure of $\downarrow\uparrow\text{-Strat}$.

B.2.3. *Monoidal closure.* First, for $--$ -arenas A, B and C , there is a clear isomorphism

$$(A \otimes B) \multimap C \cong A \multimap (B \multimap C)$$

which, applied event-wise, yields $\Lambda(-) : \downarrow\uparrow\text{-Strat}(A \otimes B, C) \simeq \downarrow\uparrow\text{-Strat}(A, B \multimap C)$. Then,

$$\text{ev}_{A, B} = \Lambda^{-1}(\text{id}_{A \multimap B}) : (A \multimap B) \otimes A \multimap B$$

is the **evaluation** strategy. It is a verification akin to the neutrality of copycat that for all $\sigma : A \multimap (B \multimap C)$, we have $\text{ev}_{A, C} \odot (\sigma \otimes B) = \Lambda^{-1}(\sigma)$.

Proposition B.21. *$\downarrow\uparrow\text{-Strat}$ is a cartesian symmetric monoidal closed category.*

B.3. **Symmetry.** We now develop the structure pertaining to symmetry.

B.3.1. *Basic structure.* The extension of the construction above with symmetry and uniformity unfolds essentially as in AJM games [AJM00]. We describe the main steps.

First of all, we ensure all structural morphisms are uniform. This is the purpose of:

Lemma B.22. *For A and B $--$ -arenas and $f : A \rightarrow B$ a renaming, $\vec{f} \approx \vec{f}$.*

Proof. Straightforward by *symmetry-preserving* and *strong-receptivity* of renamings. \square

Next, we show operations on strategies are compatible with \approx . The delicate case is composition. Fix A, B and C $--$ -arenas, and write $I = (A \multimap B) \multimap C$. We shall give to events of $(A \multimap B) \multimap C$ a *polarity* corresponding to their role in an interaction: a move m is *negative* if it is in A or C and is negative for $A \multimap C$, and has *polarity* p otherwise.

The main tool is the following lifting of Definition 2.13 to interactions:

Definition B.23. Consider $\sigma, \sigma' : A \multimap B$ and $\tau, \tau' : B \multimap C$. We write $\tau \otimes \sigma \approx \tau' \otimes \sigma'$ iff

$$\text{-->simulation: } \forall um^p \in \tau \otimes \sigma, v \in \tau' \otimes \sigma', u \cong_I v \implies \exists n^p, vn \in \tau' \otimes \sigma' \& um \cong_I vn$$

$$\text{--<simulation: } \forall u \in \tau \otimes \sigma, vn^p \in \tau' \otimes \sigma', u \cong_I v \implies \exists m, um \in \tau \otimes \sigma \& um \cong_I vn$$

$$\text{-->receptive: } \forall um^- \in \tau \otimes \sigma, v \in \tau' \otimes \sigma', um^- \cong_I vn^- \implies vn^- \in \tau' \otimes \sigma'$$

$$\text{--<receptive: } \forall u \in \tau \otimes \sigma, vn^- \in \tau' \otimes \sigma', um^- \cong_I vn^- \implies um^- \in \tau \otimes \sigma$$

Lemma B.24. *Consider $\sigma, \sigma' : A \multimap B$ and $\tau, \tau' : B \multimap C$ such that $\sigma \approx \sigma'$ and $\tau \approx \tau'$.*

Then, $\tau \otimes \sigma \approx \tau' \otimes \sigma'$.

$$\begin{aligned}
\varepsilon \uparrow i &= \varepsilon \\
s[(l, (i, a))]_{(i,b)} \uparrow i &= (s \uparrow i)[(l, a)]_b \\
s[(l, (j, a))]_{(k,b)} \uparrow i &= s \uparrow i & (i \neq j) \\
s[(r, (i, b))] \uparrow i &= (s \uparrow i)[(r, b)] \\
s[(r, (j, b))] \uparrow i &= s \uparrow i & (i \neq j)
\end{aligned}$$

Figure 70: Partial restrictions for !

Proof. \rightarrow -simulation. Consider $um^p \in \tau \otimes \sigma$ and $v \in \tau' \otimes \sigma'$, s.t. $u \cong_v v$. Necessarily, m is positive for σ or τ , *w.l.o.g.* say σ . Again, we distinguish cases whether m is in A or B . We consider B , which is the most interesting case – so that $m = [(m, b)]_c$ for some $c \in \min(C)$.

We project $u[(m, b)]_c \uparrow A, B = (u \uparrow A, B)[(r, b)]^+ \in \sigma$, $v \uparrow A, B \in \sigma'$ with $u \uparrow A, B \cong_{A \rightarrow B} v \uparrow A, B$. By \rightarrow -simulation for σ , there is $[(r, b')]$ s.t. $(v \uparrow A, B)[(r, b')] \in \sigma'$ and

$$(u \uparrow A, B)[(r, b)] \cong_{A \rightarrow B} (v \uparrow A, B)[(r, b')]. \quad (\text{B.1})$$

However, we also have $(u \uparrow B, C)[(l, b)]_c^- \in \tau$, $v \uparrow B, C \in \tau'$. We also have $u \uparrow B, C \cong_{B \rightarrow C} v \uparrow B, C$. Necessarily, $(r, c) \in |u \uparrow B, C|$ and there is a symmetric $(r, c') \in |v \uparrow B, C|$. Then, from (B.1) and the definition of symmetry on $A \rightarrow B$ and $B \rightarrow C$,

$$(u \uparrow B, C)[(l, b)]_c \cong_{B \rightarrow C} (v \uparrow B, C)[(l, b')]_{c'}$$

so by receptivity, $(v \uparrow B, C)[(l, b')]_{c'} \in \tau'$. It follows that $v[(m, b')]_{c'} \in \tau' \otimes \sigma'$ with $u[(m, b)]_c \cong_I v[(m, b')]_{c'}$ as required. The condition \leftarrow -simulation is symmetric. \square

Finally, \rightarrow -receptive and \leftarrow -receptive are similar but simpler. \square

Finally, we may deduce compatibility of \approx with composition:

Corollary B.25. *Consider $\sigma, \sigma' : A \rightarrow B$ and $\tau, \tau' : B \rightarrow C$ such that $\sigma \approx \sigma'$ and $\tau \approx \tau'$. Then, $\tau \odot \sigma \approx \tau' \odot \sigma'$.*

Proof. Direct from Lemmas B.5 and B.24. \square

Lemma B.5 is crucial: this fails if we do not have a unique witness. This is the main reason why this approach to uniformity does not extend to $\mathcal{U}\text{-Strat}$. Other operations on strategies are easily seen to be compatible with \approx . Therefore, considering $\downarrow\uparrow\text{-Strat}$ as having as only morphisms the uniform strategies (*i.e.* self-equivalent for \approx), it is equipped with an additional equivalence relation \approx with respect to which all operations are compatible.

B.4. Seely category. We now provide the last ingredients to the Seely category.

First, we need a functor $! : \downarrow\uparrow\text{-Strat} \rightarrow \downarrow\uparrow\text{-Strat}$. The construction is similar to the tensor and defined by a suitable restriction, given in Figure 70. Armed with this, we set:

Definition B.26. Consider $\sigma : A \rightarrow B$ an alternating strategy. Then, we set:

$$!\sigma = \{s \in \downarrow\uparrow\text{-Plays}(!A \rightarrow !B \mid \forall i \in \mathbb{N}, s \uparrow i \in \sigma)\},$$

implying in particular that for each $s \in !\sigma$ and $i \in \mathbb{N}$, $s \uparrow i$ is defined.

This is really an infinitary tensor of σ . That this yields $!\sigma : !A \rightarrow !B$ an alternating strategy, along with functoriality, are as for the tensor. We skip the details.

Proposition B.27. *The construction ! extends to a functor $! : \downarrow\uparrow\text{-Strat} \rightarrow \downarrow\uparrow\text{-Strat}$.*

$$\begin{array}{ll}
 \eta_A : & A^\perp \rightarrow (!A)^\perp \\
 & a \mapsto (i, a) \\
 \mu_A : & (!!A)^\perp \rightarrow (!A)^\perp \\
 & (i, (j, a)) \mapsto (\langle i, j \rangle, a) \\
 \text{see}_{A,B}^{\rightarrow} & !A \otimes !B \rightarrow !(A \& B) \\
 & (1, (i, a)) \mapsto (2i, (1, a)) \\
 & (2, (i, b)) \mapsto (2i + 1, (2, a)) \\
 \text{see}_{A,B}^{\leftarrow} & !(A \& B) \rightarrow !A \otimes !B \\
 & (i, (1, a)) \mapsto (1, (i, a)) \\
 & (i, (2, b)) \mapsto (2, (i, b))
 \end{array}$$

Figure 71: Comonad renamings

Figure 72: Seely renamings

The Seely structure also includes structural morphisms, all defined through lifting. We define renamings in Figures 71 and 72 (with $\langle -, - \rangle : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$ any bijection), set

$$\begin{array}{ll}
 \text{der}_A = \overleftarrow{\eta}_A & : \quad !A \multimap A \\
 \text{dig}_A = \overleftarrow{\mu}_A & : \quad !A \multimap !!A \\
 \text{mon}_{A,B} = \overrightarrow{\text{see}}_{A,B} & : \quad !A \otimes !B \multimap !(A \& B) \\
 \text{mon}_{A,B}^{-1} = \overleftarrow{\text{see}}_{A,B} & : \quad !(A \& B) \multimap !A \otimes !B,
 \end{array}$$

and all coherence and naturality properties follow from Lemma B.18 with direct verifications. A number of those only hold up to \approx : for instance, $\text{mon}_{A,B}$ and $\text{mon}_{A,B}^{-1}$ are inverses up to \approx , but not up to equality. Finally, $\text{mon}^0 : \mathbf{1} \multimap \mathbf{1}$ is the empty strategy. Altogether:

Proposition B.28. *The category $\downarrow\uparrow\text{-Strat}$ is a Seely category.*

APPENDIX C. CONSTRUCTION OF NON-ALTERNATING STRATEGIES

Next, we construct the category $\mathcal{U}\text{-Strat}$ of $--$ -arenas and non-alternating strategies.

