

Causal Investigations in Interactive Semantics



Mémoire présenté par

Pierre Clairambault

en vue d'obtenir

l'Habilitation à Diriger des Recherches
(Informatique)

le 21 février 2024 devant le jury composé de:

Dan Ghica	Huawei, University of Birmingham	Rapporteur
Delia Kesner	IRIF, Université Paris Cité	Examinatrice
Ugo Dal Lago	University of Bologna	Examineur
Damiano Mazza	CNRS, LIPN, Université Paris Nord	Président
Guy McCusker	University of Bath	Rapporteur
Paul-André Melliès	CNRS, IRIF, Université Paris Cité	Rapporteur
Laurent Regnier	Aix-Marseille Université	Examineur

Remerciements

Tout d’abord, merci à tous mes collaborateurs passés sur ce formidable thème des jeux concurrents. Merci à Glynn pour m’avoir fait découvrir les structures d’événements et avoir accompagné mes premiers pas dans cette belle direction de recherche. Merci à Simon C.: nos séances de travail sur les jeux concurrents fins et l’innocence parallèle resteront parmi mes plus beaux souvenirs de collaboration scientifique. Merci à Aurore pour m’avoir accompagné sur les traces de Herbrand. Merci à Hugo pour une touche de probabilité, et pour m’avoir convaincu que les bicatégories, c’est après tout peut-être pas complètement inutile. Merci à Marc d’avoir insisté pour faire du quantique, menant à, à mon avis, une des plus belles applications des jeux concurrents. Merci à Lison et Lionel grâce à qui j’ai enfin pu comprendre le développement de Taylor (c’est de la sémantique des jeux). Merci à Simon F., grâce à qui on peut maintenant faire des jeux concurrents fins sans jeux concurrents.

Merci à tous les collègues de l’équipe Plume du LIP de m’avoir patiemment écouté parler de structures d’événements pendant toutes ces années, et en particulier Colin qui ne pouvait pas s’enfuir puisque nous partageons le même bureau. Merci à tous les collègues marseillais, du LIS et de l’I2M, d’avoir pris le relais, de m’avoir si bien accueilli lors de mon arrivée à Marseille, et de contribuer à un cadre de vie et de travail exceptionnel – en particulier Raphaëlle, qui a brillamment pris la suite de Colin.

Merci à Dan, Guy et Paul-André d’avoir accepté de rapporter ce travail, malgré son volume déraisonnable: vous êtes sûrement les trois chercheurs dont le travail a le plus influencé ma vision scientifique, et je vous suis reconnaissant de votre relecture approfondie de ma contribution. Merci à Delia, Ugo et Damiano d’avoir bien voulu être membres de ce jury, et merci à Laurent de m’avoir guidé dans cette aventure de l’habilitation.

Merci à tous ceux que j’oublie et qui regrettent de ne pas trouver leur nom ici.

Et pour tout le reste, merci à Marie-Laure.

Résumé

La *sémantique des jeux* est un cadre mathématique puissant pour raisonner de façon compositionnelle sur les programmes. Contrairement aux méthodes plus classiques de sémantique dénotationnelle, elle ne voit pas les programmes simplement comme des fonctions mathématiques prenant une valeur en entrée et rendant une valeur en sortie, mais comme des processus dynamiques et interactifs. Elle capture de façon compositionnelle le flot de contrôle même pour des programmes disposants de primitives calculatoires complexes: ordre supérieur, état mutable, opérateurs de contrôle, etc.

Mais la sémantique des jeux n'est pas une *théorie mathématique*: c'est plus une constellation de modèles disparates sans réel liant autre que le folklore de quelques experts du domaine. L'un de ces modèles disparates est celui des *jeux concurrents*, qui hérite des jeux concurrents d'Abramsky et Melliès et des jeux asynchrones de Melliès. Dans sa formulation récente proposée par Rideau et Winskel, ce modèle est basé sur les *structures d'événements*, un modèle de *concurrency vraie* qui évite de supposer l'existence d'une horloge globale et représente plutôt les programmes par leur structure causale.

Bien entendu, les jeux concurrents permettent de modéliser la programmation concurrente. Mais dans ce manuscrit, je montre que plus qu'un modèle comme un autre, leur expressivité permet de mettre en avant des structures implicites dans plusieurs autres modèles de la littérature pour des langages concurrents, ou pas. Ce faisant, sur les traces de la vision d'Abramsky et Melliès, ils permettent une synthèse jusqu'alors absente et contribuent à faire de la sémantique des jeux une théorie mathématique.

Concrètement, ce manuscrit commence en Partie I par une introduction détaillée à différentes sémantiques des jeux traditionnelles *à pointeurs*: d'abord pour le langage fonctionnel paradigmatique PCF, son extension IA avec mémoire mutable, puis le modèle de jeux non alternants pour son extension concurrente IA_{\parallel} . En Partie II, je donne une présentation complète des jeux concurrents et de leur extension avec symétrie. En partie III, je relie les jeux concurrents aux différents modèles de la Partie I – mais aussi d'autres, tels que le modèle relationnel – via des foncteurs préservant l'interprétation. Finalement, en Partie IV, je passe en revue d'autres développements en jeux concurrents, et conclus sur des problèmes ouverts et perspectives futures.

Abstract

Game Semantics is a powerful mathematical framework to reason compositionally on programs. In contrast to other more classical methods in denotational semantics, game semantics sees programs not merely as mathematical functions taking inputs (which may be functions themselves) and producing outputs, but rather, sees them as dynamic and interactive processes. It captures compositionally the control flow even for programs with complex computational primitives such as higher-order, mutable state, control operators, etc.

However, game semantics is not a *mathematical theory*: rather, it is a loose collection of separated technical settings, sometimes only loosely connected by the folklore of a handful of experts. One of those technical setting is *concurrent games*, inheriting from the concurrent games of Abramsky and Melliès and from Melliès' asynchronous games. In its recent technical set-up proposed by Rideau and Winskel, this model is based on *event structures*, a *true concurrency* model that avoids assuming the existence of a global clock, representing instead programs via their causal structure.

Of course, concurrent games can handle the semantics of concurrent programs. But in this monograph, I show that more than yet another new model, their expressivity reveals structures implicitly present in several other models from the literature, for concurrent languages, or not. In doing so, walking in the footsteps of Abramsky and Melliès' vision, they give us the means for a synthesis between various games or other denotational models, bringing game semantics one step closer to a mathematical theory.

Concretely, this monograph starts in Part I with a detailed introduction to several traditional *pointer* game semantics: first for the paradigmatic purely functional language PCF, its extension IA with mutable state, and the non-alternating variant for its concurrent extension IA_{\parallel} . In Part II, I give a complete presentation of concurrent games and their extension with symmetry. In Part III, I link concurrent games to the different games models introduced in Part I – but also other models, such as the relational model – via interpretation-preserving functors. Finally, in Part IV, I survey other developments in concurrent games, and I conclude with some open problems and perspectives.

Contents

1	Introduction	12
1.1	Operational Semantics	12
1.2	Toward a Mathematical Semantics	13
1.3	The Next 700 Denotational Semantics	16
1.4	Full Abstraction and Intentional Intensionality	17
1.5	Plays Considered Harmful	18
1.6	A Case for Causal Game Semantics	21
1.7	Contributions and Outline	23
2	General Preliminaries	26
2.1	Mathematical Language	26
2.2	IA_{\parallel} and its Fragments	27
2.2.1	Syntax of IA_{\parallel}	28
2.2.2	Operational Semantics of IA_{\parallel}	30
2.2.3	Expressiveness of IA_{\parallel}	30
2.3	Generalities on Semantics	34
I	Pointer Game Semantics	37
3	First Steps in Game Semantics	41
3.1	Executions, Plays, and Strategies	41
3.1.1	Dialogues, Plays, and Simple Games	41
3.1.2	Replication and Thread Indexing	44
3.1.3	Pointer Games	45
3.2	Carving Out Innocent Strategies	48
3.2.1	The Strategies of PCF	48
3.2.2	Well-Bracketing	50
3.2.3	Visibility and Innocence	52
3.2.4	Finite Definability and Full Abstraction	54
3.3	Non-Innocence and Effects	57
3.3.1	Non-Local State	57
3.3.2	Higher-Order State	58

3.3.3	The “Semantic Cube”	59
3.4	An Alternative to Pointers: Copy Indices	63
3.5	Conclusions and Historical Notes	66
4	The Category of Alternating Strategies	67
4.1	The Ambient Cartesian Closed Category $\Downarrow\text{-Strat}$	67
4.1.1	More Arena Constructions	67
4.1.2	Composition of Strategies	68
4.1.3	A Category of Arenas and Strategies	72
4.1.4	Cartesian Closed Structure	74
4.1.5	Recursion	76
4.2	Interpretation of \mathbf{IA}	76
4.2.1	Interpretation of PCF	77
4.2.2	Interpretation of State	77
4.3	Complements on Alternating Strategies	82
4.3.1	Complements on Single-Threadedness	82
4.3.2	Factorization and Definability	84
4.3.3	Summary of Results	85
4.4	Conclusions and Historical Notes	85
5	Non-Alternating Game Semantics	87
5.1	Concurrency and Non-Alternation	87
5.1.1	Non-Alternating Plays	87
5.1.2	Non-Alternating Strategies	90
5.1.3	Single-Threadedness	92
5.1.4	Full Abstraction for \mathbf{IA}_{\parallel}	93
5.2	The Ambient Cartesian Closed Category $\mathcal{O}\text{-Strat}$	96
5.2.1	Constructing $\mathcal{O}\text{-Strat}$	96
5.2.2	Cartesian Closed Structure	98
5.2.3	Recursion	99
5.3	Interpretation of \mathbf{IA}_{\parallel}	100
5.3.1	Interpretation of PCF $_{\parallel}$	100
5.3.2	Interpretation of State	102
5.3.3	Summary of Results	103
5.4	Conclusions and Historical Notes	103

II Thin Concurrent Games

104

6	Basic Concurrent Games	107
6.1	Games and Strategies as Event Structures	107
6.1.1	Basic Intuitions on Concurrent Strategies	107
6.1.2	Formalizing Concurrent Strategies	109
6.1.3	Basic Properties of Event Structures and Maps	113
6.2	Composition of Prestrategies	118
6.2.1	Interaction of prestrategies	119

6.2.2	Composition of prestrategies	127
6.2.3	The associator	129
6.3	Composition of Strategies	131
6.3.1	Composition of Strategies	131
6.3.2	Strategies and $+$ -covered configurations	132
6.3.3	Composition and $+$ -coveredness	133
6.4	The Bicategory CG	135
6.4.1	Copycat	135
6.4.2	Copycat and Strategies	137
6.4.3	Horizontal Composition and Bicategorical Structure	138
6.4.4	Characterization of Strategies	139
6.5	Conclusions and Historical Notes	140
7	Thin Concurrent Games	142
7.1	Symmetry in Games and Strategies	142
7.1.1	Event Structures with Symmetry	143
7.1.2	Saturated Games with Symmetry	146
7.1.3	Thin Concurrent Games	149
7.1.4	Constructions on Thin Concurrent Games	152
7.2	Mediating Between Strategies	155
7.2.1	A Zoology of Morphisms and Equivalences	155
7.2.2	A Study of Positive Morphisms	157
7.2.3	All Equivalences Coincide	160
7.2.4	Constructing Positive Morphisms	162
7.3	Composition and Copycat	164
7.3.1	Interactions with Symmetry	165
7.3.2	Composition with Symmetry	168
7.3.3	Copycat	171
7.4	The Bicategory TCG	174
7.4.1	Synchronization up to Symmetry	176
7.4.2	Horizontal Composition	179
7.4.3	Bicategorical Laws	181
7.4.4	A \sim -category	185
7.5	History and Related Work	186
8	Constructing Games and Strategies	188
8.1	A Compact closed \sim -Category	189
8.1.1	A bifunctor	189
8.1.2	Lifting Maps to Strategies	191
8.1.3	Compact Closed Structure	192
8.2	Negative Winning Games	194
8.2.1	Relative Seely categories	195
8.2.2	Playing Board Games	196
8.2.3	Negative Winning Strategies	201
8.2.4	Symmetric Monoidal Structure	202
8.2.5	Cartesian Product of Strict $--$ -Boards	204

8.2.6	Relative Closure	205
8.3	Non-Linear Structure	208
8.3.1	The Resource Modality	208
8.3.2	Recursion	215
8.4	History and Related Work	217

III Disentangling Parallelism and State 219

9 The Causal Semantics of IA_{\parallel} and its Unfolding 222

9.1	Interpretation of IA_{\parallel}	222
9.1.1	Interpretation of PCF_{\parallel}	222
9.1.2	Semantics of state	226
9.1.3	Discussion on the Interpretation	229
9.2	Unfolding Causal to Non-Alternating Strategies	233
9.2.1	A First Unfolding	233
9.2.2	Mixed Boards and Pointifixion	235
9.2.3	Unfoldings of a Causal Strategy	240
9.3	\rightarrow - Strat and Full Abstraction	242
9.3.1	Unfolding the Categorical Structure	242
9.3.2	Unfolding the Interpretation of IA_{\parallel}	247
9.4	History and Related Work	248