C.1. Basic Categorical Structure. We first study composition and copycat. For composition, the development is essentially a simpler version of $\downarrow\uparrow\text{-Strat}$ as there is no determinism requirement. On the other hand, copycat requires slightly more care.

C.1.1. Composition. Consider fixed A, B and C three $--$ -arenas, and $\sigma : A \multimap B$, $\tau : B \multimap C$ non-alternating prestrategies. Re-using pre-interactions and restrictions from before, we set:

Definition C.1. An **interaction** $u \in \tau \circledast \sigma$, is a pre-interaction u such that

$$u \upharpoonright A, B \in \sigma, \quad u \upharpoonright B, C \in \tau,$$

and we automatically have $u \upharpoonright A, C \in \mathcal{U}\text{-Plays}(A \multimap C)$.

The **composition** of σ and τ is $\tau \odot \sigma = \{u \upharpoonright A, C \mid u \in \tau \circledast \sigma\}$.

It is clear this yields a non-alternating prestrategy, which only requires *non-empty* and *prefix-closed*. In composing non-alternating (pre)strategies, there is no analogue of Lemma B.4: all players can move anytime. The unique witness property is lost: $s \in \tau \odot \sigma$ if there is a witness $u \in \tau \circledast \sigma$ such that $s = u \upharpoonright A, C$, but u is in general not unique²⁹.

That composition preserves strategies will follow from Propositions C.3 and C.5.

²⁹Causal strategies in $\rightarrow\text{-Strat}$, with their explicit branching information, recover a unique witness property.

C.1.2. *Associativity.* Fix $\sigma : A \multimap B$, $\tau : B \multimap C$ and $\delta : C \multimap D$ non-alternating prestrategies. As above, an **interaction** of σ, τ and δ is a 3-pre-interaction $w \in |((A \multimap B) \multimap C) \multimap D|^*$ such that

$$w \upharpoonright A, B \in \sigma, \quad w \upharpoonright B, C \in \tau, \quad w \upharpoonright C, D \in \delta,$$

from which it follows automatically that $w \upharpoonright A, D \in \mathcal{C}\text{-Plays}(A \multimap D)$.

Again, associativity relies on a “zipping lemma”.

Lemma C.2 (Zipping). *Consider $u \in \delta \circledast (\tau \circledast \sigma)$ and $v \in \tau \circledast \sigma$ such that $u \upharpoonright A, C = v \upharpoonright A, C$. Then, there is $w \in \delta \circledast \tau \circledast \sigma$ such that $w \upharpoonright A, C, D = u$ and $w \upharpoonright A, B, C = v$.*

Again, w is not unique. The interactions u and v impose causal constraints on $|w|$, and w may be chosen as any ternary interaction respecting those. As in the alternating case there is a mirror lemma, and associativity follows as in Proposition B.9.

Proposition C.3. *We have $(\delta \circledast \tau) \circledast \sigma = \delta \circledast (\tau \circledast \sigma)$.*

Associativity works for prestrategies, *i.e.* it does not rely on *receptive* and *courteous*.

C.1.3. *Copycat.* Non-alternating strategies are intended to model *asynchronous* concurrency – accordingly, the non-alternating copycat, still written $\mathfrak{c}_A : A \multimap A$, is an asynchronous forwarder. We first describe the configurations on $A \vdash A$ reached by copycat: those are $x \parallel y$ for $x, y \in \mathcal{C}(A)$, such that $y \sqsubseteq x$, where \sqsubseteq is the “Scott order” [Win13] defined as

$$y \sqsubseteq x \iff y \supseteq^- x \cap y \subseteq^+ x$$

with polarities taken on A . Whenever \mathfrak{c}_A receives a negative event, it forwards it to the other side, but asynchronously: y may contain negative events not yet forwarded to x , and x may contain positive events (*i.e.* negative for $A \vdash A$) not yet forwarded to y .

But copycat plays on $A \multimap A$, not on $A \vdash A$. Remember from Section 2.1.3 that there is

$$\chi_{A,B} : (A \multimap B) \rightarrow (A \vdash B)$$

satisfying the axioms of a map of event structures. We may set:

Definition C.4. For A any --arena, we set \mathfrak{c}_A to comprise all $s \in \mathcal{C}\text{-Plays}(A \multimap A)$ s.t.:

- balanced:* for all $1 \leq i \leq n$, writing $\chi_{A,A} |s_1 \dots s_n| = x \parallel y$, we have $y \sqsubseteq x$,
- well-linked:* for all $[(l, a)]_{a'}$, we have $a = a'$,

where $s = s_1 \dots s_n$.

It is direct that copycat is *receptive* and *courteous*. In fact, it turns out that *receptive* and *courteous* are *exactly* the required conditions for copycat to be neutral for composition:

Proposition C.5. *Consider A and B two --arenas, and $\sigma : A \multimap B$ a prestrategy.*

Then, σ is a strategy if and only if $\mathfrak{c}_B \circledast \sigma \circledast \mathfrak{c}_A = \sigma$.

Sketch. *If.* It is direct that all non-alternating prestrategies $\mathfrak{c}_B \circledast \sigma \circledast \mathfrak{c}_A$ are *receptive* and *courteous*. *Only if.* Considering $w \in \mathfrak{c}_B \circledast \sigma \circledast \mathfrak{c}_A$, using that σ is *receptive* and *courteous*, w can be transformed by permuting into w' with the same outer restriction, but where all moves are immediately forwarded by copycat. It follows that the outer restriction is in σ . \square

This was noticed in [GM08], and also holds for causal strategies [RW11]. We deduce:

Corollary C.6. *The $-$ -arenas and non-alternating strategies form a category $\mathcal{C}\text{-Strat}$.*

Proof. Follows from Propositions C.3 and C.5. \square

C.1.4. *Further structure.* As for $\downarrow\uparrow\text{-Strat}$, we have $-$ with the same constructions:

Proposition C.7. *$\mathcal{C}\text{-Strat}$ is a cartesian symmetric monoidal closed category.*

We could easily adapt the developments of Section 6.1 to show that there is a subcategory of *sequential* non-alternating strategies, which maps functorially to $\downarrow\uparrow\text{-Strat}$ (without uniformity). However, our attempts to endow $\mathcal{C}\text{-Strat}$ with symmetry failed, because of the lack of *unique witness* in compositions. It seems it could be done by using a different approach to symmetry, namely Melliès' group-theoretic formulation of uniformity [Mel03]; this is however left for future work. We conclude this section with:

Proposition C.8. *There is a symmetric monoidal closed category with products $\mathcal{C}\text{-Strat}^{\text{wb}}$, with objects $-$ -arenas, and morphisms well-bracketed non-alternating strategies on $A \multimap B$.*

Proof. It remains to prove that copycat strategies are well-bracketed and that well-bracketing is preserved by operations on strategies, which is a routine verification. \square

APPENDIX D. THIN CONCURRENT GAMES

In this section, we give some proof for thin concurrent games. Notably, we detail the proofs for the characterisations of configurations of interaction and composition used in this paper: they do not appear in [CCW19], as they were elaborated more recently.

D.1. Representable functions. We investigate the functions between domains of configurations representable as maps of event structures.

D.1.1. *Representable functions between configurations.* We will be interested in those functions between configurations arising from maps of event structures:

Definition D.1. For A, B es, a function $f : \mathcal{C}(A) \rightarrow \mathcal{C}(B)$ is **representable** if there is a (necessarily unique) map of event structure $\hat{f} : A \rightarrow B$ s.t. for all $x \in \mathcal{C}(A)$, $\hat{f}(x) = f(x)$.

Proof. For uniqueness, if $f_1, f_2 : A \rightarrow B$ have the same image for configurations, we have

$$f_1([a]_A) \xrightarrow{f_1(a)} \subset f_1([a]_A), \quad f_2([a]_A) \xrightarrow{f_2(a)} \subset f_2([a]_A),$$

where $[a]_A = \{a' \in |A| \mid a' <_A a\}$.

Since $f_1([a]_A) = f_2([a]_A)$ and $f_1([a]_A) = f_2([a]_A)$ this entails $f_1(a) = f_2(a)$. \square

We shall use the following characterisation of representable functions:

Lemma D.2. *For A, B es, a function $f : \mathcal{C}(A) \rightarrow \mathcal{C}(B)$ is representable iff it:*

- preserves the empty set: $f(\emptyset) = \emptyset$
- preserves covering: for $x, y \in \mathcal{C}(A)$, if $x \multimap y$, then $f(x) \multimap f(y)$
- preserves unions: for $x, y \in \mathcal{C}(A)$, if $x \cup y \in \mathcal{C}(A)$ then $f(x \cup y) = f(x) \cup f(y)$.

Proof. If. For any $a \in |A|$, write $[a]_A = [a]_A \setminus \{a\}$. It is always a configuration of A , and $[a]_A \prec [a]_A$. Because f preserves covering, there is $b \in |B|$ such that

$$f([a]_A) \xrightarrow{b} f([a]_A),$$

i.e. $f([a]_A) = f([a]_A) \cup \{b\}$, write $b = \hat{f}(a)$.

This defines a function on events. Next, we show that its direct image of configurations agrees with f . We first claim that for all $x, y \in \mathcal{C}(A)$ such that $x \prec y$ and $y = x \cup \{a\}$, we have $f(y) = f(x) \cup \{\hat{f}(a)\}$ as well. To prove this, consider the diagram

$$\begin{array}{ccccccc} [a]_A & \prec & x_1 & \prec & \dots & \prec & x_{n-1} & \prec & x \\ \downarrow & & \downarrow & & & & \downarrow & & \downarrow \\ [a]_A & \prec & y_1 & \prec & \dots & \prec & y_{n-1} & \prec & y \end{array}$$

where, for each $1 \leq i < n$, $y_i = x_i \cup \{a\}$. Because f preserves covering and unions, f preserves covering squares, hence by straightforward induction, for all $1 \leq i < n$,

$$f(x_i) \xrightarrow{\hat{f}(a)} f(y_i),$$

so $f(y) = f(x) \cup \{\hat{f}(a)\}$ as required.

Now, we can prove that for all $x \in \mathcal{C}(A)$, we have $f(x) = \hat{f}(x)$. Indeed, there is

$$\emptyset \xrightarrow{a_1} \dots \xrightarrow{a_n} x$$

a covering chain, which as f preserves the empty set and by the property just above, entails

$$\emptyset \xrightarrow{\hat{f}(a_1)} \dots \xrightarrow{\hat{f}(a_n)} f(x),$$

as well, showing that $f(x) = \hat{f}(x)$ as needed. From this point it is trivial that \hat{f} is a map of ess: preservation of configurations follows from $f : \mathcal{C}(A) \rightarrow \mathcal{C}(B)$, and local injectivity from the preservation of covering. *Only if.* Immediate verification. \square

D.1.2. *Representable functions with symmetry.* We extend this construction with symmetry.

Definition D.3. Let A, B be ess. A function $f : \mathcal{C}(A) \rightarrow \mathcal{C}(B)$ is **representable** if there is a (necessarily unique) map of ess $\hat{f} : A \rightarrow B$ such that for all $x \in \mathcal{C}(A)$, $\hat{f}(x) = f(x)$.

As above, we give a characterization of representable functions between configurations.

Lemma D.4. *Let A, B be ess. Then, a function $f : \mathcal{C}(A) \rightarrow \mathcal{C}(B)$ is **representable** iff it:*

- preserves the empty set: $f(\emptyset) = \emptyset$
- preserves covering: for $x, y \in \mathcal{C}(A)$, if $x \prec y$, then $f(x) \prec f(y)$
- preserves unions: for $x, y \in \mathcal{C}(A)$, if $x \cup y \in \mathcal{C}(A)$ then $f(x \cup y) = f(x) \cup f(y)$.
- preserves symmetry: there is a (necessarily unique) monotone function $\tilde{f} : \mathcal{S}(A) \rightarrow \mathcal{S}(B)$ such that $\text{dom} \circ \tilde{f} = f \circ \text{dom}$ and $\text{cod} \circ \tilde{f} = f \circ \text{cod}$.

Proof. If. If f satisfies the first three axioms, it is representable without symmetry. By Lemma D.2 there is a map of event structures $\hat{f} : A \rightarrow B$ s.t. for all $x \in \mathcal{C}(A)$, $\hat{f}(x) = f(x)$.

Now, assuming that f preserves symmetry, we show that actually we must have

$$\tilde{f}(\theta) = \{(\hat{f}(a), \hat{f}(a')) \mid (a, a') \in \theta\},$$

i.e. $\hat{f}(\theta)$, for all $\theta \in \mathcal{S}(A)$. Indeed, consider a covering chain for θ , i.e. a sequence in $\mathcal{S}(A)$:

$$\emptyset = \theta_0 \xrightarrow{(a_1, a'_1)} \theta_1 \xrightarrow{(a_2, a'_2)} \dots \xrightarrow{(a_n, a'_n)} \theta_n = \theta.$$

We show by induction on $0 \leq i \leq n$ that $\tilde{f}(\theta_i) = \hat{f}(\theta_i)$. First, $\text{dom} \circ \tilde{f}(\emptyset) = f(\emptyset) = \emptyset$, so $\tilde{f}(\emptyset) = \emptyset = \hat{f}(\emptyset)$. For $0 \leq i \leq n$, by IH, $\tilde{f}(\theta_i) = \hat{f}(\theta_i)$. We then have:

$$\begin{aligned} \text{dom}(\tilde{f}(\theta_{i+1})) &= f(\text{dom}(\theta_{i+1})) \\ &= f(\text{dom}(\theta_i) \cup \{a_{i+1}\}) \\ &= f(\text{dom}(\theta_i)) \cup \{\hat{f}(a_{i+1})\}, \end{aligned}$$

and the symmetric reasoning shows $\text{cod}(\tilde{f}(\theta_{i+1})) = \text{cod}(\tilde{f}(\theta_i)) \cup \{\hat{f}(a'_{i+1})\}$ as well. But finally, we also have $\tilde{f}(\theta_i) \subseteq \tilde{f}(\theta_{i+1})$ since \tilde{f} is monotone. So we must have

$$\begin{aligned} \tilde{f}(\theta_{i+1}) &= \tilde{f}(\theta_i) \cup \{(\hat{f}(a_{i+1}), \hat{f}(a'_{i+1}))\} \\ &= \hat{f}(\theta_i) \cup \{(\hat{f}(a_{i+1}), \hat{f}(a'_{i+1}))\} \\ &= \hat{f}(\theta_{i+1}), \end{aligned}$$

as required. Now this exactly means that for all $\theta \in \mathcal{S}(A)$ we have $\hat{f}(\theta) \in \mathcal{S}(B)$, and thus \hat{f} is a map of event structures with symmetry. *Only if.* Obvious. \square

D.2. Interaction and Composition. Fix $\sigma : A \vdash B$ and $\tau : B \vdash C$ (pre)strategies.

D.2.1. Interaction. First, we characterise interaction in terms of its configurations.

We start by recalling that by Lemma 3.12 of [CCW19], there is a pullback

$$\begin{array}{ccc} & \tau \circledast \sigma & \\ \Pi_\sigma \swarrow & \downarrow & \searrow \Pi_\tau \\ \sigma \parallel C & & A \parallel \tau \\ \partial_\sigma \parallel C \searrow & & \swarrow A \parallel \partial_\tau \\ & A \parallel B \parallel C & \end{array}$$

in the category of event structures with symmetry, the *interaction pullback*. First, we have:

Lemma D.5. *For any $x \in \mathcal{C}(\tau \circledast \sigma)$, writing*

$$\Pi_\sigma x = x^\sigma \parallel x_C^\tau \in \mathcal{C}(\sigma \parallel C) \quad \Pi_\tau x = x_A^\sigma \parallel x^\tau \in \mathcal{C}(A \parallel \tau),$$

then we have $x^\sigma \in \mathcal{C}(\sigma)$ and $x^\tau \in \mathcal{C}(\tau)$ causally compatible.

Proof. First, x^σ and x^τ are matching. For causal compatibility, we consider

$$\varphi_{x^\sigma, x^\tau} : x^\sigma \parallel x_C^\tau \xrightarrow{\partial_\sigma \parallel x_C^\tau} x_A^\sigma \parallel x_B \parallel x_C^\tau \xrightarrow{x_A^\sigma \parallel \partial_\tau^{-1}} x_A^\sigma \parallel x^\tau$$

from Definition 4.5, with $(m, n) \triangleleft (m', n')$ iff $m <_{\sigma \parallel C} m'$ or $n <_{A \parallel \tau} n'$ which we must prove acyclic. Now, the function which to $c \in x$ associates $(\Pi_\sigma(c), \Pi_\tau(c))$ is a bijection

$$\psi : x \simeq \varphi_{x^\sigma, x^\tau},$$

and we now claim that for all $c, c' \in x$, if $\psi(c) \triangleleft \psi(c')$, then $c <_{\tau \circledast \sigma} c'$. Indeed, say $c = (m, n)$ and $c' = (m', n')$, with *w.l.o.g.* $m <_{\sigma \parallel C} m'$. Of course then, $m = \Pi_\sigma(c)$ and $m' = \Pi_\sigma(c')$.

Now, we use that Π_σ is a map of event structures, and those always locally reflect causality (Lemma A.2). Therefore, we must have $c <_{\tau \otimes \sigma} c'$ as required.