10 Parallel Innocence 249

10.1	Defining Parallel Innocence	250
10.1.1	Causal Determinism	250
10.1.2	Pre-innocence	253
10.1.3	Visibility	254
10.2	Composition of Visibility	256
10.2.1	Justifiers in causal strategies	256
10.2.2	Justifiers in interactions	257
10.2.3	Views of gccs	257
10.2.4	The Relative Seely \sim -Category NTCG-Vis	259
10.3	Composition of Innocence	259
10.3.1	The “forking lemma”	260
10.3.2	Stability of pre-innocence	261
10.3.3	The Relative Seely Category NTCG-Inn	262
10.3.4	Complement: the “Bang Lemma”	263
10.4	Relational Collapse	265
10.4.1	Stopping Positions and Witnesses	265
10.4.2	Composition and Deadlocks	267
10.4.3	The Deadlock-Free Lemma	269
10.4.4	A Relative Seely \sim -Functor	273
10.5	Globularity	278
10.5.1	Motivation and Definition	278
10.5.2	Composition of Globularity	280

10.6 History and Related Work	281
11 Sequentiality	282
11.1 Sequentiality	283
11.1.1 Definition of Sequentiality	283
11.1.2 Categorical Structure	286
11.1.3 A Relative Seely \sim -Category	290
11.1.4 \rightarrow -Seq and Interpretation of IA	291
11.2 The Alternating Unfolding	292
11.2.1 Alternating Unfolding on a Mixed Board	292
11.2.2 Alternating Unfoldings of Sequential Strategies	295
11.2.3 Unfolding the Basic Categorical Structure	297
11.2.4 Unfolding the Interpretation of IA	299
11.3 Sequential Innocence	301
11.3.1 Causal Analysis of Sequential Innocence	301
11.3.2 The Unfolding Preserves Innocence	304
11.3.3 Sequential Globularity	306
11.3.4 Intensional Full Abstraction	307
11.4 History and Related Work	308
12 Finite Definability for PCF_{//}	309
12.1 Meager Form	310
12.1.1 Updating Mixed Boards	311
12.1.2 Meager Innocent Strategies	313
12.1.3 The Meager Form of Sequential Globularity	321
12.1.4 Finite Tests Suffice	322
12.2 Factorization	323
12.2.1 The Flow Substrategy	324
12.2.2 The Argument Substrategies	326
12.2.3 Recomposition	330
12.3 Finite Definability and Full Abstraction	333
12.3.1 First-Order Definability	333
12.3.2 Positional First-Order Definability	334
12.3.3 Finite Globular Definability	337
12.3.4 Intensional Full Abstraction	338
IV Other Developments and Openings	339
13 Further Work	341
13.1 Effects and Concurrency	341
13.1.1 Non-canonical causal presentation	343
13.1.2 Non-angelic concurrent games	343
13.1.3 Resource-tracking concurrent games	345
13.2 Quantitative Concurrent Games	346
13.2.1 Probabilistic PCF	347

13.2.2	Quantum strategies	349
13.3	Quantitative Relational Collapses	352
13.3.1	Counting witnesses and quantitative collapse	352
13.3.2	Generalized species of structure	355
13.4	Game Semantics as Taylor Expansion	357
13.4.1	Pointer concurrent games and positional injectivity	357
13.4.2	Isogmentations and the resource calculus	359
13.4.3	Taylor expansion, extensional resource terms	360
13.5	Operational Concurrent Games	360
13.6	Miscellaneous	363
13.6.1	Full Abstraction for Parallel-Or	363
13.6.2	A Semantic Proof of Herbrand's Theorem	366
13.6.3	Revisiting Games Models of MALL	368
14	Conclusions	370
14.1	The Work Done so Far	370
14.2	Open Problems	371
14.2.1	Non-deterministic parallel innocence	371
14.2.2	Proper determinism	373
14.2.3	All well-bracketings	374
14.2.4	Observing causality	374
14.3	Perspectives and Future Directions	375
V	Appendices	393
A	Complements on Relative Seely Categories	394
A.1	Definition and Kleisli Category	395
A.2	Relative Seely Functors	396

Preface

This document is my HDR – “Habilitation à Diriger les Recherches”. As such, its principal aim is to motivate and present the heart of my research work so far.

Of course, I was not the only one in this adventure, this is the product of collaborations with a number of people. Firstly with Glynn Winskel: I started working with Glynn during my postdoc in Cambridge from September 2011 to September 2013, and our collaboration kept going strong long after that. My work owes a lot to him; many results presented here may not exist at all if not for his encouragements and insights.

Secondly, almost of my further work in this direction is collaborative, and was carried out with PhD students. First with Simon Castellan (2014 – 2017), with whom we constructed thin concurrent games and parallel innocence. This forms the core of the material that I have chosen to present in this monograph, though I believe the presentation has since then significantly gained in maturity. Then, with Aurore Alcolei (2016 – 2019) and Marc de Visme (2017 – 2020). Though the work done with them will not be covered therein, this document does greatly benefit from the intuitions gained with them – their work will be mentioned in the outline of further work in Chapter 13. Finally, the story currently continues with Lison Blondeau-Patissier (2021 – 2024).

No research is done in a vacuum and this work, as usual, stands on the shoulders of giants. First and foremost, the giants include the pioneers of game semantics: Martin Hyland and Luke Ong; Samson Abramsky, Radha Jagadeesan and Pasquale Malacaria; Hanno Nickau; and Thierry Coquand. Beyond those, the giants also include those who realized the strength of game semantics as a means to represent computational effects, for instance Guy McCusker, Jim Laird, Dan Ghica and Andrzej Murawski. And finally, the giants include those who kept deciphering the foundations of game semantics and maintained the sometimes thinning connection with linear logic: for instance, Pierre-Louis Curien, Laurent Regnier and Paul-André Melliès. The present work builds on their inheritance (and surely that of many others), and to an extent is a synthesis of their ideas. It is not at all obvious that such a synthesis was possible: it was first foreseen by Paul-André Melliès that this was possible, building on techniques from concurrency theory, and to a large extent this was the intent behind his line of work on *asynchronous games*. This line of work, which for simplicity (and although it can be slightly misleading) I will often refer to as *concurrent games*, inherits Paul-André’s vision. Many of our technical steps parallels his, with the hindsight provided by 15 more years of work.

Beyond presenting this work, this document has two further objectives:

Firstly, I intend it as an entry point to my work on concurrent games. The technical underpinning of concurrent games is fairly elaborate: event structures are much more mathematically demanding than sequential structures (traces, trees, etc). The foundations are spread over a number of papers, written over the years. These papers rely on definitions and notations that are not quite consistent with each other, as our conceptual understanding progressed over time. In this manuscript, I present a complete and detailed construction of *Thin Concurrent Games*, a bicategory of concurrent strategies which plays a pivotal role in my further work – this is the purpose of Part II. I also present in Part III a detailed presentation of two developments at the core of the theory of Thin Concurrent Games, which I will introduce and motivate in Section 1.7.

Secondly, I needed an introduction to game semantics for programming languages, beyond merely concurrent games. In part, this is because I needed a precise set of definitions to refer to for the work presented in Part III. But beyond that, over the years I had come to dread the question “Where can I find a good introduction to game semantics”? While good introductory texts do exist, I felt that none gave a satisfactory answer. Indeed, many were written in the late 90s or earlier 00s and are not very modern or up-to-date. Others give a good survey, focusing on intuitions, but do not offer precise technical definitions; or they do offer a technical introduction, but to one particular technical setting, ignoring the wider picture. Hence, Part I is the result of my attempt at an accessible introduction to game semantics for programming languages.

Chapter 1

Introduction

Attempts to formally prove properties of the execution of programs (*e.g* functional correctness, safety, termination, *etc*) must rest on *formal semantics* turning the execution into a well-specified mathematical object. This is the main purpose of the field of *semantics of programming languages*, for which a reference textbook is [Winskel, 1993].

1.1 Operational Semantics

Typically, this is done via the methodology of **operational semantics**: by means of formal rules operating on syntax, showing how the source code is transformed throughout computation. For instance, assuming we have a value \underline{n} for every $n \in \mathbb{N}$, the rules

$$\frac{e_1 \rightsquigarrow e'_1}{e_1 + e_2 \rightsquigarrow e'_1 + e_2} \qquad \frac{e_2 \rightsquigarrow e'_2}{\underline{n} + e_2 \rightsquigarrow \underline{n} + e'_2} \qquad \frac{}{\underline{n} + \underline{m} \rightsquigarrow \underline{n + m}}$$

inductively define an evaluation relation \rightsquigarrow between additive arithmetic expressions, specifying the operational behaviour of the sum: here, it has a left-to-right evaluation strategy: given a sum expression $e_1 + e_2$, we first evaluate e_1 to a value \underline{n} – which may take many steps – then evaluate e_2 to a value \underline{m} . Once both operands are values, the expression $\underline{n} + \underline{m}$ is rewritten to $\underline{n + m}$. For instance, we have the reduction sequence

$$(2 + 3) + (7 + (1 + 2)) \rightsquigarrow 5 + (7 + (1 + 2)) \rightsquigarrow 5 + (7 + 3) \rightsquigarrow 5 + 10 \rightsquigarrow 15,$$

omitting the underlining. Each step above is justified by a derivation, *e.g.*

$$\frac{\frac{\frac{}{1 + 2 \rightsquigarrow 3}}{7 + (1 + 2) \rightsquigarrow 7 + 3}}{5 + (7 + (1 + 2)) \rightsquigarrow 5 + (7 + 3)}}$$

for the second step. This would be usually called a *small-step operational semantics*; *big-step semantics* also exist, where one axiomatises directly the evaluation relation $e \Downarrow \underline{n}$ of an expression to a value. More generally, there is a large body of literature studying variations of operational semantics; including a mathematical theory of operational semantics initiated by Turi and Plotkin [Turi and Plotkin, 1997].

Operational semantics is a powerful methodology for formal semantics of programming languages. It is very robust to the addition of syntactic constructs or programming features, and has proved capable of handling realistic programming languages. As it builds on simple inductive structures, it is amenable to mechanization in a proof assistant – as powerfully illustrated in the celebrated CompCert project by Leroy *et al* [Leroy, 2009, Krebbers et al., 2014]. However, operational semantics also has some drawbacks. By design, it is very much tied to syntax. More importantly, most frameworks are exclusively designed to reason on complete programs, operating in a closed world. They are ill-equipped to reason about *open* programs, that may include free variables or appeal to external libraries, and consequently are not compositional¹.

1.2 Toward a Mathematical Semantics

Another approach to the formal semantics of programs is that of **denotational semantics**, initiated by Scott and Strachey under the name of *mathematical semantics* [Scott and Strachey, 1971]. Following the methodology of denotational semantics, any program M is sent to a mathematical object $\llbracket M \rrbracket$, its *meaning* or **denotation**. It is usually an invariant of evaluation, in the sense that if $M \rightsquigarrow M'$, then $\llbracket M \rrbracket = \llbracket M' \rrbracket$. Thus a closed program of ground type is typically sent directly to the value it returns, if any, giving no insight on the evaluation process itself (at least explicitly). In contrast with operational semantics, denotational semantics focuses on *open programs*. To those it also gives a proper meaning, often as a function, as in the example below:

$$\begin{array}{rcl} \llbracket x + (x + 5) \rrbracket & : & \mathbb{N} \rightarrow \mathbb{N} \\ n & \mapsto & 2n + 5 \end{array}$$

In contrast with operational semantics, compositionality is at the core of the methodology of denotational semantics. A denotation is assigned to each syntactic construct, and the denotation of the program is obtained by replacing modularly each piece of syntax by its denotation, as illustrated by the computation below:

$$\begin{aligned} \llbracket x + (x + 5) \rrbracket(n) &= \llbracket + \rrbracket(\llbracket x \rrbracket(n), \llbracket x + 5 \rrbracket(n)) \\ &= \llbracket + \rrbracket(n, \llbracket + \rrbracket(\llbracket x \rrbracket(n), \llbracket 5 \rrbracket(n))) \\ &= \llbracket + \rrbracket(n, \llbracket + \rrbracket(n, 5)) \\ &= \llbracket + \rrbracket(n, n + 5) \\ &= 2n + 5 \end{aligned}$$

¹There are, of course, formalisms for compositional operational semantics, used for instance to make CompCert compositional [Stewart et al., 2015]. By and large, those put into play structures strongly related to denotational semantics; for instance [Stewart et al., 2015] acknowledge similarity with Ghica and Tzevelekos' *system-level game semantics* [Ghica and Tzevelekos, 2012].

where $\llbracket + \rrbracket : \mathbb{N}^2 \rightarrow \mathbb{N}$ is the usual arithmetical addition operation.

The traditional core of denotational semantics is *domain theory*, in which types are represented by certain posets satisfying completeness properties useful to represent infinite data and partial computation, and open or higher-order programs are represented as (continuous) functions. But this traditional approach to denotational semantics has limitations, the focus on representing programs as functions being in tension with computational effects such as non-deterministic choice, probabilistic choice, references, exceptions, *etc* (though many of those effects can be accommodated via adequate monads [Moggi, 1991]). The modern understanding of *denotational semantics* is often taken more generally to mean a compositional interpretation of programs, usually structured categorically (typically as a cartesian closed category [Lambek and Scott, 1988]). In this wider understanding, denotational semantics may be regarded as the art of reasoning compositionally about program behaviour, which is of paramount importance for the purpose of having formal semantic techniques that scale [O’Hearn, 2015].

Over the years, a wealth of denotational models have appeared for programming languages, according to which programs are represented as other structures than functions. For instance, in the *relational model* [Girard, 1988], we have

$$\begin{aligned} \llbracket x + (x + 5) \rrbracket &\subseteq \mathcal{M}_f(\mathbb{N}) \times \mathbb{N} \\ &= \{([n, m], n + m + 5) \mid n, m \in \mathbb{N}\} \end{aligned}$$

where $\mathcal{M}_f(X)$ is the set of *finite multisets* of elements of X – in this semantics, we record that we may obtain value $n + m + 5$ provided we have *two* successful evaluations of the argument, one yielding value n , the other yielding value m . In contrast to the functional models mentioned earlier, this one is *quantitative*: we see in the interpretation that the program makes *two* calls to its argument, an information lost in domain semantics.

We also account for the behaviour of the program against a potentially non-deterministic environment: for instance, while in functional semantics the program $x \mathbf{xor} x$ will return \mathbf{ff} for any value of x (writing \mathbf{tt} and \mathbf{ff} the two boolean values), we have

$$([\mathbf{tt}, \mathbf{ff}], \mathbf{tt}) \in \llbracket x \mathbf{xor} x \rrbracket$$

in relational semantics, accounting for an execution of $x \mathbf{xor} x$ in the context of a non-deterministic x . Note that here, $[\mathbf{tt}, \mathbf{ff}] = [\mathbf{ff}, \mathbf{tt}]$. There is no sense in which an element of the multiset may be unambiguously assigned to an occurrence of x – we do not know if the program returns \mathbf{tt} because the first x returned \mathbf{tt} and the second \mathbf{ff} , or the other way around. We can enrich the model to account for the fact that $([\mathbf{tt}, \mathbf{ff}], \mathbf{tt})$ is realized by these *two* executions, by assigning a *weight* to every point of the relational model

$$\llbracket x \mathbf{xor} x \rrbracket([\mathbf{tt}, \mathbf{ff}], \mathbf{tt}) = 2$$

and obtain in this way a *weighted relational model* [Laird et al., 2013], in which two terms of PCF are sent to the same representation iff they cannot be distinguished by PCF with probabilistic choice [Ehrhard et al., 2018]!

But the above still forgets much about the program, for instance it tells us nothing about the behaviour of the program under a stateful execution environment. For that,

we must yet again move up in the intensionality ladder, into the realm of *game semantics* [Hyland and Ong, 2000, Abramsky et al., 2000], and record the behaviour of the program over time as illustrated in the diagram below

$$\begin{array}{c}
 \llbracket x \text{ xor } x \rrbracket \quad : \quad x : \mathbb{B} \quad \vdash \quad \mathbb{B} \\
 \qquad \qquad \qquad \mathbf{q}^+ \\
 \qquad \qquad \qquad \mathbf{tt}^- \\
 \qquad \qquad \qquad \mathbf{q}^+ \\
 \qquad \qquad \qquad \mathbf{ff}^- \\
 \qquad \qquad \qquad \mathbf{tt}^+
 \end{array}$$

breaking the symmetry between the two calls by explicitly showing them over time.

There are multiple ways to enrich denotational models with time, leading to many mathematical settings for game semantics: for instance, one can interpret a PCF program in Abramsky and McCusker’s model [Abramsky and McCusker, 1996] which characterizes programs undistinguishable with ground state, in Berry and Curien’s sequential algorithms which characterize indistinguishability in the presence of control operators [Cartwright et al., 1994], in Murawski’s model for indistinguishability by state and control [Murawski, 2007], in Ghica and Murawski’s model for indistinguishability by shared state and concurrency [Ghica and Murawski, 2008]... That is just the tip of the iceberg, and only listing game semantics models. Overall, a term of PCF can be interpreted into maybe a dozen different denotational models, sometimes similar and sometimes different, each recording a different aspect of its interactive behaviour².

In our tour of denotational models we should certainly also mention bicategorical models such as *generalized species of structure* [Fiore et al., 2008] or Melliès’ *template games* [Melliès, 2019a], which assign to each point of the interpretation of a type in the relational model a *set* (or a category) of witnesses, presenting the different executions realizing that point – so rather than merely stating $\llbracket x \text{ xor } x \rrbracket([\mathbf{tt}, \mathbf{ff}], \mathbf{tt}) = 2$, the interpretation would yield a particular two-element set. Finally, all these developments phrased explicitly as denotational models fit into a more general landscape of tools and methods to approximate program behaviour (including idempotent or non-idempotent intersection types, resource or differential calculi, the Taylor expansion of programs, *etc*). There is an impressive breadth here: denotational semantics go way beyond the somewhat boring formalization of the “intended” denotation of a term as an official mathematical function, it is in fact the quite varied art of extracting compositional information from programs, and of reasoning about program identity.

²Game semantics models, along with others such as Girard’s *geometry of interaction* [Girard, 1989], are often collectively referred to as *interactive semantics*. But this might be a misnomer: in reality, all the above semantics are interactive semantics in the sense that they all collect information about the interactive or compositional behaviour of the program, as opposed to the mechanics of its evaluation.

1.3 The Next 700 Denotational Semantics

It is fascinating that so many subtly different compositional accounts can be given, even if one considers only a single computational object such as PCF. What is also striking, is that there are comparatively surprisingly few results formally *relating* all these different models – even though many of those largely rely on similar intuitions, notably inherited from Linear Logic [Girard, 1987]. One explanation is that the community has focused more on building new models, taming new languages and effects, than on linking existing models. There is also a social aspect, individual researchers sticking with their favourite tools and methods. But besides those reasons, it remains true that such connections have in the past proved much more elusive than anticipated.

This is well-illustrated by the *relational collapse* of game semantics: we introduced above game semantics as relational semantics plus time, so it is only natural to attempt to send game semantics to the relational model by simply forgetting time:

$$\begin{array}{ccc}
 x : \mathbb{B} & \vdash & \mathbb{B} \\
 & & \mathbf{q}^- \\
 \mathbf{q}^+ & & \\
 \mathbf{tt}^- & & \\
 \mathbf{q}^+ & & \rightsquigarrow ([\mathbf{tt}, \mathbf{ff}], \mathbf{tt}) \\
 \mathbf{ff}^- & & \\
 & & \mathbf{tt}^+
 \end{array}$$

but this almost simple-minded idea hides two major conceptual hurdles: firstly, the traditional notion of plays in game semantics has time too wired in, and does not support an adequate notion of *position*. Secondly, composition of strategies also involves time; two interacting strategies may agree on a common position but not on a common *trajectory*, preventing synchronization. Understanding this prompted the work of many researchers, spanning over two decades [Baillot et al., 1997b, Melliès, 2005, Boudes, 2009, Calderon and McCusker, 2010, Clairambault and Paquet, 2021].

There are a few landmark results relating different models, such as Ehrhard’s *extensional collapse* theorem [Ehrhard, 2012], which links the qualitative and quantitative views of resource replication. But by and large, a lot of seemingly simple questions turn out to hide a surprising amount of conceptual and technical depth³.

Despite this relative rarity of actual formal connections, experts in the field do have a strong intuition as to what all these models record, and where they intuitively stand in relation to each other. After all, intuitively all those models are not independent: they all refer to, and present a part of, the interactive behaviour of the program. Over time, the working semanticist builds up an intuition of this full interactive behaviour, through the lens of their favourite technical tools. This “full interactive behaviour” is, arguably,

³As a more recent example, generalized species of structure are often presented – as we did above – as refining the weighted relational model, replacing mere cardinality information by an explicit of witnesses. This suggests that the coefficients computed by the weighted relational model are simply the cardinal of the set of witnesses computed by generalized species; but that is simply not true! (see Section 13.3).

the actual object of study of our field. In a scientific meeting on denotational semantics, there might be some talks about relational models, vectorial semantics, game semantics, Taylor expansion, categorical models, and others; but in essence, all these talks are about this common “full interactive behaviour”. Yet the “full interactive behaviour” is not a formal object. It is intuitively there, but it seems that it can only be accessed indirectly through the various models – one cannot look at it in the eye.

But what if one could, at least to an extent? If nothing else, this would be a precious help in understanding rigorously how the models relate to each other, transporting results, *etc.* It would be a cornerstone of a unified theory of denotational semantics.

1.4 Full Abstraction and Intentional Intensionality

Of course, a model presenting the “full intensional behaviour” would be by nature extremely intensional⁴, and representing the intensional behaviour of programs has not been the historical focus of denotational semantics. Instead, the driving question has long been that of *full abstraction*: the problem of capturing observational equivalence.

Let us make things a bit more precise. Informally, two programs M and N in a fixed language \mathcal{L} are **observationally equivalent** if their behaviour cannot be separated within \mathcal{L} : there is no context $C[]$ – a program of \mathcal{L} with a hole – such that, for instance, $C[M]$ converges while $C[N]$ diverges. In other words, within \mathcal{L} , M and N are completely interchangeable. In that case, we write $M \simeq N$. In many ways, observational equivalence appears to be the “ideal” semantic equivalence between programs; with implications in both theory and practice (for instance, one could hope that an optimization pass in a compiler replaces a program with an optimized, but observationally equivalent, variant). But observational equivalence is hard to establish, due to the infinitary universal quantification on contexts. Accordingly, the driving question in denotational semantics has long been to seek characterizations of observational equivalence: an interpretation $\llbracket - \rrbracket$ is **fully abstract** when we have $M \simeq N$ iff $\llbracket M \rrbracket = \llbracket N \rrbracket$.

In the 90s, much research on denotational semantics was driven by the *full abstraction* problem, of building a fully abstract model for the purely functional language PCF. This was not a practical question, but rather a quest to understand higher-order sequentiality. As an extensional solution remained elusive, this prompted the development of *game semantics for PCF* [Hyland and Ong, 2000, Abramsky et al., 2000, Nickau, 1994]: rather than to find conditions for functions to be sequentially computable, the idea became to find conditions for certain sequential interactive processes called *strategies* to compute *functions*; in other words, to find conditions on strategies to represent the interactive behaviour of purely functional programs. In Hyland-Ong games [Hyland and Ong, 2000], two were necessary: *innocence* and *well-bracketing*. It is debatable whether this did solve the full abstraction problem for PCF⁵, but game semantics had a much better sur-

⁴Meaning, recording information about the *dynamics* of programs, which might not be observable by the execution environment – as opposed to the *extensional* behaviour, *i.e.* input/output.

⁵Nowadays, it seems that the consensus is that it *did not*: the additional intensional behaviour carried by strategies is eliminated by an undecidable quotient. On the other hand, it appears that we cannot do better: there is no effectively presentable fully abstract model for PCF, because observational equivalence in PCF is

prise in store: in the subsequent years, it quickly turned out that lifting *innocence* gave a model of PCF plus local references of ground type [Abramsky and McCusker, 1996, McCusker, 2003], lifting *well-bracketing* gave a model of PCF plus **call/cc**, while lifting both gave a model of PCF with those two effects combined [Murawski, 2007], all fully abstract *and* effectively presentable! In other words, we had *one single* semantic framework for all combinations of state and control, with each condition on strategies corresponding to (the absence of) a certain effect. This was not missed by the community, eventually acknowledged by the *2017 Alonzo Church Award for Outstanding Contributions to Logic and Computation* awarded to the authors of the three *full abstraction* for PCF papers, with the following excerpt from the announcement⁶.

“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.”

To explain this success, we argue that what game semantics has stumbled upon here while seeking full abstraction for PCF, is a glimpse of this “full interactive behaviour” – *at least in a world where everything is sequential deterministic*. This made it possible to understand computational effects in terms of their induced intensional behaviour, and to present four fully abstract models, which could otherwise have appeared completely independent from each other, as four projections of this interactive behaviour.

Of course, game semantics did not stop at the sequential deterministic picture outlined above: there are also fully abstract models for non-deterministic programming languages [Harmer, 1999], probabilistic [Danos and Harmer, 2000] or concurrent languages [Laird, 2001b, Ghica and Murawski, 2008] among others; but they are defined by separate, seemingly incompatible sets of definitions, and they no longer fit in a unified framework as above. For instance the models for non-deterministic or probabilistic computation do not support a notion of *innocence* eliminating state; there is no clear path to a probabilistic extension of the game semantics of concurrent computation, *etc.*

So why do we get this appealing unified picture for sequential deterministic computation, only to lose it immediately after? Where does this come from?

1.5 Plays Considered Harmful

We argue that the limitation causing this is hidden in the very basic definition:

Definition 1.5.1. Consider A a game, inducing a set of moves M_A and **plays** $P_A \subseteq M_A^*$. A **strategy** on A is a non-empty, prefix-closed $\sigma \subseteq P_A$, subject to further conditions.

This seems natural: the interpretation of a type specifies the valid executions, the *plays* of the game; and σ lists the paths that the strategy is prepared to take.

undecidable even with finite basic datatypes and without recursion [Loader, 2001].

⁶<https://eatcs.org/index.php/component/content/article/1-news/2473-the-2017-alonzo-church-award>

But this choice of representation has a cost.