Therefore a cycle for $<$ in $\varphi_{x^\sigma, x^\tau}$ would induce a cycle for $<_{\tau \otimes \sigma}$ in x , contradiction. \square

Proposition D.6. *There is a pre-interaction $\tau \otimes \sigma$, unique up to iso, such that there are*

$$\begin{aligned} (- \otimes -) & : \{(x^\tau, x^\sigma) \in \mathcal{C}(\tau) \times \mathcal{C}(\sigma) \mid x^\sigma, x^\tau \text{ causally compatible}\} \simeq \mathcal{C}(\tau \otimes \sigma) \\ (- \otimes -) & : \{(\theta^\tau, \theta^\sigma) \in \mathcal{S}(\tau) \times \mathcal{S}(\sigma) \mid \theta^\sigma, \theta^\tau \text{ causally compatible}\} \simeq \mathcal{S}(\tau \otimes \sigma) \end{aligned}$$

order-isomorphisms commuting with dom and cod, and satisfying

$$\partial_{\tau \otimes \sigma}(\theta^\tau \otimes \theta^\sigma) = \theta_A^\sigma \parallel \theta_B \parallel \theta_C^\tau$$

for all $\theta^\sigma \in \mathcal{S}(\sigma)$ and $\theta^\tau \in \mathcal{S}(\tau)$ causally compatible.

Proof. Existence. We must provide the two order-isomorphisms announced. Take $x^\sigma \in \mathcal{C}(\sigma)$ and $x^\tau \in \mathcal{C}(\tau)$ causally compatible. By Definition 4.5, the bijection

$$\varphi_{x^\sigma, x^\tau} : x^\sigma \parallel x_C^\tau \xrightarrow{\partial_\sigma \parallel x_C^\tau} x_A^\sigma \parallel x_B \parallel x_C^\tau \xrightarrow{x_A^\sigma \parallel \partial_\tau^{-1}} x_A^\sigma \parallel x^\tau$$

is secured, *i.e.* the relation $(m, n) < (m', n') \Leftrightarrow m \leq_{\sigma \parallel C} m' \vee n \leq_{A \parallel \tau} n'$ defined on the graph of $\varphi_{x^\sigma, x^\tau}$ is acyclic, so its reflexive transitive closure \leq_{x^σ, x^τ} is a partial order. This turns $\varphi_{x^\sigma, x^\tau}$ into an event structure, and in fact an ess with identity symmetries. Moreover, there are obvious maps of ess

$$\pi_\sigma : \varphi_{x^\sigma, x^\tau} \rightarrow \sigma \parallel C \quad \pi_\tau : \varphi_{x^\sigma, x^\tau} \rightarrow A \parallel \tau$$

commuting with display maps, so by the universal property, $\langle \pi_\sigma, \pi_\tau \rangle : \varphi_{x^\sigma, x^\tau} \rightarrow \tau \otimes \sigma$ and

$$x^\tau \otimes x^\sigma = \langle \pi_\sigma, \pi_\tau \rangle(|\varphi_{x^\sigma, x^\tau}|) \in \mathcal{C}(\tau \otimes \sigma)$$

concludes the definition of the action of $(- \otimes -)$ on causally compatible pairs. Reciprocally, if $x \in \mathcal{C}(\tau \otimes \sigma)$ then its projections yield $\Pi_\sigma(x) = x^\sigma \parallel x_C^\tau$ and $\Pi_\tau(x) = x_A^\sigma \parallel x^\tau$, and by Lemma D.5, x^σ and x^τ are causally compatible. Finally, for x^σ and x^τ causally compatible,

$$\Pi_\sigma(x^\tau \otimes x^\sigma) = x^\sigma \parallel x_C^\tau \quad \Pi_\tau(x^\tau \otimes x^\sigma) = x_A^\sigma \parallel x^\tau$$

by construction and if $x \in \mathcal{C}(\tau \otimes \sigma)$, $x = x^\tau \otimes x^\sigma$ by universal property of the pullback. The projections are monotone, and the monotonicity of $(- \otimes -)$ follows from the universal property. For symmetries, causal compatibility of $\theta^\sigma \in \mathcal{S}(\sigma)$ and $\theta^\tau \in \mathcal{S}(\tau)$ amounts to

$$\begin{aligned} x^\sigma \parallel x_C^\tau & \xrightarrow{\Pi_\sigma^{-1}} x^\tau \otimes x^\sigma \xrightarrow{\Pi_\tau} x_A^\sigma \parallel x^\tau \\ \theta^\sigma \parallel \theta_C^\tau & \parallel \parallel \parallel \theta_A^\sigma \parallel \theta^\tau \\ y^\sigma \parallel y_C^\tau & \xrightarrow{\Pi_\sigma} y^\tau \otimes y^\sigma \xrightarrow{\Pi_\tau} y_A^\sigma \parallel y^\tau \end{aligned}$$

commuting, inducing $\theta^\otimes : x^\tau \otimes x^\sigma \simeq y^\tau \otimes y^\sigma$. But symmetries on $\tau \otimes \sigma$ are precisely those bijections $x^\tau \otimes x^\sigma \simeq y^\tau \otimes y^\sigma$ projecting to $\mathcal{S}(S \parallel C)$ and $\mathcal{S}(A \parallel T)$ as above, so $\theta^\otimes \in \mathcal{S}(\tau \otimes \sigma)$. Reciprocally, $\theta^\otimes \in \mathcal{S}(\tau \otimes \sigma)$ induces $\theta^\sigma \in \mathcal{S}(\sigma)$ and $\theta^\tau \in \mathcal{S}(\tau)$ by projections. That this yields an order-iso compatible with dom and cod is direct.

Uniqueness. Two event structures with symmetry satisfying the hypotheses obviously have isomorphic domains of configurations (and isomorphic domains of symmetries, in a compatible manner). But such order-isomorphisms between domains of configurations and symmetries are automatically representable in the sense of Definition D.3. Therefore, by Lemma D.4 the isomorphisms are generated by isomorphisms of event structures with

symmetry, as required. Preservation of display maps holds for configurations by hypothesis; and by uniqueness in Lemma D.4, two maps of ess with the same action on configurations must be equal. Therefore, the isomorphism commutes with display maps and is an isomorphism of pre-interactions as required. \square

D.2.2. Composition. We aim to prove Propositions 4.10 and 4.12, exploiting Proposition 4.8, along with a few extra lemmas. First, we notice a connection between *minimality* of causally compatible pairs and that maximal events of interactions are visible.

Lemma D.7. *For $x^\sigma \in \mathcal{C}(\sigma)$ and $x^\tau \in \mathcal{C}(\tau)$ causally compatible, they are minimal causally compatible iff the maximal events of $x^\tau \otimes x^\sigma$ are **visible**, e.g. occur in A or C .*

Proof. If. Consider $y^\sigma \in \mathcal{C}(\sigma)$ and $y^\tau \in \mathcal{C}(\tau)$ causally compatible such that $y^\sigma \subseteq x^\sigma$, $y^\tau \subseteq x^\tau$, while $x_A^\sigma = y_A^\sigma$ and $x_C^\tau = y_C^\tau$. Assume, seeking a contradiction, that we have $y_B \subset x_B$ a strict inclusion. Necessarily, we have $y^\tau \otimes y^\sigma \subset x^\tau \otimes x^\sigma$. Take $m \in (y^\tau \otimes y^\sigma) \setminus (x^\tau \otimes x^\sigma)$, w.l.o.g. we can assume that m is maximal for $\leq_{\tau \otimes \sigma}$ in $x^\tau \otimes x^\sigma$. By hypothesis, m occurs in A or C . But this immediately contradicts the hypothesis that $x_A^\sigma = y_A^\sigma$ and $x_C^\tau = y_C^\tau$.

Only if. Assume x^σ, x^τ are minimal causally compatible, and take $m \in x^\tau \otimes x^\sigma$ maximal. Seeking a contradiction, assume that m occurs in B . So, projecting

$$\Pi_\sigma(x^\tau \otimes x^\sigma) = x^\sigma \parallel x_C^\tau, \quad \Pi_\tau(x^\tau \otimes x^\sigma) = x_A^\sigma \parallel x^\tau,$$

we have $\Pi_\sigma m = (1, s)$ with $s \in x^\sigma$ and $\Pi_\tau m = (2, t)$ with $t \in x^\tau$. As maps of event structures locally reflect causality (Lemma A.2), s is maximal in x^σ and t is maximal in x^τ . Hence, $y^\sigma = x^\sigma \setminus \{s\} \in \mathcal{C}(\sigma)$ and $y^\tau = x^\tau \setminus \{t\} \in \mathcal{C}(\tau)$. By construction they are causally compatible with the same projections to A and C , contradicting minimality of x^σ and x^τ . \square

We write $\mathcal{C}^v(\tau \otimes \sigma)$ for the configurations whose maximal events are visible and likewise for symmetries. Then, the lemma above means that we can refine Proposition 4.8 to:

Proposition D.8. *The order-isomorphisms of Proposition 4.8 restrict to*

$$\begin{aligned} (- \otimes -) &: \{(x^\tau, x^\sigma) \in \mathcal{C}(\tau) \times \mathcal{C}(\sigma) \mid x^\sigma \text{ and } x^\tau \text{ minimal causally compatible}\} \simeq \mathcal{C}^v(\tau \otimes \sigma) \\ (- \otimes -) &: \{(\theta^\tau, \theta^\sigma) \in \mathcal{S}(\tau) \times \mathcal{S}(\sigma) \mid \theta^\sigma \text{ and } \theta^\tau \text{ minimal causally compatible}\} \simeq \mathcal{S}^v(\tau \otimes \sigma) \end{aligned}$$

Now, it remains to link configurations of $\tau \odot \sigma$ with configurations of $\tau \otimes \sigma$ with visible maximal events, and likewise for symmetries. First, we fix a few notations. Any configuration $x \in \mathcal{C}(\tau \otimes \sigma)$ yields a configuration of the composition, its **hiding**, defined as $x \downarrow = x \cap |\tau \odot \sigma| \in \mathcal{C}(\tau \odot \sigma)$. Reciprocally, if $x \in \mathcal{C}(\tau \odot \sigma)$ is a configuration of the composition, its **witness** is defined as

$$[x]_{\tau \otimes \sigma} = \{m \in |\tau \otimes \sigma| \mid \exists n \in x, m \leq_{\tau \otimes \sigma} n\} \in \mathcal{C}(\tau \otimes \sigma).$$

The next point we make, is that interactions with visible maximal events are exactly those arising as witnesses of configurations of the composition. More precisely, we have:

Lemma D.9. *There are order-isomorphisms compatible with dom, cod and display maps:*

$$\begin{aligned} (-) \downarrow &: \mathcal{C}^v(\tau \otimes \sigma) \simeq \mathcal{C}(\tau \odot \sigma) &: [-]_{\tau \otimes \sigma} \\ (-) \downarrow &: \mathcal{S}^v(\tau \otimes \sigma) \simeq \mathcal{S}(\tau \odot \sigma) &: [-]_{\tau \otimes \sigma} \end{aligned}$$

Proof. For configurations, to $x \in \mathcal{C}^v(\tau \circledast \sigma)$ we associate its hiding $x \downarrow \in \mathcal{C}(\tau \odot \sigma)$. Reciprocally, to $x \in \mathcal{C}(\tau \odot \sigma)$, we associate its witness $[x]_{\tau \circledast \sigma} \in \mathcal{C}^v(\tau \circledast \sigma)$. Those operations preserve inclusion, and it is an elementary verification that they are inverses.

For symmetries, any $\theta : x \cong_{\tau \circledast \sigma} y$ must preserve visible events, so it induces by *hiding*

$$\theta \downarrow : x \downarrow \cong_{\tau \odot \sigma} y \downarrow$$

a symmetry on the composition; and hiding preserves inclusion. Reciprocally, if $\theta : x \cong_{\tau \odot \sigma} y$ then by definition there is $\theta \subseteq \theta' : x' \cong_{\tau \circledast \sigma} y'$. Necessarily, $[x]_{\tau \circledast \sigma} \subseteq x'$ and $[y]_{\tau \circledast \sigma} \subseteq y'$. Since θ' is an order-iso, by *restriction* we may assume *w.l.o.g.* $\theta' : [x]_{\tau \circledast \sigma} \cong_{\tau \circledast \sigma} [y]_{\tau \circledast \sigma}$.

But the witness θ' is *unique*: by Lemma 3.33 of [CCW19], if $\theta'' : [x]_{\tau \circledast \sigma} \cong_{\tau \circledast \sigma} [y]_{\tau \circledast \sigma}$ is such that $\theta'' \downarrow = \theta' \downarrow = \theta$, then $\theta' = \theta''$. So to $\theta : x \cong_{\tau \odot \sigma} y$ we associate this unique $\theta' : [x]_{\tau \circledast \sigma} \cong_{\tau \circledast \sigma} [y]_{\tau \circledast \sigma}$. Monotonicity and that these operations are inverses also follow immediately from uniqueness of the witness symmetry, *i.e.* Lemma 3.33 of [CCW19]. \square

Composing the order-isomorphisms of Proposition D.8 and Lemma D.9, we get:

Proposition D.10. *Consider $\sigma : A \vdash B$, and $\tau : B \vdash C$ causal strategies.*

There is a causal strategy $\tau \odot \sigma$, unique up to iso, s.t. there are order-isos:

$$\begin{aligned} (- \odot -) & : \{(x^\tau, x^\sigma) \in \mathcal{C}(\tau) \times \mathcal{C}(\sigma) \mid x^\sigma, x^\tau \text{ minimal causally compatible}\} \simeq \mathcal{C}(\tau \odot \sigma) \\ (- \odot -) & : \{(\theta^\tau, \theta^\sigma) \in \mathcal{S}(\tau) \times \mathcal{S}(\sigma) \mid \theta^\sigma, \theta^\tau \text{ minimal causally compatible}\} \simeq \mathcal{S}(\tau \odot \sigma) \end{aligned}$$

commuting with dom and cod; s.t., for $\theta^\sigma \in \mathcal{S}(\sigma), \theta^\tau \in \mathcal{S}(\tau)$ minimal causally compatible,

$$\partial_{\tau \odot \sigma}(\theta^\tau \odot \theta^\sigma) = \theta_A^\sigma \parallel \theta_C^\tau.$$

Proof. Existence. Direct by Proposition D.8 and Lemma D.9.

Uniqueness. As in the proof of Proposition 4.8. \square

This means that any configuration $x \in \mathcal{C}(\tau \odot \sigma)$ may be written uniquely as $x^\tau \odot x^\sigma$ for $x^\sigma \in \mathcal{C}(\sigma)$ and $x^\tau \in \mathcal{C}(\tau)$ minimal causally compatible – note that from the construction of the order-isomorphism in the proposition above, we then have $x^\tau \circledast x^\sigma = [x^\tau \odot x^\sigma]_{\tau \circledast \sigma}$.

D.2.3. *+covered case.* Finally, it remains to prove Proposition 4.12.

In essence, Proposition 4.12 is a specialization of Proposition 4.10 to strategies. Consider therefore from now on that $\sigma : A \vdash B$ and $\tau : B \vdash C$ are *strategies*. First of all, we show that for *+covered* configurations we can omit the minimality assumption.

Lemma D.11. *Consider $x^\sigma \in \mathcal{C}^+(\sigma)$ and $x^\tau \in \mathcal{C}^+(\tau)$ causally compatible *+covered*.*

Then, x^σ and x^τ are minimal causally compatible.

Proof. Seeking a contradiction, assume y^σ and y^τ are causally compatible such that $y^\tau \circledast y^\sigma \subseteq x^\tau \circledast x^\sigma$ with $x_A^\sigma = y_A^\sigma$ and $x_C^\tau = y_C^\tau$. Without loss of generality, consider $m \in (x^\tau \circledast x^\sigma) \setminus (y^\tau \circledast y^\sigma)$ maximal for $\leq_{\tau \circledast \sigma}$. By hypothesis, m occurs in B . Therefore, projecting

$$\Pi_\sigma(m) = (1, s) \quad \Pi_\tau(m) = (2, t),$$

maximality of m in $x^\tau \circledast x^\sigma$ entails via Lemma A.2 that s is maximal in x^σ and t maximal in x^τ . But necessarily, s and t has dual polarities: *w.l.o.g.* say that $\text{pol}_\sigma(s) = +$ and $\text{pol}_\tau(t) = -$. So, t is negative maximal in x^τ , contradicting that $x^\tau \in \mathcal{C}^+(\tau)$ is *+covered*. \square

So in a synchronization between *+covered* configurations, the maximal events are visible as if they are synchronized, they will be both maximal and negative for one of the two players. Resulting configurations of the composition are automatically *+covered*:

Lemma D.12. For $x^\sigma \in \mathcal{C}^+(\sigma)$, $x^\tau \in \mathcal{C}^+(\tau)$ causally compatible, $x^\tau \odot x^\sigma \in \mathcal{C}^+(\tau \odot \sigma)$.

Proof. Consider $m \in x^\tau \odot x^\sigma$ maximal. This means that m is also maximal in $x^\tau \otimes x^\sigma = [x^\tau \otimes x^\sigma]_{\tau \otimes \sigma}$. Necessarily, m occurs in A or C , *w.l.o.g.* assume it occurs in C . Then, $\Pi_\tau(m)$ has the form $(2, t)$ where using Lemma A.2, necessarily t is maximal (for \leq_τ) in x^τ . But then, since x^τ is $+$ -covered, t is positive – hence, m is positive as well. \square

So causally compatible $+$ -covered $x^\sigma \in \mathcal{C}^+(\sigma)$ and $x^\tau \in \mathcal{C}^+(\tau)$ are minimal, and their composition yields $x^\tau \odot x^\sigma \in \mathcal{C}^+(\tau \odot \sigma)$ $+$ -covered. We prove the converse:

Lemma D.13. Consider $x^\sigma \in \mathcal{C}(\sigma)$ and $x^\tau \in \mathcal{C}(\tau)$ minimal causally compatible.

If $x^\tau \odot x^\sigma \in \mathcal{C}^+(\tau \odot \sigma)$ is $+$ -covered, so are $x^\sigma \in \mathcal{C}^+(\sigma)$ and $x^\tau \in \mathcal{C}^+(\tau)$.