Bounding Opponent's behaviour. When setting up a game semantics, one must usually start by specifying the plays. Plays form the basic backbone of a game semantics – strategies are formed out of plays, composition starts with operations on plays, *etc.*

This may not sound like much, but this choice of plays is perhaps *the* aspect that keeps the most game semantics settings apart and incompatible. Should plays be alternating? Or just locally alternating, as in AJM games? Should they be well-bracketed? Visible? Or even innocent? Some developments carry easily from one choice to another, for instance, essentially the same well-bracketed innocent strategies can be defined on many of those options. But in general, this choice has heavy consequences: it essentially specifies the computational effect available to the ambient environment. Picking a notion of plays is picking a computational universe, it is choosing the *programming language* that the *rest of the world* is written in. If a program was interpreted in a model based on well-bracketed plays, we can say nothing about its execution against **call/cc**. If a program was interpreted in a model based on alternating plays, then we can say nothing about its execution in a concurrent environment.

Should we just adopt once and for all the most general available setting, say based on non-alternating plays? That may be appealing in principle, but many key constructions in game semantics (such as the notion of *P-view*) only make sense for restricted plays.

Losing the branching structure. A play describes a *single* interactive execution, in isolation from all the other ways in which history could have unfolded. A strategy is a set of such plays, with an induced tree structure, that may *branch*. If the first moves that differ are Opponent moves, this simply corresponds to different evaluation choices by the execution environment. In contrast, if the first moves that differ are Player moves, then this witnesses a non-deterministic choice by the program. For instance, the plays

$$\begin{array}{c} \mathbb{B} \\ \mathbf{q}^- \\ \mathbf{tt}^+ \end{array} \quad , \quad \begin{array}{c} \mathbb{B} \\ \mathbf{q}^- \\ \mathbf{ff}^+ \end{array}$$

together inform a strategy **coin** : \mathbb{B} implementing non-deterministic choice.

But the point where branches of the tree of plays separate is only the point where a non-deterministic choice becomes observable, not necessarily where the choice actually happens! For instance, the two terms (for \mathbb{U} a unit type with only one value)

$$x : \mathbb{U} \vdash \mathbf{if\ coin\ then\ } x; \mathbf{tt\ else\ } x; \mathbf{ff} : \mathbb{B}, \quad x : \mathbb{U} \vdash x; \mathbf{coin} : \mathbb{B}$$

allow exactly the same plays, which are all the prefixes of the one below:

$$\begin{array}{c} x : \mathbb{U} \vdash \mathbb{B} \\ \mathbf{q}^- \\ \mathbf{q}^+ \\ \checkmark^- \\ b^+ \end{array}$$

In reality, for the first program, at q^+ the coin has already landed and the result is already determined, but we cannot see that. This has a number of consequences: firstly, being unable to *see* the point of non-deterministic branching means that we are stuck with an *angelic* view of non-determinism, where **if coin then x else \perp** is equivalent to x and the possibility of divergence (with \perp a looping primitive) is ignored⁷.

Secondly, while this missing branching information turns out unnecessary to obtain a fully abstract model for a higher-order language with non-deterministic choice [Harmer and McCusker, 1999], this model no longer fits into a unified landscape as in the sequential deterministic case. In particular, attempting to ban state by a naive extension of *innocence* fails [Harmer, 2004]. Intuitively, *innocence* works by identifying a well-behaved fragment within the model that looks like syntax; but syntax *does* recall the branching information (as illustrated by the two terms above) while strategies as sets of plays cannot. There is no evident way to ban state in Harmer and McCusker’s model [Harmer and McCusker, 1999] and obtain full abstraction for non-deterministic PCF; in this monograph we say that non-determinism and state are *entangled*.

While non-determinism is one possible branching structure, there are other “branching effects” that similarly change the very geometry of execution: two main examples are probabilistic choice and parallelism. As in the non-deterministic case, probabilistic choice and state are entangled in [Danos and Harmer, 2000], while parallelism and state are entangled in [Ghica and Murawski, 2008]; in either case there is no clear notion of innocence banning state while still allowing probabilistic choice and parallelism.

Prior to the line of work reported here, it had already appeared that adopting notions of strategies with explicit branching information, one could indeed obtain a notion of non-deterministic innocence which as desired captures a non-deterministic syntax [Castellan et al., 2014, Tsukada and Ong, 2015]. So it seems that extending the unified picture offered by game semantics beyond the sequential deterministic world is possible in principle, provided we replace plays with adequate mathematical objects giving a true account of various branching structures – we shall see later on that such a mathematical setting may be provided by *event structures* [Winskel, 1986].

Losing positions. Games are incredibly widespread beyond semantics; but on the surface their basic definitions often look rather different from Definition 1.5.1. Typically (say *e.g.* in *parity games* [Arnold and Niwiński, 2001] – or more generally, in games used in verification), a *game* is a graph of *positions*, which are split into *Player* and *Opponent* transitions. A *play* is a path in the graph starting in some initial position, and a *strategy* is a function that to any path ending in a Player position may associate a Player transition. A crucial notion is that of a *positional* strategy, whose behaviour does not depend on the whole play but only on the current position.

Curiously, this notion of *position* is missing from the very basic definitions of traditional game semantics, where instead the fundamental object is that of *plays*. While we

⁷Information about *must-convergence* can be retained by additionally recording *divergences*, as in [Harmer and McCusker, 1999] – but this does not extend to more elaborate resolutions of non-determinism, such as bisimulation or fair-testing equivalences.

may *a posteriori* observe that the two plays

$$\begin{array}{ccc}
 \mathbb{B} & \otimes & \mathbb{B} \\
 \mathbf{q}^- & & \mathbf{q}^- \\
 \mathbf{tt}^+ & & \mathbf{tt}^+ \\
 & \mathbf{q}^- & \mathbf{q}^- \\
 & \mathbf{tt}^+ & \mathbf{tt}^+
 \end{array} \quad (1.1)$$

reach the same *state*, *i.e.* the same set of moves where the two booleans have been evaluated and both have returned \mathbf{tt} , the strategies used in semantics are typically not positional, and positional strategies are not stable under composition.

This is more than a peculiarity – at least if one hopes to link game semantics to other denotational models, since as hinted at in Section 1.3, positions are the key notion to establish a link with the relational model. Alternative, positional foundations for game semantics have existed for a while, such as the *graph games* of [Hyland and Schalk, 2002] and the *asynchronous games* of [Melliès, 2004a]. But by and large, they have not been developed beyond models of fragments of linear logic.

1.6 A Case for Causal Game Semantics

But if we are to remove plays, by what should they be replaced?

At first sight, *removing plays from game semantics* sounds a bit like removing malt from brewing, or removing wool from knitting. Of course committing to a notion of plays will constrain what you may do with them, but you have to live with that, after all plays *are* the basic material with which the game semanticist builds.

However, there is an alternative: *causal game semantics*. The developments presented in this monograph result from a paradigm shift: rather by presenting computational events in sequence, as they could be observed by a specific environment – however powerful it may be, we propose to organize them *causally*: one event appears before another only if this is imposed by a causal dependency, typically internal to the program.

Recovering positions. For instance, the pair $(\mathbf{tt}, \mathbf{tt})$ would be represented by:


$$\begin{array}{ccc}
 \mathbb{B} & \otimes & \mathbb{B} \\
 \mathbf{q}^- & & \mathbf{q}^- \\
 \downarrow & & \downarrow \\
 \mathbf{tt}^+ & & \mathbf{tt}^+
 \end{array}$$

where below the first line that presents the two type components, we show the Hasse diagram of a *partial order*, read from top to bottom. In each type component, the program is prepared to return \mathbf{tt} , under the unique cause that Opponent starts computation by playing \mathbf{q}^- . The comparison with the two diagrams of (1.1) is clear: we have removed *time* and the induced necessity of a choice for Opponent's scheduling. The diagram

only describes what actions the program is prepared to take, and under which conditions. The *positions* are easily accessible as certain down-closed (so *up-closed* in the diagram) sets of observable events; and we should see in the course of the monograph that this allows for a crisp link with relational semantics.

Recovering the branching structure(s). Adopting such partial orders allows us to record explicitly *parallel* branching information. The following diagram represents a strategy evaluating its free variables of unit type \mathbb{U} (admitting only one value) in parallel

$$x : \mathbb{U}, y : \mathbb{U} \vdash \mathbb{U}$$




where the branching on the second line with the two q^+ is *not* a non-deterministic choice (although it might lead to a non-deterministic choice if that parallelism is resolved by a scheduler), but a *parallel* branching: the events corresponding to the two threads for x and y are, as far as the intention of the program is concerned, *causally independent*.

A warning: this diagram does *not* mean that the program *enforces* the independence of x and y , it does not mean that the evaluation of x and y *will* be independent once the context for x and y is known. It only means that the program does not itself impose causal constraints between the two threads. It is possible that the calls to x and y trigger the access to a shared resource, making them interfere. But this interference is then due to the interaction of the program with its environment, not to the program itself.

Much of concurrency theory is based on *interleavings*: an execution of two processes in parallel $P \parallel Q$ is formalized by considering all interleavings of executions of P and Q , hence reducing concurrency to non-determinism by integrating the non-determinism from the scheduler. According to that interpretation, the program above looks non-deterministic as either x or y may be called first in a sequential execution. In contrast, for us (following the so-called *truly concurrent* approach to concurrency theory), this program is considered deterministic: all schedulings eventually reach the same state.

But if that is so, and if branching is parallelism, how can we express actual non-deterministic choice? We do so by adding a *conflict* relation, shown as a wiggly line



indicating that the two moves are *incompatible*, and cannot both occur in a single execution. Altogether, the partial order and the conflict relation form an *event structure* [Winskel, 1986], a well-established notion in concurrency theory.

Lifting constraints on Opponent. Those causal structures, which we call *concurrent strategies*, are developed in Part II. By design, concurrent strategies do not put *any* constraint on Opponent’s behaviour: unlike all play-based game semantics, concurrent strategies do not record how a program is observed by a fixed environment, but only what observable actions the program is prepared to take, and under which conditions.

The purpose of concurrent games is to record very intensional information on the interactive behaviour of programs; information that is typically not observable, however strong the environment may be. In that, it is at odds with much of the earlier work on game semantics for programming languages, which focuses on obtaining full abstraction results and as such, on capturing the information about program behaviour that is observable by specific programming features. In adopting concurrent games, we exchange that for a crisp description of the *full interactive behaviour* of the programs from which all other denotational interpretations can, at least in principle, be obtained as certain projections – a claim that we shall aim to substantiate in this monograph.

Despite our overarching point about the necessity of capturing the full interactive behaviour, some will surely question the point of a semantics not designed to reason on observational equivalence, and recording interactive behaviour that is unobservable, even by a wildly powerful environment. To those, we oppose the following question: is it really reasonable to consider that a program may only interact with an environment written in a *single* programming language, with a fixed observational power? This “fixed world hypothesis”, followed in much of denotational semantics, is in tension with the heterogeneous and distributed nature of modern software. Like any denotational semantics, game semantics is by nature modular with respect to program construction. But the fixed world hypothesis makes it *less* modular in another direction, that of programming language extensions.

One also could wonder if history will remember the full abstraction results as the most important contribution of game semantics. Arguably, there are already signs that probably not. In the recent past, a few developments have been proposed by groups without prior knowledge and experience in game semantics, but leveraging ideas or techniques for game semantics [Stewart et al., 2015, Xia et al., 2020, Chappe et al., 2023, Koenig and Shao, 2020, Vale et al., 2023]. By and large they are motivated by the need for more compositional structures in program certification – they need game semantics for its compositional description of interactive execution, *not* its full abstraction results!

Thus it seems healthy to part with the “fixed world hypothesis” – but it would be better if that did not mean throwing away decades of work and developments in denotational semantics, and the myriad of results and techniques developed. This is, we claim, what concurrent games may eventually achieve, and this monograph contributes some steps.

1.7 Contributions and Outline

In our “concurrent games” line of work, quite a few developments fit into that general line; we survey a few of those at the end of this monograph, in Chapter 13. To form the heart of this document, we had to select a precious few contributions:

Concurrent games. In the published literature, the technical construction of our “concurrent games” framework spans quite a few papers. We first build on Rideau and Winskel’s original set of definitions for a bicategory where both games and strategies are event structures [Rideau and Winskel, 2011]. Our work is based on an extension of this bicategory with symmetry, a framework called *thin concurrent games*, that first appeared (without details and proofs) in a conference paper [Castellan et al., 2015] and then detailed in [Castellan et al., 2019], though depending on an earlier reconstruction of the basic bicategory of concurrent games in [Castellan et al., 2017a]. This makes the chain of dependencies of specific results and lemmas a bit challenging to track, especially given that the notations and terminology have matured over time.

Thus the first contribution of this monograph is a complete, self-contained and hopefully pedagogical presentation of thin concurrent games. It comprises a new presentation of the basic bicategory **CG** without symmetry, the central bicategory **TCG** of *thin concurrent games*; and the description of a number of additional constructions useful to define and study the interpretation of programming languages.

Disentangling parallelism and state. We use the above to give a causal version of one of the most expressive games models out there: Ghica and Murawski’s non-alternating games, fully abstract for *Idealized Concurrent Algol* – a concurrent higher-order language with (local) ground state references [Ghica and Murawski, 2008].

In fact, *ICA* may be regarded as PCF enriched with two programming features: *parallelism* adds the ability to perform computations in parallel, a priori with no interaction or side-channel communication; and *state* is given by ground state local references (and semaphores). We regard *concurrency* not as a primitive computing feature, but as a phenomenon emerging from the combination of parallelism and state. Hence rather than *ICA*, we call the language *Idealized Parallel Algol* (\mathbf{IA}_{\parallel}), *i.e.* Idealized Algol plus parallelism – but it is indeed a concurrent language.