Proof. Consider $s \in x^\sigma$ maximal. Necessarily, there is a unique $m \in x^\tau \otimes x^\sigma$ such that $\Pi_\sigma(m) = (1, s)$. Assume first that m is maximal in $x^\tau \otimes x^\sigma$. As $x^\tau \otimes x^\sigma = [x^\tau \odot x^\sigma]_{\tau \otimes \sigma}$, if m is maximal it must be visible and maximal in $x^\tau \odot x^\sigma$. Therefore, it is positive by hypothesis, and s is positive.

Otherwise, assume m is not maximal, so there is some $m \rightarrow_{\tau \otimes \sigma} n$. By Lemma 4.21,

$$\Pi_\sigma(m) \rightarrow_{\sigma \parallel C} \Pi_\sigma(n) \quad \text{or} \quad \Pi_\tau(m) \rightarrow_{A \parallel \tau} \Pi_\tau(n).$$

If this is the former, then there is $s \rightarrow_\sigma s'$ with $s' \in x^\sigma$, absurd by maximality of s . If this is the latter, then two cases arise. If m occurs in A , then $\Pi_\tau(m) = (1, a)$ and $\Pi_\tau(n) = (1, a')$ with $a \rightarrow_A a'$. Likewise, $\Pi_\sigma(m) = (1, s)$ and $\Pi_\sigma(n) = (1, s')$. But then by Lemma A.2, we must have $s <_\sigma s'$ contradicting again the maximality of s . Finally, if m occurs in B , then $\Pi_\tau(m) = (2, t)$ and $\Pi_\tau(n) = (2, t')$ with $t \rightarrow_\tau t'$. We split cases one last time, depending on the polarity of t in τ . If t is negative, then s is positive in σ as required. Otherwise, by courtesy of τ we must have $\partial_\tau(t) \rightarrow_{B \vdash C} \partial_\tau(t')$. In particular, t' must also occur in B and we must have $\Pi_\sigma(n) = (1, s')$ for $s' \in x^\sigma$, with moreover $\partial_\sigma(s) \rightarrow_{A \vdash B} \partial_\sigma(s')$. Therefore, again by Lemma A.2, we must have $s <_\sigma s'$ contradicting the maximality of s .

The symmetric reasoning shows that any $t \in x^\tau$ maximal in x^τ is positive. \square

We are almost in position to prove Proposition 4.12 – the only missing piece is *uniqueness*:

Lemma D.14. Consider σ, τ : A two causal strategies. Assume there are

$$\psi : \mathcal{C}^+(\sigma) \simeq \mathcal{C}^+(\tau) \quad \psi : \mathcal{S}^+(\sigma) \simeq \mathcal{S}^+(\tau)$$

order-isomorphisms compatible with dom , cod , and display maps.

Then, σ and τ are isomorphic.

Proof. We extend ψ to all configurations and all symmetries, and conclude via Lemma D.4.

Let $x \in \mathcal{C}(\sigma)$. Consider $x^+ \in \mathcal{C}(\sigma)$ minimal such that $x^+ \subseteq^- x$. Necessarily, $x^+ \in \mathcal{C}^+(\sigma)$, so we may take $\psi(x^+) \in \mathcal{C}^+(\tau)$. Now, since ψ is compatible with display maps,

$$\partial_\tau(\psi(x^+)) \subseteq^- \partial_\sigma(x),$$

therefore by receptivity and courtesy (see Lemma 3.13 from [CCRW17]), there is a unique $y \in \mathcal{C}(\tau)$ such that $\psi(x^+) \subseteq^- y$ and $\partial_\tau(y) = \partial_\sigma(x)$; we set $\psi(x) = y$.

We must show this extended ψ preserves inclusion; we show it preserves covering, and distinguish the positive and negative cases. First, consider configurations in $\mathcal{C}(\sigma)$:

$$x \overset{s^+}{\text{---}} y,$$

which means $x^+ \subseteq^- z \overset{s^+}{\dashv\!\!-\!\!-} y^+$. Now, by hypothesis $\psi(x^+) \subseteq \psi(y^+)$, so $\psi(x^+) \subseteq \psi(y)$. Moreover, this inclusion must contain exactly one positive event, write it $t^+ \in \psi(y) \setminus \psi(x^+) -$ necessarily, $\partial_\tau(t) = \partial_\sigma(s)$. Now, notice $\partial_\tau(t)$ is maximal in $\partial_\tau(\psi(y))$: indeed,

$$\partial_\tau(\psi(x)) = \partial_\sigma(x) \overset{\partial_\sigma(s)}{\dashv\!\!-\!\!-} \partial_\sigma(y) = \partial_\tau(\psi(y))$$

by hypothesis. So, by courtesy, t is maximal in $\psi(y)$ as well. So, we have $z = \psi(y) \setminus \{t\} \in \mathcal{C}(\tau)$. Moreover, $\partial_\tau(z) = \partial_\tau(\psi(y)) \setminus \{\partial_\tau(t)\} = \partial_\sigma(x)$. So, $\psi(x^+) \subseteq^- z$ with $\partial_\tau(z) = \partial_\sigma(x)$, therefore $z = \psi(x)$; and $\psi(x) \dashv\!\!-\!\!- \psi(y)$ as required. Considering now a negative extension, *i.e.*

$$x \overset{s^-}{\dashv\!\!-\!\!-} y,$$

it follows that $x^+ = y^+$ and $\psi(x) \dashv\!\!-\!\!- \psi(y)$ is immediate by receptivity. Altogether we have

$$\psi : \mathcal{C}(\sigma) \rightarrow \mathcal{C}(\tau)$$

compatible with display maps and preserving inclusion. Likewise we construct $\psi^{-1} : \mathcal{C}(\tau) \rightarrow \mathcal{C}(\sigma)$ preserving inclusion from its action on $+$ -covered configurations. That they are inverses is immediate from ψ being a bijection between $+$ -covered configurations, and receptivity.

Now, we consider the action of ψ on symmetries. If $\theta \in \mathcal{S}(\sigma)$, consider $\theta^+ \in \mathcal{C}(\sigma)$ minimal s.t. $\theta^+ \subseteq^- \theta$ – recall that as a symmetry, θ is an order-iso preserving polarities, so this is well-defined. Now, $\theta^+ \in \mathcal{S}^+(\sigma)$, so that we may take $\psi(\theta^+) \in \mathcal{S}^+(\tau)$ as for configurations. Again, since ψ is compatible with display maps, we have

$$\partial_\tau(\psi(\theta^+)) \subseteq^- \partial_\sigma(\theta).$$

By receptivity and courtesy of τ , there are unique extensions of $\text{dom}(\psi(\theta^+))$ and $\text{cod}(\psi(\theta^+))$ to $\psi(\text{dom}(\theta))$ and $\psi(\text{cod}(\theta))$, projecting via ∂_τ to $\text{dom}(\partial_\sigma(\theta))$ and $\text{cod}(\partial_\sigma(\theta))$ respectively. We get $\psi(\theta^+) \subseteq^- \Omega \in \mathcal{S}(\tau)$ by iterating \sim -receptivity for τ , such that $\partial_\tau(\Omega) = \partial_\sigma(\theta)$ – which characterises Ω as the composition

$$\psi(\text{dom}(\theta)) \overset{\tau}{\cong} \text{dom}(\partial_\sigma(\theta)) \overset{\partial_\sigma(\theta)}{\cong_A} \text{cod}(\partial_\sigma(\theta)) \overset{\tau}{\cong} \psi(\text{cod}(\theta)),$$

so a *unique* extension of $\psi(\theta^+)$ matching $\partial_\sigma(\theta)$ and compatible with display maps – we fix $\psi(\theta) = \Omega$. Monotonicity is immediate from compatibility with dom and cod and that ψ preserves covering $\dashv\!\!-\!\!-$ on configurations. Likewise we extend ψ^{-1} to all symmetries; that ψ and ψ^{-1} are inverses follows as they preserve covering and are inverses on configurations. \square

We may now conclude our final characterization of composition:

Proposition D.15. *Consider $\sigma : A \vdash B$ and $\tau : B \vdash C$ causal strategies.*

Then, there is a strategy $\tau \odot \sigma : A \vdash C$, unique up to iso, such that there are order-isos:

$$\begin{aligned} (- \odot -) & : \{(x^\tau, x^\sigma) \in \mathcal{C}^+(\tau) \times \mathcal{C}^+(\sigma) \mid x^\sigma \text{ and } x^\tau \text{ causally compatible}\} \simeq \mathcal{C}^+(\tau \odot \sigma) \\ (- \odot -) & : \{(\theta^\tau, \theta^\sigma) \in \mathcal{S}^+(\tau) \times \mathcal{S}^+(\sigma) \mid \theta^\sigma \text{ and } \theta^\tau \text{ causally compatible}\} \simeq \mathcal{S}^+(\tau \odot \sigma) \end{aligned}$$

commuting with dom and cod, and s.t., for $\theta^\sigma \in \mathcal{S}^+(\sigma)$ and $\theta^\tau \in \mathcal{S}^+(\tau)$ causally compatible,

$$\partial_{\tau \odot \sigma}(\theta^\tau \odot \theta^\sigma) = \theta_A^\sigma \parallel \theta_C^\tau.$$

Proof. Existence. Simply a restriction of the isomorphisms of Proposition 4.10. By Lemma D.11, causally compatible $x^\sigma \in \mathcal{C}^+(\sigma)$ and $x^\tau \in \mathcal{C}^+(\tau)$ are automatically minimal, and by Lemma D.12, $x^\tau \odot x^\sigma \in \mathcal{C}^+(\tau \odot \sigma)$ is $+$ -covered. Reciprocally, if $x^\tau \odot x^\sigma \in \mathcal{C}^+(\tau \odot \sigma)$ is $+$ -covered, then by Lemma D.14, so are $x^\sigma \in \mathcal{C}^+(\sigma)$ and $x^\tau \in \mathcal{C}^+(\tau)$. Since the isomorphism

$(- \odot -)$ is compatible with dom and cod and symmetries are order-isomorphisms, there observations apply to symmetries. *Uniqueness.* By Lemma D.14. \square

D.3. Charactering immediate causality.

Lemma D.16. *For $\sigma : A \vdash B, \tau : B \vdash C$ causal prestrategies, for $m, m' \in |\tau \otimes \sigma|$, if $m \rightarrow_{\tau \otimes \sigma} m'$, then $m_\sigma \rightarrow_\sigma m'_\sigma$, or $m_\tau \rightarrow_\tau m'_\tau$, where m_σ, m_τ are defined whenever used.*

Proof. If $m \rightarrow_{\tau \otimes \sigma} m'$, then $x = [m]_{\tau \otimes \sigma} \setminus \{m, m'\} \in \mathcal{C}(\tau \otimes \sigma)$. Then we have

$$x = x^\tau \otimes x^\sigma \xrightarrow{m} y^\tau \otimes y^\sigma \xrightarrow{m'} z^\tau \otimes z^\sigma = [m]_{\tau \otimes \sigma}$$

in $\mathcal{C}(\tau \otimes \sigma)$, inlining the order-isomorphism of Proposition 4.7. Let us focus first on

$$x^\tau \otimes x^\sigma \xrightarrow{m} y^\tau \otimes y^\sigma.$$

Since $(- \otimes -)$ is an order-iso, this yields a covering in the partial order of causally compatible pairs, ordered by pairwise inclusion. By compatibility with display maps these inclusions add exactly one event in A , in B , or in C . This yields three cases: (a) m occurs in A , $x^\tau = y^\tau$, and $x^\sigma \prec y^\sigma$ adds one event $s \in |\sigma|$; (b) m occurs in C , $x^\sigma = y^\sigma$ and $x^\tau \prec y^\tau$ adds one event $t \in |\tau|$; or (c) m occurs in B , $x^\sigma \prec y^\sigma$ adds one event s and $x^\tau \prec y^\tau$ adds one event t with $\partial_\sigma(s) = (2, b)$ and $\partial_\tau(t) = (1, b)$.

Now, back to considering m and m' . If both occur in C , then we have

$$x^\tau \otimes x^\sigma \xrightarrow{m} y^\tau \otimes x^\sigma \xrightarrow{m'} z^\tau \otimes x^\sigma,$$

with $x^\tau \xrightarrow{t} y^\tau \xrightarrow{t'} z^\tau$. If $t <_\tau t'$, then $t \rightarrow_\tau t'$ and we are done. Otherwise, we also have

$$x^\tau \xrightarrow{t'} u^\tau \xrightarrow{t} z^\tau,$$

and as t' occurs in C we also have x^σ and u^τ causally compatible. Therefore

$$x^\tau \otimes x^\sigma \xrightarrow{n'} u^\tau \otimes x^\sigma \xrightarrow{n} z^\tau \otimes x^\sigma$$

for some $n, n' \in |\tau \otimes \sigma|$ since $(- \otimes -)$ is an order-isomorphism, so $\{n, n'\} = \{m, m'\}$. But since $u^\tau \otimes x^\sigma \neq y^\tau \otimes x^\sigma$ we must have $n = m$ and $n' = m'$, contradicting $m <_{\tau \otimes \sigma} m'$.

The case where m and m' both occur in A is symmetric. If m occurs in A and m' in C ,

$$x^\tau \otimes x^\sigma \xrightarrow{m} x^\tau \otimes z^\sigma \xrightarrow{m'} z^\tau \otimes z^\sigma$$

with $x^\sigma \xrightarrow{s} z^\sigma$ and $x^\tau \xrightarrow{t} z^\tau$. But then x^σ and z^τ are also causally compatible, and

$$x^\tau \otimes x^\sigma \prec z^\tau \otimes x^\sigma \prec z^\tau \otimes z^\sigma$$

which as above contradicts $m <_{\tau \otimes \sigma} m'$. The case where m occurs in C and m' occurs in A is symmetric. Now, assume m occurs in A and m' occurs in B . Then we have

$$x^\tau \otimes x^\sigma \xrightarrow{m} x^\tau \otimes y^\sigma \xrightarrow{m'} z^\tau \otimes z^\sigma$$

where $x^\sigma \xrightarrow{s} y^\sigma$, $x^\tau \xrightarrow{t'} z^\tau$, and $y^\sigma \xrightarrow{s'} z^\sigma$. If $s <_\sigma s'$, then $s \rightarrow_\sigma s'$ and we are done. Otherwise, we also have $x^\sigma \prec u^\sigma \prec z^\sigma$, u^σ and z^τ are also causally compatible, and

$$x^\tau \otimes x^\sigma \prec z^\tau \otimes u^\sigma \prec z^\tau \otimes z^\sigma,$$

contradiction. All cases with one event occurring in B and the other in A or C are symmetric. Finally, assume both m and m' occur in B . In that case, we have

$$x^\sigma \xrightarrow{s} y^\sigma \xrightarrow{s'} z^\sigma, \quad \text{and} \quad x^\tau \xrightarrow{t} y^\tau \xrightarrow{t'} z^\tau.$$

If we have $s <_\sigma s'$ then $s \rightarrow_\sigma s'$ and we are done, and likewise for $t \rightarrow_\tau t'$. Otherwise, $x^\sigma \xrightarrow{u} u^\sigma \xrightarrow{v} z^\sigma$ and $x^\tau \xrightarrow{u'} u'^\tau \xrightarrow{v'} z^\tau$, and $x^\tau \otimes x^\sigma \xrightarrow{u} u^\tau \otimes u^\sigma \xrightarrow{v} z^\tau \otimes z^\sigma$, contradiction. \square

Lemma D.17. *Consider $m, m' \in |\tau \otimes \sigma|$ such that $m \rightarrow_{\tau \otimes \sigma} m'$.*

Then, if $m_\sigma <_\sigma m'_\sigma$ we have $m_\sigma \rightarrow_\sigma m'_\sigma$, and likewise for τ .

Proof. If $m \rightarrow_{\tau \otimes \sigma} m'$, then there is $x^\tau \otimes x^\sigma, y^\tau \otimes y^\sigma$ and $z^\tau \otimes z^\sigma$ in $\mathcal{C}(\tau \otimes \sigma)$ such that

$$x^\tau \otimes x^\sigma \xrightarrow{m} y^\tau \otimes y^\sigma \xrightarrow{m'} z^\tau \otimes z^\sigma,$$

but if m_σ and m'_σ are defined then $x^\sigma \xrightarrow{m_\sigma} y^\sigma \xrightarrow{m'_\sigma} z^\sigma$. If $m_\sigma <_\sigma m'_\sigma$, $m_\sigma \rightarrow_\sigma m'_\sigma$. \square

D.4. The Bang Lemma. Now, we prove the *bang lemma* from AJM games [AJM00]. Fix A and B two concrete arenas with B pointed, and $\sigma : !A \vdash !B$ a causal strategy.

By *receptive*, for each $i \in \mathbb{N}$, there is a unique $q_i \in \min(\sigma)$ such that $\partial_\sigma(q_i) = (2, (i, b))$ the unique minimal move of B of copy index i . Let us write

$$|\sigma_i| = \{m \in |\sigma| \mid q_i \leq_\sigma m\}$$

for the set of events of σ above q_i . Since σ is pointed, it follows that for $i, j \in \mathbb{N}$ distinct, $|\sigma_i|$ and $|\sigma_j|$ are disjoint. Likewise, since arenas are concrete, moves in immediate conflict have the same predecessor – it follows that if $m \in |\sigma_i|$ and $n \in |\sigma_j|$, the negative dependencies of m and n are compatible, hence m and n are compatible by *determinism*. Therefore, we have $\sigma \cong \prod_{i \in \mathbb{N}} \sigma_i$ where σ_i has a structure of ess directly imported from σ . Furthermore, $\sigma_i : !A \vdash B$ with the display map ∂_{σ_i} defined in the obvious way. The key argument is:

Lemma D.18. *For any $i, j \in \mathbb{N}$, we have $\sigma_i \approx \sigma_j$.*

Proof. We exploit Lemma D.2 and build a (necessarily representable) iso between the domains of configurations, compatible with symmetry. Consider $x_i \in \mathcal{C}(\sigma_i)$, with

$$\partial_\sigma(x_i) = x_A \parallel \{i\} \times x_B.$$

From Definition 7.19, $x_A \parallel \{i\} \times x_B \cong_{!A \vdash !B}^- x_A \parallel \{j\} \times x_B$ with the obvious symmetry $\theta_{i,j}^-$. We must transport x_i along this negative symmetry $\theta_{i,j}^-$. By Lemma B.4 from [CCW19], there are unique $x_j \in \mathcal{C}(\sigma)$ and $\psi : x_i \cong_\sigma x_j$ s.t. $\partial_\sigma \psi = \theta^+ \circ \theta_{i,j}^-$ for some

$$\theta^+ : x_A \parallel \{j\} \times x_B \cong_{!A \parallel !B}^+ y_A \parallel \{j\} \times y_B$$

where we know that j is unchanged, because of condition *+transparent* of Definition 7.19. Therefore, $x_j \in \mathcal{C}(\sigma_j)$ as required. Monotonicity of this operation and the fact that it is a bijection between configurations follow from the *uniqueness* clause for Lemma B.4 of [CCW19]; compatibility with symmetry follows from composition with the symmetry ψ .