Resting on **TCG** we construct $\rightarrow\text{-}\mathbf{Strat}$, a cartesian closed category in which we describe an interpretation of \mathbf{IA}_{\parallel} . On strategies in $\rightarrow\text{-}\mathbf{Strat}$, we introduce two conditions: *sequentiality*, and *globularity*, which induce two sub-cartesian closed categories $\rightarrow\text{-}\mathbf{Glob}$ (supporting the interpretation of *parallelism*, but not *state*) and $\rightarrow\text{-}\mathbf{Seq}$ (supporting the interpretation of *state*, but not *parallelism*), and a third sub-cartesian closed category $\rightarrow\text{-}\mathbf{SeqGlob}$ (supporting the interpretation of PCF only) where the two conditions are required. We construct **interpretation-preserving functors**:

$$\begin{array}{lll} \rightarrow\text{-}\mathbf{Strat} & \rightarrow & \mathcal{O}\text{-}\mathbf{Strat} & (\text{Theorem 9.3.8}) \\ \rightarrow\text{-}\mathbf{Seq} & \rightarrow & \Downarrow\text{-}\mathbf{Strat} & (\text{Theorem 11.2.18}) \\ \rightarrow\text{-}\mathbf{SeqGlob} & \rightarrow & \Downarrow\text{-}\mathbf{InnWB} & (\text{Theorem 11.3.10}) \end{array}$$

where $\mathcal{O}\text{-}\mathbf{Strat}$, $\Downarrow\text{-}\mathbf{Strat}$ and $\Downarrow\text{-}\mathbf{InnWB}$ are the reference fully abstract models for \mathbf{IA}_{\parallel} , \mathbf{IA} and PCF respectively, phrased in terms of traditional game semantics – so that $\rightarrow\text{-}\mathbf{Strat}$, $\rightarrow\text{-}\mathbf{Seq}$ and $\rightarrow\text{-}\mathbf{SeqGlob}$ are also fully abstract. Finally, we prove a new **finite definability result** for $\rightarrow\text{-}\mathbf{Glob}$ with respect to PCF plus *parallelism*.

Altogether, these results **disentangle** parallelism and state, in the sense above.

Relational collapse. Setting up the tools required for this is quite the journey. A particularly significant step encountered along the way is that *visibility*, one of the constituents of *globularity*, ensures that the composition of strategies induces no deadlocks so that we have an **interpretation-preserving functor to the relational model** (Corollary 10.4.15) – a modern account of the link between game semantics with the relational model, in line with our motto of bridging denotational models.

A uniform presentation of pointer games. Of course, a challenge in constructing the interpretation-preserving functors above to traditional games models, is that they refer to models constructed in a number of different papers, working with different notations, definitions that differ in subtle but significant ways... To achieve our goals, it was necessary to first recast the various games models of interest, which we attempted to do in an introductive and pedagogical way.

Outline. In the next chapter, we introduce a few mathematical conventions and notations used throughout the monograph.

Part I contains our introduction to pointer games. It focuses on alternating pointer games, covering innocence, well-bracketing, the full abstraction results for PCF and Idealized Algol. Then, it also presents non-alternating pointer games and the fully abstract model for IA_{\parallel} . We mention a few other topics along the way: *e.g.* the *semantic cube*, and the alternative approach to uniformity via *copy indices*.

Part II contains our detailed construction of concurrent games. We start with the basic bicategory **CG**, then add symmetry to obtain *thin concurrent games* (**TCG**). Finally, we construct additional structure aiming for (essentially) a cartesian closed category, as required by the interpretation of (call-by-name) programming languages.

Part III contains all the developments required to disentangle parallelism and state. It contains the interpretation of IA_{\parallel} , the developments of *sequentiality* and *globularity* along with all the proofs of stability under composition. It contains all our interpretation-preserving functors, including that to the relational model. Finally, it contains a proof of finite definability for globular strategies.

Part IV surveys other developments in concurrent games fitting in our general pattern of bridging denotational models; as well as perspectives and open problems.

Finally, Part V contains some appendices.

Chapter 2

General Preliminaries

In this first technical chapter, we introduce some *preliminaries* to the *preliminaries*: first, we fix some of the general notations and conventions used in this monograph. Then, we introduce our programming languages of interest. Finally, we will recall definitions for a few of the basic denotational semantics concepts used later on.

2.1 Mathematical Language

To start off, let us introduce a few notations that we will use throughout this monograph.

Sets. If X is a set, then we write $\mathcal{P}(X)$ for the **powerset** of X , *i.e.* the set comprising all subsets of X . If X and Y are sets, we write $X + Y$ for the **tagged disjoint union** of X and Y , defined as $X + Y = \{(1, x) \mid x \in X\} \cup \{(2, y) \mid y \in Y\}$. We are careful to distinguish this with $X \uplus Y$ which is the usual set-theoretic union, when it is known or assumed to be disjoint – for instance, we may write $X + Y = \{1\} \times X \uplus \{2\} \times Y$.

Sequences. If X is a set, then X^* denotes the set of **words** or **finite sequences** of elements of X . We write $\varepsilon \in X^*$ for the empty sequence, and \sqsubseteq denotes the **prefix ordering** on finite sequences.

Functions. If X and Y are sets, we write $f : X \rightarrow Y$ to mean that f is a function from X to Y (if X and Y are understood from the context to be objects in a specific category \mathcal{C} , this might mean that f is a morphism of \mathcal{C} from X to Y instead). Likewise, we write $f : X \rightharpoonup Y$ to mean that f is a partial function from X to Y . In both cases, we write the application of f to $x \in X$ as $f(x)$ or sometimes simply as fx . If $f : X \rightarrow Y$ is a bijection, we write $f : X \simeq Y$.

Categories. We assume familiarity with the basic notions of category theory and categorical logic. We write **Set** for the category of sets and functions. If \mathcal{C} is a category, we

write C_0 for its set of objects; if A and B are objects of C , we write $C(A, B)$ for the set of morphisms from A to B . If $f \in C(A, B)$ is an isomorphism, we write $f : A \cong B$. For an isomorphism in **Set**, we tend to prefer the term *bijection* and the notation $f : A \simeq B$, but use *isomorphism* and write $f : A \cong B$ in the presence of further structure – for instance if A and B are partial orders, and f is an order-isomorphism.

We write \top for the terminal object of C , if it exists. If C is a cartesian category, then we write $A \& B$ for the cartesian product of A and B . We write $\pi_A : A \& B \rightarrow A$ and $\pi_B : A \& B \rightarrow B$ for the two projections, and $\langle f, g \rangle : X \rightarrow A \& B$ for the pairing of $f : X \rightarrow A$ and $g : X \rightarrow B$. We will also use n -ary versions of these constructions.

If C is cartesian closed, we write $A \Rightarrow B$ for the arrow object of A and B and $\text{ev}_{A,B} : (A \Rightarrow A) \& A \rightarrow B$ for the evaluation morphism. If $h : X \times A \rightarrow B$, we write $\Lambda(h) : X \rightarrow A \Rightarrow B$ for its currying.

Common abbreviations. We shall use the following abbreviations in the text: *iff* for “if and only if”, *lhs* for “left hand side”, *rhs* for “right hand side”, *s.t.* for “such that” and *w.r.t.* for “with respect to”.

2.2 IA_{\parallel} and its Fragments

The contributions of this monograph are not about a programming language in particular: we view them as methodological contributions in denotational semantics, that may apply, in principle, to a variety of languages – the past body of work on game semantics has indeed proved that the insights gained when studying programming features in a specific setting could, to a large extent, be transported to other settings (*e.g.* the notion of *innocence*, originally intended to ban state in a call-by-name language, turned out to apply just as well to ban state in a call-by-value language [Honda and Yoshida, 1999]).

But we do need to fix target programming languages to illustrate our game semantics contributions, and to serve as target for our technical developments. For that, we settled on *Idealized Parallel Algol* (IA_{\parallel}), a call-by-name concurrent higher-order programming language with ground local references. The main reason we adopted IA_{\parallel} is because among the languages that are well-studied in game semantics, this is one of the most expressive. Removing parallelism and/or state allows one to recover languages that are *also* very well-studied, though the corresponding models do not fit in a uniform semantic picture. If we want to prove full abstraction for the different fragments of IA_{\parallel} , it “suffices” to relate to the existing fully abstract models, saving us some work.

The core of IA_{\parallel} is PCF, the prototypical higher-order sequential language introduced by Plotkin [Plotkin, 1977]: a call-by-name simply-typed λ -calculus with booleans, natural numbers and their combinators, plus recursion. We consider two extensions of PCF: firstly, *Idealized Algol* (IA) is PCF extended with *shared state*, under the form of (mainly) ground type local references. Secondly, *Parallel PCF* (PCF_{\parallel}) extends PCF with a primitive for pure parallel evaluation, without communication. IA_{\parallel} , obtained by putting all these together, is rather complex: in particular it is a *concurrent* language, as the combination of parallelism and shared memory induce *racy behaviour*.

In the rest of this section, we introduce formally the syntax of IA_{\parallel} and its fragments. A disclaimer: IA_{\parallel} and its fragments are not claimed to be realistic or in any way faithful to actual programming languages; nor do we think that realistic programming languages should be inspired by IA_{\parallel} . It is a toy language, whose only purpose is to put together certain programming aspects in a controlled environment where they can be studied in isolation; like a physicist studying properties of matter in a vacuum. Likewise, the reader should not see this work as a “study of the semantics of IA_{\parallel} ” or its fragments. Again, our contribution is semantic and methodological, IA_{\parallel} being merely a way to illustrate it and provide a target that is challenging but achievable.

2.2.1 Syntax of IA_{\parallel}

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

$$A, B ::= \mathbb{U} \mid \mathbb{B} \mid \mathbb{N} \mid A \rightarrow B \quad \text{PCF} \\ | \mathbb{V} \mid \mathbb{S} \quad \text{+state}$$

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 state, \mathbb{V} is a type for *references* storing natural numbers, while \mathbb{S} 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.

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 form $\Gamma \vdash M : A$ with Γ a context and A a type. In addition to Figure 2.1, we consider present an exchange rule letting us 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 \mathbb{V} may be obtained as

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

The **bad variable** and **bad semaphore** constructs **mkvar** and **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 \mathbb{V} but not behaving as actual variables. Full abstraction results in the concerned games models [Abramsky and McCusker, 1996, Ghica and Murawski, 2008] require a corresponding syntactic construct **mkvar** allowing one to *form* bad variables by appending arbitrary read and write methods¹. The same holds for semaphores.

¹A very nice result by McCusker is that in fact, the fully abstract games model for IA is also fully abstract without bad variables [McCusker, 2003]!

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

Figure 2.1: Typing rules for \mathbf{IA}_{\parallel}

2.2.2 Operational Semantics of IA_{\parallel}

In this monograph, we shall not refer explicitly to the operational semantics for IA_{\parallel} , as we will only interact with those through already established denotational models. However, for the sake of completeness, we include here the definition.

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 give a small-step operational semantics, following [Ghica and Murawski, 2008]. 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. A **configuration** is a term together with a store, formally a pair $\langle M, s \rangle$ where s is a store with $\text{dom}(s) = \{\ell_1, \dots, \ell_n\}$ and $\Sigma \vdash M : A$ with $\Sigma = \ell_1 : \mathbb{V}, \dots, \ell_i : \mathbb{V}, \ell_{i+1} : \mathbb{S}, \dots, \ell_n : \mathbb{S}$.

Execution is formalized via reductions between configurations, with shape

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

where $\text{dom}(s) = \text{dom}(s')$; we write \rightsquigarrow^* for the reflexive transitive closure. If $\vdash M : \mathbb{X}$, we write $M \Downarrow v$ if $\langle M, \emptyset \rangle \rightsquigarrow^* \langle v, \emptyset \rangle$ for some value v or simply $M \Downarrow$ leaving v implicit. We give in Figure 2.2 the reduction rules; separated in *basic rules* listing the basic cases, and the *context rules* specifying the next available redexes. For rules not interacting with the state, we omit the state component – it is simply left unchanged by stateless basic reductions, and propagated upwards by stateless context rules.

2.2.3 Expressiveness of IA_{\parallel}

We shall illustrate the expressiveness of the language via its fragments.

The core language: PCF. Introduced in [Plotkin, 1977], PCF is the paradigmatic (call-by-name) language for higher-order recursive computation.

PCF is of course Turing-complete, and may be used to program any (higher-order) computable function. For instance, the Ackermann function may be defined with

$$\mathcal{Y}(\lambda A^{\mathbb{N} \rightarrow \mathbb{N} \rightarrow \mathbb{N}} n^{\mathbb{N}} m^{\mathbb{N}}. \text{if}(\text{iszero } n)(\text{succ } m)(\text{iter}(A(\text{pred } n)) m))$$

with $\text{iter} = \lambda f^{\mathbb{N} \rightarrow \mathbb{N}} x^{\mathbb{N}}. \mathcal{Y}(\lambda I^{\mathbb{N} \rightarrow \mathbb{N}} n^{\mathbb{N}}. \text{if}(\text{iszero } n) x (f(I(\text{pred } n))))$; showing how we can use recursion on higher-order functions in PCF. We define additional combinators in PCF that will eventually be used in the monograph. First, there is a **divergence**

$$\overline{\Gamma \vdash \perp_A : A}$$

for any type A , defined as any looping program, for instance via $\perp_A = \mathcal{Y}_A(\lambda x^A. x)$. We also add syntactic sugar for an **equality test** for all ground types, with rule

$$\frac{\Gamma \vdash M : \mathbb{X} \quad \Gamma \vdash N : \mathbb{X}}{\Gamma \vdash M =_{\mathbb{X}} N : \mathbb{B}}$$

<p>Basic red. for PCF</p> $ \begin{aligned} (\lambda x^A. M) N &\rightsquigarrow M[N/x] \\ \text{skip}; N &\rightsquigarrow N \\ \text{if tt } M N &\rightsquigarrow M \\ \text{if ff } M N &\rightsquigarrow N \\ \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>Basic reductions for state</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>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>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>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} \quad \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>Stateful context rules</p> $ \begin{aligned} &\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{aligned} $	

Figure 2.2: Operational semantics of $\text{IA}_{//}$

returning **tt** iff the two terms yield the same value. For $\mathbb{X} = \mathbb{U}$ we may define $M =_{\mathbb{U}} N$ simply as $M; N; \mathbf{tt}$. Likewise, we set $M =_{\mathbb{B}} N$ as **if** $M N$ (**if** N **ff** **tt**), and $\Gamma \vdash M =_{\mathbb{N}} N : \mathbb{B}$ similarly, with the obvious recursive program.

Next we introduce a n -ary case construct branching on all values of ground types. If $V = \{v_1, \dots, v_n\}$ is a finite set of values of ground type \mathbb{X} , we set

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

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

This makes crucial use of the **let** construct, ensuring that M is evaluated only once. Note that the historical presentation of PCF has no **let** construct; however including it makes for a more complete language of study with some control over the evaluation order, and also allows for a neater correspondence with the game semantics.

Idealized Algol. *Idealized Algol*, or **IA** for short, corresponds in Figure 2.1 to PCF + state. It is used in the game semantics community (starting with the seminal paper [Abramsky and McCusker, 1996]) as the (call-by-name) paradigmatic programming language with ground type references: one can store an integer (though nothing much would change if one could store other finite datatypes), but one cannot store a function, or a reference. **IA** subsumes PCF, so all the programs and syntactic sugar mentioned above may still be defined here. But **IA** also supports imperative programming, *e.g.*

$$\begin{array}{l} \lambda n^{\mathbb{N}}. \text{newref } r \text{ in} \\ \quad \text{newref } i \text{ in} \\ \quad \quad r := 1; \\ \quad \quad i := n; \\ \quad \quad \text{while (not (iszero } !i)) \text{ do} \\ \quad \quad \quad r := (!i) * (!r); \\ \quad \quad \quad i := \text{pred } (!i) \\ \quad \quad \text{done;} \\ \quad \quad !r \end{array}$$

where **while** M **do** N **done** is syntactic sugar for $\mathcal{Y}(\lambda x. \text{if } M (N; x) \text{ skip})$; and **not** and the arithmetic operation $*$ are defined in the obvious way.

Beyond first-order imperative programming, **IA** supports a combination of higher-order and state letting us write programs such as the following *strictness test*

$$\begin{array}{l} \lambda f^{\mathbb{U} \rightarrow \mathbb{U}}. \text{newref } r \text{ in } f(r := 1); \\ \quad \text{if (iszero } (!r)) \text{ ff tt} \\ : (\mathbb{U} \rightarrow \mathbb{U}) \rightarrow \mathbb{B} \end{array}$$

replying **ff** when applied to a constant function, and **tt** if the function calls its argument.

Our presentation of IA differs from the historical one in two inessential ways. First, as for PCF the **let** construct is usually missing from IA. But our **let** evaluates only terms of ground type, which is definable² in IA by letting **let** $x = M$ **in** N be sugar for

$$\mathbf{newref} \, r \, \mathbf{in} \, r := M; (\lambda x. N) (!r),$$

illustrating an important aspect of IA references: they store *values* and not *computations*, so an expression $r := M$ must evaluate M fully before storing the value.

The second inessential difference is that our version of IA includes semaphores – which we need, because they are unavoidable to get full abstraction in the presence of parallelism also. For sequential IA, semaphores of course are definable via:

$$\begin{aligned} \mathbf{newsem} \, s \, \mathbf{in} \, M &= \mathbf{newref} \, s \, \mathbf{in} \, M \\ \mathbf{grab} \, M &= \mathbf{if} \, (\mathbf{iszero} \, !M) \, (M := 1) \, \perp \\ \mathbf{release} \, M &= \mathbf{if} \, (\mathbf{iszero} \, !M) \, \perp \, (M := 0), \end{aligned}$$

which is fine only because IA is sequential: if a program attempts to grab a semaphore which is not free (or dually, to release a semaphore which is free), then computation will hang as there is no other thread to free the semaphore while the program waits.

Parallel PCF. *Parallel PCF*, or PCF_{\parallel} for short, corresponds in Figure 2.1 to PCF + parallelism. To PCF it adds mere parallel computation. This is made official by the operational semantics of Figure 2.2, which will indeed evaluate both branches of a parallel **let in parallel**. But keep in mind that this parallelism is not observable as there is no primitive in PCF_{\parallel} allowing one to observe the evaluation order. Because of that, giving denotational semantics to PCF_{\parallel} in isolation is a fairly trivial endeavour: any denotational semantics for standard PCF will do, interpreting the parallel **let** as

$$\mathbf{let} \, x_1 = M_1 \, \mathbf{in} \, (\mathbf{let} \, x_2 = M_2 \, \mathbf{in} \, N),$$

but in this monograph we do not wish one denotational semantics for each fragment of IA_{\parallel} , we want all fragments to be interpreted in the *same* denotational universe!

Idealized Parallel Algol. *Idealized Parallel Algol*, or IA_{\parallel} for short, corresponds in Figure 2.1 to PCF + parallelism + state. It is used in the game semantics community [Ghica and Murawski, 2008] as the (call-by-name) paradigmatic higher-order concurrent programming language with shared memory and semaphores.

To IA, our version of IA_{\parallel} adds parallelism, under the form of

$$\frac{\Gamma, x_1 : \mathbb{X}, x_2 : \mathbb{X} \vdash M : \mathbb{Y} \quad \Gamma \vdash N_1 : \mathbb{X} \quad \Gamma \vdash N_2 : \mathbb{X}}{\Gamma \vdash \mathbf{let} \left(\begin{array}{l} x_1 = N_1 \\ x_2 = N_2 \end{array} \right) \mathbf{in} \, M : \mathbb{Y}}$$

²Up to observational equivalent; interestingly the two do not have the same causal behaviour!

a parallel let operation which evaluates N_1 and N_2 independently. We refer to this as *parallelism* and not *concurrency* because this **let** only adds parallel evaluation; in our view concurrency emerges from parallel computation plus shared resources.

The usual presentation of IA_{\parallel} instead has a parallel composition operator

$$\frac{\Gamma \vdash M : \mathbb{U} \quad \Gamma \vdash N : \mathbb{U}}{\Gamma \vdash M \parallel N : \mathbb{U}}$$

which may be defined as syntactic sugar, via $M \parallel N = \text{let } \begin{pmatrix} x & = & M \\ y & = & N \end{pmatrix} \text{ in skip}$.

In this work, we adopt the parallel let rather than merely parallel composition because we also consider parallelism without shared state; but without state $M \parallel N$ is not a relevant construction as M, N produce no value and can only communicate with the ambient program via the shared memory. Again, this is not a significant change with respect to the usual formulation of IA_{\parallel} , as the same behaviour may be obtained with

$$\text{let } \begin{pmatrix} x_1 & = & N_1 \\ x_2 & = & N_2 \end{pmatrix} \text{ in } M = \text{newref } r_1, r_2 \text{ in } (r_1 := N_1) \parallel (r_2 := N_2); \\ (\lambda x_1 x_2. M) (!r_1) (!r_2)$$

if \parallel is primitive rather than the parallel let.

IA_{\parallel} is a very expressive language. The most impactful thing to notice is that the combination of parallelism and shared state make it non-deterministic: one can define

$$\vdash \text{choice} = \text{newref } r \text{ in let } \begin{pmatrix} x & = & (r := 1) \\ y & = & !r \end{pmatrix} \text{ in iszero } y : \mathbb{B}$$

exploiting a race in the memory to generate non-deterministic behaviour.

2.3 Generalities on Semantics

Finally, we include some reminders and settle on some terminology and notations regarding observational equivalence, and denotational semantics. In this section, we use \mathcal{L} as the *ambient programming language* – it can refer to PCF, IA, PCF_{\parallel} , or IA_{\parallel} .

Observational equivalence. A \mathcal{L} -context for the judgment $\Gamma \vdash A$ is a term $C[]$ of \mathcal{L} with a hole written $[]$, such that for any $\Gamma \vdash M : A$ in \mathcal{L} , we have $\vdash C[M] : \mathbb{U}$ obtained by the (non capture-avoiding) substitution of $[]$ with M .

Two terms $\Gamma \vdash M, N : A$ of \mathcal{L} are **\mathcal{L} -observationally equivalent** iff

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

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} .

Denotational semantics. Observational equivalence is ultimately an *operational* notion of program identity; it rests on the evaluation relation $M \Downarrow$, itself defined via the operational semantics. An alternative to study program identity is given by *denotational semantics*, embedding programs in a syntax-independent mathematical universe.

Denotational semantics is usually formulated *categorically*: the language \mathcal{L} is interpreted in a category \mathcal{C} equipped with sufficient structure. More precisely, a denotational semantics in \mathcal{C} consists in several components. First, there is an interpretation function

$$\llbracket - \rrbracket : \mathbf{Types} \rightarrow \mathcal{C}_0$$

for *types* – this usually exploits the categorical structure, with *e.g.* $\llbracket A \rightarrow B \rrbracket = \llbracket A \rrbracket \Rightarrow \llbracket B \rrbracket$ resting on a cartesian closed structure. This interpretation is extended to

$$\llbracket - \rrbracket : \mathbf{Contexts} \rightarrow \mathcal{C}_0$$

an interpretation of contexts, defined simply via $\llbracket \Gamma \rrbracket = \&_{(x_i : A_i) \in \Gamma} \llbracket A_i \rrbracket$. Now, we have

$$\llbracket - \rrbracket : \mathbf{Terms}(\Gamma \vdash A) \rightarrow \mathcal{C}(\llbracket \Gamma \rrbracket, \llbracket A \rrbracket)$$

the interpretation of terms, defined by structural induction on the typing rules of \mathcal{L} – for the simply-typed λ -calculus this interpretation typically follows a cartesian closed structure on \mathcal{C} , while the primitives of \mathcal{L} (*e.g.* natural numbers, booleans, manipulations of state, *etc*) are interpreted by well-chosen morphisms in \mathcal{C} . By a **denotational model** of \mathcal{L} , we mean a category \mathcal{C} along with the interpretation function $\llbracket - \rrbracket$ – sometimes focusing on \mathcal{C} and leaving the latter implicit.

In this monograph, we assume the reader is familiar with basic categorical logic, including the interpretation of the simply-typed λ -calculus in a cartesian closed category. For details, the reader is referred to [Lambek and Scott, 1988].

Adequacy and full abstraction. A denotational model should reflect the operational evaluation. We say that an interpretation in a model \mathcal{C} is **computationally adequate**, or simply *adequate* for short, if for all $\vdash M : \mathbb{X}$ a closed term of ground type,

$$M \Downarrow v \quad \Leftrightarrow \quad \llbracket M \rrbracket = \llbracket v \rrbracket,$$

i.e. the denotational model faithfully predicts the evaluation of a program to a value. For general reasons, an adequate model is automatically *sound* with respect to observational equivalence, *i.e.* if $\llbracket M \rrbracket = \llbracket N \rrbracket$ then $M \simeq_{\mathcal{L}} N$; it is **fully abstract** if the converse also holds: if $M \simeq_{\mathcal{L}} N$ then we actually have $\llbracket M \rrbracket = \llbracket N \rrbracket$.

Intensional full abstraction. Given a denotational model \mathcal{C} , we may replay in \mathcal{C} the syntactic definition of observational equivalence: we fix an object for observations, typically $\llbracket \mathbb{U} \rrbracket$ (still written \mathbb{U} to alleviate notations). An *observation* on A is then any

$$\alpha \in \mathcal{C}(A, \mathbb{U})$$

playing the role of \mathcal{L} -contexts: $f, g \in C(A, B)$ are **observationally equivalent** iff³

$$f \simeq_C g \quad \Leftrightarrow \quad \text{for all } \alpha \in C(A \Rightarrow B, \mathbb{U}), \alpha \circ f = \alpha \circ g$$

like syntactic observational equivalence. Quotienting C by this yields a new model, and C is **intensionally fully abstract** when this quotiented model is fully abstract.

The term “intensional full abstraction” was originally proposed by Abramsky, Jagadeesan and Malacaria [Abramsky et al., 2000] to better describe the “full abstraction” results for PCF. Intensional full abstraction is *not* full abstraction: it is full abstraction of the quotiented model, but this quotiented model is usually not very useful as it is obtained non-constructively – in particular, a model constructed in this way rarely helps directly in reasoning on observation equivalence.