This induces an isomorphism of ess $\varphi : \sigma_i \approx \sigma_j$ which we must still check is a positive isomorphism. But for $x_i \in \mathcal{C}(\sigma_i)$, the symmetry θ^+ above entails

$$\partial_{\sigma_i}(x_i) = x_A \parallel x_B \cong_{!A \vdash B}^+ \varphi^+ y_A \parallel y_B = \partial_{\sigma_j}(x_j)$$

ensuring that the triangle commutes up to positive symmetry as required. \square

Next we lift a positive isomorphism on one copy index to the whole strategy:

Lemma D.19. *Consider A and B concrete $--$ arenas with B pointed, and $\sigma, \tau : !A \vdash !B$. If $\sigma_0 \approx \tau_0$, then $\sigma \approx \tau$.*

Proof. By Lemma D.18, for $i \in \mathbb{N}$, $\sigma_i \approx \sigma_0 \approx \tau_0 \approx \tau_i$, we conclude by parallel composition. \square

We may finally deduce the bang lemma:

Lemma D.20. *For concrete arenas A, B with B pointed and $\sigma \in \rightarrow\text{-Strat}(!A, !B)$,*

$$(\text{der}_B \odot \sigma)^\dagger \approx \sigma.$$

Proof. By Proposition 4.12 and a reasoning analogous to Proposition 4.19, we have $\text{der}_B \odot \sigma \approx \sigma_0$. Therefore, by Lemma D.19, for any two $\sigma, \tau : !A \vdash !B$, if $\text{der}_B \odot \sigma \approx \text{der}_B \odot \tau$, then $\sigma \approx \tau$. The lemma follows directly from that and the Seely category laws. \square

D.5. Expansion of Meager Strategies. Consider $\sigma : A$ parallel innocent. We first observe that any move $m \in |\sigma|$ is determined by $\text{mf}(m)$, along with the copy index of its negative dependencies. An **exponential slice** for $y \in \mathcal{C}(\text{mf}(\sigma))$ normal is an assignment of copy indices for all negative questions of y – or more precisely,

$$\alpha : y^{\mathcal{Q},-} \rightarrow \mathbb{N},$$

with $y^{\mathcal{Q},-}$ the negative questions of y . To any $x \in \mathcal{C}(\sigma)$ we have associated $\text{mf}(x) \in \mathcal{C}(\text{mf}(\sigma))$. To complete $\text{mf}(x)$, we also associate to x an exponential slice for $\text{mf}(x)$:

$$\begin{aligned} \text{slice}(x) & : (\text{mf}(x))^{\mathcal{Q},-} \rightarrow \mathbb{N} \\ n^- & \mapsto \text{ind}(\partial_\sigma(\theta_x^{-1}(n))). \end{aligned}$$

We now show how to reconstruct events of σ from $\text{mf}(\sigma)$ and an exponential slice.

Lemma D.21. *Consider A a concrete arena, $\sigma : A$ a causal strategy, and $y \in \mathcal{C}(\text{mf}(\sigma))$.*

For any $\alpha : y^{\mathcal{Q},-} \rightarrow \mathbb{N}$, there is a unique $x \in \mathcal{C}(\sigma)$ such that $y = \text{mf}(x)$, $\alpha = \text{slice}(x)$, and $x \cong_\sigma y$. Moreover, the symmetry $\varphi_y : y \cong_\sigma x$ is unique.

Proof. Same proof as for Lemma 7.24, setting up indices following α instead of 0. \square

Just as Lemma 7.24, for $\sigma : A$ parallel innocent, Lemma D.21 can be used to assemble an event $m \in |\text{mf}(\sigma)|$ and an exponential slice $\alpha : [m]^{\mathcal{Q},-} \rightarrow \mathbb{N}$ into $n \in |\sigma|$ s.t. $\text{mf}(n) = m$ and $\text{slice}([n]_\sigma) = \alpha$. So together, Lemmas 7.24 and D.21 establish a bijection between $|\sigma|$ and pairs (m, α) of $m \in |\text{mf}(\sigma)|$ and an exponential slice $\alpha : [m]_\sigma^{\mathcal{Q},-} \rightarrow \mathbb{N}$. From this it seems clear how to reconstruct σ from $\text{mf}(\sigma)$: first, we reconstruct a partial order.

Proposition D.22. *For A a concrete arena and $\sigma : A^+$, we define a partial order $\text{exp}(\sigma)$:*

$$\begin{aligned} |\text{exp}(\sigma)| & = \{(m, \alpha) \mid m \in |\sigma| \wedge \alpha : [m]_\sigma^{\mathcal{Q},-} \rightarrow \mathbb{N}\} \\ (m_1, \alpha_1) \leq_{\text{exp}(\sigma)} (m_2, \alpha_2) & \Leftrightarrow m_1 \leq_\sigma m_2 \wedge \forall n^- \leq_\sigma m_1, \alpha_1(n) = \alpha_2(n). \end{aligned}$$

Then, for any $\sigma : A$, σ and $\text{exp}(\text{mf}(\sigma))$ are isomorphic partial orders.

Proof. Direct consequence of Lemmas 7.24 and D.21. \square

It remains to complete $\exp(\sigma)$ into a causal strategy. The display map is determined by a choice of copy index for positive questions. Hence, for any $m^{\mathcal{Q},+} \in |\sigma|$, assume fixed some

$$f_m : \mathbb{N}^{[m]_{\sigma}^{\mathcal{Q},-}} \rightarrow \mathbb{N}$$

specifying, for each Player question, its copy index depending on indices of Opponent questions in its causal dependency. For simplicity, we assume this choice is globally injective, *i.e.* f_m, f_n have disjoint codomains for distinct $m, n \in |\sigma|$. The resulting strategy will not depend on the choice of the family $(f_m)_{m \in |\sigma|^{\mathcal{Q},+}}$ up to positive isomorphism.

Proposition D.23. *Consider A a concrete arena, $\sigma : A^+$ parallel innocent causally well-bracketed. We define a display map $\partial_{\exp(\sigma)}$ for $\exp(\sigma)$ by induction, with image*

$$\begin{aligned} \partial_{\exp(\sigma)}(m^A, \alpha) &= a, \\ \partial_{\exp(\sigma)}(m^{\mathcal{Q},-}, \alpha) &= q \quad \text{s.t.} \quad \text{ind}(q) = \alpha(m), \\ \partial_{\exp(\sigma)}(m^{\mathcal{Q},+}, \alpha) &= q \quad \text{s.t.} \quad \text{ind}(q) = f_m(\alpha), \end{aligned}$$

where a, q is the unique event of A with label $\text{lbl}(m)$, predecessor $\partial_{\exp(\sigma)}(\text{just}(m, \alpha))$ with $\text{just}(m, \alpha)$ defined as $\text{just}(m)$ with slice the restriction of α , label $\text{lbl}(m)$; satisfying the additional constraint given. Further components are:

$$\begin{aligned} (m, \alpha) \#_{\exp(\sigma)} (n, \beta) &\Leftrightarrow m \#_{\sigma} n, \text{ and } \alpha \text{ and } \beta \text{ coincide on their common domain} \\ \theta \in \mathcal{S}(\exp(\sigma)) &\Leftrightarrow \theta : x \cong y \text{ order-iso s.t. } \pi_1 = \pi_1 \circ \theta. \end{aligned}$$

Then, $\exp(\sigma) : A$ is parallel innocent causally well-bracketed.

The only difficulty is in handling conflict. A positive iso $\varphi : \sigma_1 \approx \sigma_2$ directly lifts to $\exp(\sigma_1) \approx \exp(\sigma_2)$, applying φ to moves while keeping copy indices unchanged. Finally:

Corollary D.24. *For A a concrete arena, the operations $\text{mf}(-)$ and $\exp(-)$ yield a bijection between (positive isomorphism classes of) causally well-bracketed strategies on A and A^+ .*

Proof. For $\sigma : A^+$, $\text{mf}(\exp(\sigma)) \approx \sigma$ by construction. For $\sigma : A$, $\exp(\text{mf}(\sigma)) \approx \sigma$ is obtained from the iso of Proposition D.22 and a verification that this preserves further structure. \square