In contrast, *intensional full abstraction* keeps the focus on the model pre-quotient (which, at least in the case of game semantics, turned out to be more important). There we understand it as meaning that the “intensional behaviour” of morphisms of C captures that of \mathcal{L} , or at least is sufficiently close that the morphisms of C have the *same distinguishing power* as contexts of \mathcal{L} through the interpretation⁴ – there is no “abstraction leaks”. Interestingly, it is in that sense that the term “full abstraction” is often understood in compilation – see *e.g.* [New et al., 2016].

In this monograph, we try to be careful in distinguishing *full abstraction* from *intensional full abstraction*, though we might occasionally slip and omit *intensional*.

³This definition is adequate when $C(\top, \mathbb{U})$ has only two elements, for convergence and divergence. In some settings such as in concurrent games, this must be adapted to match syntactic observational equivalence.

⁴Curien calls this property *definable separation* [Curien, 2007].

Part I

Pointer Game Semantics

Introduction to Part I




Figure 2.3: As a low-pressure system moves in from the west, we can expect widespread pointer showers to develop by Chapter 3, and these showers are likely to continue on and off throughout Part I.

The first part of this monograph is an introduction to “pointer game semantics”, a phrasing intended to include both the traditional (sequential deterministic) Hyland-Ong games [Hyland and Ong, 2000], and the non-alternating games by Ghica and Muraski [Ghica and Murawski, 2008], which of course have much in common with the former.

There are two pedagogical difficulties when learning game semantics. The first is coming up with an intuition as to what *plays* and *strategies* represent in relation with the interactive behaviour of the program – the operational meaning of well-bracketing, innocence, *etc.* Understanding this is not enough to *work* with game semantics, but it suffices to see the scientific value and role of game semantics. The second difficulty is understanding the denotational machinery: the composition of strategies, the cartesian closed structure, *etc.* We organized our introduction to pointer games addressing these two difficulties separately, so that a reader wanting to get the *point* of game semantics without mastering the combinatorics of the categorical structure can do so.

Accordingly, Chapter 3 introduces pointer games and its main notions and achievements, independently of the mechanics of the interpretation of programs. The content is technical and mathematically precise, but omits the definition of the interpretation and the corresponding categorical structure. The chapter starts by introducing the basic intuitions behind game semantics, at first without committing to any particular technical setting, then motivates and introduces the main concepts of pointer games: *arenas*, *plays with pointers*, *strategies*, *etc.*; resting on operational intuitions. Still relying on operational intuitions, we show how certain patterns in plays witness the use of computational effects such as state or control, and use this to motivate the notions of well-bracketing and innocence, and showcase the associated full abstraction results. Finally, we mention an alternative to pointers, *copy indices*, as in AJM games [Abramsky et al., 2000].

Then, Chapter 4 addresses the denotational interpretation of programs. We define the composition of strategies, constructing a cartesian closed category $\Downarrow\text{-Strat}$. We define the interpretation of IA in $\Downarrow\text{-Strat}$, and wrap up with a few complements.

Finally, Chapter 5 introduces *non-alternating pointer games* in the sense of Ghica and Murawski [Ghica and Murawski, 2008]. The exposition roughly following the pattern of the previous two chapters: first we define non-alternating strategies and illustrate how they capture the interactive behaviour of programs of IA_{\parallel} ; then we build the corresponding cartesian closed category $\cup\text{-Strat}$, and define the interpretation of IA_{\parallel} .

Chapter 3

First Steps in Game Semantics

In this first chapter, we give an introduction to game semantics – focusing in particular on so-called HO or *Hyland-Ong* games¹ [Hyland and Ong, 2000] presenting a selection of developments and landmark historical results. The presentation is technical, in that we give mathematically precise definitions and statements. However, we focus on operational intuitions, (temporarily) remaining elusive on the main source of complexity in game semantics: the categorical structure and denotational definition of the interpretation of programs. This, in turn, shall be the topic of Chapter 4 – so altogether, Chapters 3 and 4 give an introduction to sequential alternating game semantics.

Of course, as an introduction to game semantics, this presentation is far from complete. Many developments in game semantics do not fit in the framework presented here; we also focus on semantics of programming languages and essentially omit the works on logics and proof systems, on the other side of the Curry-Howard correspondence.

3.1 Executions, Plays, and Strategies

This section aims to convey some intuitions, and to motivate and gradually put into place our choices of mathematical formalizations for the basic ingredients of game semantics.

3.1.1 Dialogues, Plays, and Simple Games

Game Semantics is designed as a way to give a clean mathematical description of the behaviour of *open* programs, as a back-and-forth interaction with an unspecified execution environment. More precisely, a program is interpreted as an aggregate of all

¹Those are often called Hyland-Ong/Nickau or *HON* games, also acknowledging Nickau, who came up independently with a model of PCF [Nickau, 1994] that is essentially the same as Hyland and Ong’s. That is fair, but then one should also acknowledge the work of Coquand, who also proposed independently definitions having much in common with those [Coquand, 1995] (though not for PCF). Rather than the cumbersome *HONC* games, we prefer to stick with *HO* games, which is well-recognized in the community and refers to the specific set of technical definitions introduced by Hyland and Ong.




Figure 3.1: First-order




Figure 3.2: Higher-order




Figure 3.3: Repetition

its possible interactive behaviour with an execution environment. This is formalized as a two-player game between **Player** – which plays for the program under scrutiny; and **Opponent** – which plays the execution environment. Execution unfolds as those two players exchange *moves*, each of which represents an observable computational event.

First-order. As a first example, let us consider the following open PCF program:

$$x : \mathbb{N}, y : \mathbb{N} \vdash x + y : \mathbb{N}$$

For the sake of this discussion, we assume that the addition is computed on a left-to-right basis. From the point of view of standard operational semantics, there is no computation to perform here: the program is in *normal form*. *Game semantics* start where usual operational semantics stop, by *asking the environment the value of x* . One may then imagine a dialogue between the program and its environment, as follows:

O : “What is the value of $x + y$?”
P : “Evaluation is stuck: what is the value of x ?”
O : “The value of x is 12.”
P : “Evaluation is stuck: what is the value of y ?”
O : “The value of y is 30.”
P : “The value of $x + y$ is 42.”

In game semantics, such a sequence is called a **play**, and is typically drawn in a diagram as in Figure 3.1. By convention, we assign to Player the positive polarity $+$, and to Opponent the negative polarity $-$. In this diagram and all game semantics diagrams to come, we label moves with their polarity. The flow of time is from top to bottom, and each column corresponds to a component of the interface of the program with its environment. Variable calls are represented with the symbol \mathbf{q} which stands for “Questions”. All other moves correspond to returns, and are called “Answers”.

This dialogue is one possible interaction between our program of interest and an execution environment; there are countably many of those, for each values for x and y . Such interactions only display the events observable by the environment: the program $(x + 0) + y$ has exactly the same plays as $x + y$, the internal computation is ignored.

Higher-order. Let us now consider another example:

$$f : \mathbb{N} \rightarrow \mathbb{N}, x : \mathbb{N} \vdash f(x) : \mathbb{N}$$

One could expect Player to first prompt the environment for an actual function f from integers to integers. It is a central feature of game semantics that we instead have:

O : “What is the value of $f(x)$?”
P : “Evaluation is stuck: what is the output value of f ?”
O : “What is the input value of f ?”
P : “Evaluation is stuck: what is the value of x ?”
O : “The value of x is 50.”
P : “The input value of f is 50.”
O : “The output value of f is 42.”
P : “The value of $f(x)$ is 42.”

which, in the usual style of game semantics, would appear as in Figure 3.2: we interact with f in a lazy fashion, never fully grasping the infinite object behind. In this way, game semantics reduces higher-order evaluation to an exchange of first-order tokens.

Simple games. But what are, formally, the plays of Figures 3.1 and 3.2?

In traditional game semantics, each type (or typing judgment) yields by interpretation a **game** whose rules describe the valid open executions on the interface given by that type. Exact formalizations vary wildly in the literature and we will see several of those in this monograph; as a stepping stone we first present what must be the simplest notion of game considered in the game semantics literature, **simple games** [Hyland, 1997]:

Definition 3.1.1. A *simple game* comprises $A = (|A|, \text{pol}_A, P_A)$ with $|A|$ a set of *moves*, $\text{pol}_A : |A| \rightarrow \{-, +\}$ a *polarity function*, and $P_A \subseteq |A|^*$ a set of *plays* such that:

- non-empty: $\epsilon \in P_A$,
- prefix-closed: for all $s \in P_A$, for all $s' \sqsubseteq s$, we have $s' \in P_A$,
- alternating: for all $s = s_1 \dots s_n$, for all $1 \leq i \leq n - 1$, $\text{pol}_A(s_i) \neq \text{pol}_A(s_{i+1})$,
- negative: for all $s = s_1 \dots s_n$, if $n \geq 1$ then $\text{pol}_A(s_1) = -$.

For instance, a game \mathbb{B} for the booleans would have $|\mathbb{B}| = \{\mathbf{q}, \mathbf{tt}, \mathbf{ff}\}$, with $\text{pol}_{\mathbb{B}}(\mathbf{q}) = -$, $\text{pol}_{\mathbb{B}}(\mathbf{tt}) = \text{pol}_{\mathbb{B}}(\mathbf{ff}) = +$, and $P_{\mathbb{B}} = \{\epsilon, \mathbf{q}, \mathbf{q}\mathbf{tt}, \mathbf{q}\mathbf{ff}\}$. As above, we shall often write moves with their polarities, as in $|\mathbb{B}| = \{\mathbf{q}^-, \mathbf{tt}^+, \mathbf{ff}^+\}$ and $P_{\mathbb{B}} = \{\epsilon, \mathbf{q}^-, \mathbf{q}^-\mathbf{tt}^+, \mathbf{q}^-\mathbf{ff}^+\}$. However do keep in mind that this is merely for convenience: polarities are given by pol_A , and not intrinsic to moves, *i.e.* to the elements of $|A|$.

A program $M : A$ is then interpreted as a *strategy for Player* on the game for A :

Definition 3.1.2. A *strategy* $\sigma : A$ on simple game A is a subset $\sigma \subseteq P_A$ satisfying:

- non-empty: $\varepsilon \in \sigma$,
- prefix-closed: for all $s \in \sigma$, if $s' \sqsubseteq s$, then $s' \in \sigma$,
- receptive: for all $s \in \sigma$, if $sa^- \in P_A$, then $sa^- \in \sigma$,
- deterministic: for all $sa_1^+, sa_2^+ \in \sigma$, we have $a_1 = a_2$.

A strategy aggregates all executions that the program is prepared to make against any execution environment; prescribing, in each state, the next observable action. When presented an Opponent-ending play $sa^- \in P_A$, a strategy $\sigma : A$ checks whether sa^- appears in σ ; and if so, if there is some extension sa^-b^+ . If that is the case, σ plays b^+ . For instance, the strategy for the constant \mathbf{tt} would be given by the set:

$$\{\varepsilon, \mathbf{q}^-, \mathbf{q}^-\mathbf{tt}^+\} : \mathbb{B}.$$

Simple games are appealing; they seem to capture exactly the intuitions so far, without useless mathematical arabesque. But their expressiveness is limited in various ways, making them seldom used in game semantics. We review next a major constraint, playing a central role in the design of essentially all game semantics formalisms.

3.1.2 Replication and Thread Indexing

We consider the term $x : \mathbb{N} \vdash x + x : \mathbb{N}$; whose game semantics admits:

- O** : “What is the value of $x + x$?”
- P** : “Evaluation is stuck: what is the value of x ?”
- O** : “The value of x is 12.”
- P** : “Evaluation is stuck: what is the value of x ?”
- O** : “The value of x is 30.”
- P** : “The value of $x + x$ is 42.”

drawn in game semantics style in Figure 3.3. This illustrates an important aspect of game semantics: it is *quantitative*, in that it replays separately each variable call. This may seem redundant but it is not: if the execution environment is non-deterministic or has access to state, then the two calls to x may obtain different values.

This does, however, ask the question of how this repetition is managed. Of course, the definition of simple games does not forbid simply repeating the same move. For instance, one could redefine the simple game for booleans as having the same set of moves and polarities as above, but plays given by the recursive definition:

$$P_{\mathbb{B}} ::= \varepsilon \mid \mathbf{q}^- \mid \mathbf{q}^-\mathbf{tt}^+P_{\mathbb{B}} \mid \mathbf{q}^-\mathbf{ff}^+P_{\mathbb{B}},$$

but this idea quickly fails as we climb the type-theoretic ladder. Consider the term:

$$f : \mathbb{B} \rightarrow \mathbb{B} \rightarrow \mathbb{B} \vdash f(f\mathbf{tt}\mathbf{ff})\mathbf{tt} : \mathbb{B}$$

and the play in Figure 3.4. On the right, we illustrate each position with an indicative, informal description of the corresponding operational state. Intuitively, Opponent




Figure 3.4: Ambiguity without thread-indexing

moves are requests to evaluate a certain subterm; so at Opponent moves we highlight the subterm currently being evaluated. Likewise, Player questions correspond to variable calls; so at Player moves we underline the corresponding variable occurrence.

On the first line, the initial Opponent move \mathbf{q}^- prompts the evaluation of the whole term – so the whole term $f(f \text{ tt ff}) \text{ tt}$ is highlighted. In call-by-name, this prompts a Player question \mathbf{q}^+ corresponding to requesting an output for the first occurrence of f , which is therefore underlined in the second line. In our example, Opponent then plays \mathbf{q}^- in the column corresponding to the first argument of f – so the first occurrence of f requests an evaluation of its first argument, highlighted in the third line. Yet again, the highlighted sub-term starts with an occurrence of f , so the next Player move, on the fourth line, is a move \mathbf{q}^+ asking an output for f , seemingly identical to the move on the second line – but the variable occurrence underlined is not the same.

Now, in our example, the next move is an Opponent move corresponding to f calling its *second* argument. But which occurrence of f ? We have seen two of those, on the second and fourth lines! As this last \mathbf{q}^- reacts to the \mathbf{q}^+ in the fourth line, it is natural to expect that it refers to it so that the operational state is $f(f \text{ tt ff}) \text{ tt}$. But in a sufficiently expressive language, it is also possible to find an environment in which the Opponent asks to evaluate the *first* occurrence of f , *i.e.* in which case the operational state is $f(f \text{ tt ff}) \text{ tt}$. Because there are two copies of the move \mathbf{q}^+ calling to f , Opponent has, intuitively, two available copies of the move \mathbf{q}^- requesting the evaluation of the second argument of f ; and without further structure, we cannot disambiguate².

The literature proposes two ways to solve this: *pointers*, or *copy indices*.

3.1.3 Pointer Games

The most widely used solution to this problem, used in particular in Hyland-Ong games [Hyland and Ong, 2000], is to enrich plays with *pointers* providing exactly the missing information. For instance, we show in Figure 3.5 two *plays with pointers* disambiguat-

²To be precise, the usual arrow construction of simple games [Hyland, 1997], via its so-called “switching conditions”, enforces a property called *local alternation* – roughly speaking, the projection of plays on every subtype is alternating – which only allows one of these options. But the phenomenon does occur at third-order, or in game semantics not enforcing local alternation, like for higher-order state [Abramsky et al., 1998].




Figure 3.5: Two plays with pointers disambiguating Figure 3.4.

ing Figure 3.4. On the left hand side, Player is next to play \mathbf{ff}^+ while on the right hand side, Player is to play \mathbf{tt}^+ . Pointers provide the missing *thread indexing*.

Adding pointers to plays require reworking the notion of game, which must now specify which pointers are allowed. This is achieved by switching to *arenas*:

Definition 3.1.3. An *arena* comprises $A = (|A|, \text{pol}_A, \leq_A)$ with $|A|$ a countable set of moves, $\text{pol}_A : |A| \rightarrow \{-, +\}$ a *polarity function*, \leq_A a *causality partial order*, s.t.:

- forestial: for all $a_1, a_2, a \in |A|$, if $a_1, a_2 \leq_A a$, then $a_1 \leq_A a_2$ or $a_2 \leq_A a_1$,
- well-founded: there is no infinite descending $a_1 >_A a_2 >_A a_3 >_A \dots$,
- negative: if $a \in |A|$ is minimal for \leq_A , then $\text{pol}_A(a) = -$,
- alternating: for all $a_1, a_2 \in |A|$, if $a_1 \rightarrow_A a_2$, then $\text{pol}_A(a_1) \neq \text{pol}_A(a_2)$,


where $a_1 \rightarrow_A a_2$ means $a_1 <_A a_2$ and $a_1 \leq_A a \leq_A a_2$ implies $a = a_1$ or $a = a_2$.

This shall be completed (with Question/Answer labelling) in Definition 3.2.4.

We call \rightarrow_A the **immediate causality** relation on A . Historically it has been written \vdash_A in Hyland-Ong games, and called the *enabling relation* [Hyland and Ong, 2000]; but in this monograph we prefer to consider the partial order \leq_A as primitive.

The arena for the typing judgment $f : B \rightarrow B \rightarrow B \vdash B$ appears in Figure 3.6. We have annotated the type to highlight the correspondence between components of the arena and ground type occurrences. Usually these annotations are not present in diagrams for arenas or plays, but the same information is conveyed by displaying the type, and placing each move under the corresponding type component.

Both simple games and arenas are forests; but should be read differently: while in simple games an execution amounts to exploring a branch of the forest, an arena simply presents the set of observable computational events along with their hierarchical dependencies. For instance, in Figure 3.6, Opponent may start computation with \mathbf{q}_4^- . This enables three moves: \mathbf{tt}_4^+ and \mathbf{ff}_4^+ , as Player may implement a constant function answering without evaluating its arguments. And, \mathbf{q}_3^+ if Player wishes to evaluate f . In turn, this lets Opponent return \mathbf{tt}_3^- or \mathbf{ff}_3^- , or evaluate one of its arguments – and so on.

Figure 3.6: The arena for the sequent $f : \mathbb{B} \rightarrow \mathbb{B} \rightarrow \mathbb{B} \vdash \mathbb{B}$

Clearly, executions will not be merely walks on this tree: typically a function evaluates some of its arguments, possibly several times, *then* returns a value. Accordingly, executions may hop around branches of the arena, as captured formally by:

Definition 3.1.4. A *pointing sequence* on arena A is $s = s_1 \dots s_n \in |A|^*$, where each s_i may come with a *pointer* to some s_j with $j < i$. It is **legal** if it satisfies:

- justified: if s_i points to s_j , then $s_j \rightarrow_A s_i$,
- initialized: if s_i has no pointer, then it is minimal in A .

A **(alternating) legal play** on A is a legal pointing sequence which is:

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

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

For instance, the following is an alternating legal play in the arena of Figure 3.6:

$$q_4^- \leftarrow q_3^+ \leftarrow q_1^- \quad q_3^+ \quad q_2^- \quad (3.1)$$

but it is not usually drawn like that: instead it is drawn from top to bottom instead of left to right, pointers are represented via dotted lines and without arrow heads (as they must always go back in time), and the indices 1, 2, 3, 4 are omitted, the same information being conveyed by placing moves under the corresponding type component. Following this convention, the alternating legal play above is drawn as in Figure 3.5 on the rhs.

For the rest of this chapter, unless specified otherwise, by *play* we always mean *alternating legal play*. The arguably heavy notation $\Downarrow\text{-Plays}(A)$ invites an explanation: say that a play $s \in \Downarrow\text{-Plays}(A)$ is **in state O** if it has even length, *i.e.* it is Opponent's turn to play. Likewise, an odd-length play is **in state P** . The two arrows \Downarrow suggest alternation as the control alternates between two states O and P , while the arrow \curvearrowright evokes pointers as in (3.1). This is the first instance of a scheme followed throughout this monograph for different kinds of plays – alternating or not; pointing or not.

We introduce some notations on plays. We use s, t, u, v to range over plays, and ε for the empty play. We write \sqsubseteq for the prefix ordering on plays, which includes compatibility with pointers. We often extend plays by juxtaposing moves as in sab where it is

understood that s is a play and a, b are moves, leaving their possible pointer implicit. When writing equalities such as $sab = t$, it is understood that pointers are required to coincide as well. If s is a play, we write it $s_1 \dots s_n$ and refer to s_i for individual moves. If s_i points to s_j (with $j < i$), we call s_j the **justifier** of s_i .

Plays with pointers are a powerful formalism, but one drawback is that unlike in simple games, they are no longer simply strings of moves. The reader may notice that they are referred to in english, as in “ s_j points to s_i ” rather than via a specific rigorous representation, *e.g.* conveyed by a partial function from indices to indices. Debatable as it is, this is the standard practice in Hyland-Ong games – manipulations on plays with pointers are always cumbersome, but can quickly become unwieldy especially if done formally, as pointers must be constantly reindexed or reassigned³.

We can define strategies on an arena, similarly to Definition 3.1.2:

Definition 3.1.5. For A an arena, an **alternating strategy** $\sigma : A$, or \Downarrow -**strategy** for short, is a set $\sigma \subseteq \Downarrow\text{-Plays}(A)$ satisfying the following conditions:

- non-empty: $\epsilon \in \sigma$,
- prefix-closed: $\forall s \sqsubseteq t \in \sigma, s \in \sigma$,
- receptive: $\forall s \in \sigma, sa^- \in \Downarrow\text{-Plays}(A) \Rightarrow sa^- \in \sigma$,
- deterministic: $\forall sa^-b_1^+, sa^-b_2^+ \in \sigma, sab_1 = sab_2$

We refer to an alternating strategy simply as a *strategy* when it causes no confusion.

Copy indices. We are due for a discussion on the other main solution to the thread indexing problem, *copy indices*. But we prefer to postpone this to Section 3.4, and instead we carry on with the construction of pointer game semantics.

3.2 Carving Out Innocent Strategies



Now that we know how to play, we may explore the model and examine its computational meaning – starting from PCF, the core of all languages used in this monograph.

3.2.1 The Strategies of PCF

Interpretation of types. Following Section 2.2, we consider in this monograph a version of PCF with one constructor \Rightarrow for the function space, and three base types (\mathbb{U} , \mathbb{B} and \mathbb{N}). Accordingly, there are three *basic arenas* \mathbf{U} , \mathbf{B} and \mathbf{N} , described in Figures 3.7, 3.8 and 3.9 respectively – ignoring for now the annotation of moves with \mathcal{Q} and \mathcal{A} .

In light of the discussion above, these arenas should be rather self-explanatory: Opponent prompts computation with the initial move \mathbf{q}^- , and Player answers with a value. One important common point is that all three are *well-opened*:

³There have been proposals for a more rigorous handling of pointers, notably resting on nominal sets in [Gabbay and Ghica, 2012]. The issue does not appear in AJM games, where pointers are absent altogether – though this absence limits the expressiveness of AJM games. In the model proposed later in this monograph, *concurrent games*, the issue is avoided as pointers are present, but derived rather primitive.

Figure 3.7: **U**Figure 3.8: **B**Figure 3.9: **N**

Definition 3.2.1. An arena A is **well-opened** if and only if it has exactly one minimal move, called the **initial move** of A , and written $\text{init}(A)$.

In fact, arenas arising as the interpretation of types of PCF will always be well-opened. This lets us describe an easy construction for the arrow of two arenas.

Definition 3.2.2. Consider A_1, A_2 arenas with A_2 well-opened.

Then, the (well-opened) **arrow** arena $A_1 \Rightarrow A_2$ is defined via components:

$$\begin{aligned} \text{moves:} \quad & |A_1 \Rightarrow A_2| = |A_1| + |A_2| \\ \text{causality:} \quad & (i, a) \leq_{A_1 \Rightarrow A_2} (j, a') \Leftrightarrow (i = j \ \& \ a \leq_{A_1} a') \vee (i = 2 \ \& \ a = \text{init}(A_2)) \\ \text{polarities:} \quad & \text{pol}_{A_1 \Rightarrow A_2}(i, a) = (-1)^i \text{pol}_{A_1}(a). \end{aligned}$$

In $A \Rightarrow B$, the two arenas are first put side by side, with polarities reversed in A (so that Player and Opponent exchange roles). Then, A is set to depend on the initial move of B . This means that once the initial move is played in $A \Rightarrow B$, Player has two choices: either to play in B (ignoring A), or to play some minimal move in A .

Altogether, this yields an interpretation of types of PCF as arenas, via

$$\llbracket A \rightarrow B \rrbracket = \llbracket A \rrbracket \Rightarrow \llbracket B \rrbracket$$

with $\llbracket \mathbb{U} \rrbracket = \mathbf{U}$, $\llbracket \mathbb{B} \rrbracket = \mathbf{B}$ and $\llbracket \mathbb{N} \rrbracket = \mathbf{N}$. The reader is invited to check that for instance, $\llbracket (\mathbb{B} \rightarrow \mathbb{B} \rightarrow \mathbb{B}) \rightarrow \mathbb{B} \rrbracket$ is indeed the arena shown in Figure 3.6.

Interpretation of terms. For each term $\vdash M : A$, there is a strategy

$$\llbracket M \rrbracket : \llbracket A \rrbracket,$$

its *interpretation* – it is computed by induction on M following the methodology of denotational semantics, via a cartesian closed category of arenas and strategies. This will appear in due course but we postpone it for now: it takes a while to set up this machinery, which is not needed to build up an intuition for what game semantics *do*.

Though the strategy of a term is computed denotationally, an experienced game semanticist will be able to directly list its plays, without going through the intricate definition of the interpretation. This is because plays are executions: rather than denotationally, they can be obtained directly from the term by operational means⁴. This is

⁴Though making this formal is not easy, and has been the topic of an active line of work around *operational game semantics* [Danos et al., 1996, Jaber, 2015, Ghica and Tzevelekos, 2012, Levy and Staton, 2014].



The standard way of stating that the model reflects computation is:

For each $\vdash M : \mathbb{X}$, we have $M \Downarrow$ iff $\llbracket M \rrbracket$ has an answer to the initial move.

3.2.2 Well-Bracketing

We argue informally why this *cannot* be an execution in PCF. The first action of the term is to ask its argument, so within PCF it must have the form $\lambda f^{\mathbb{U} \rightarrow \mathbb{U}}. f\ M$; we

annotate the figure with the corresponding operational state as in Figure 3.10. In the last line, \mathbf{tt} at toplevel indicates the overall computation has terminated to \mathbf{tt} . This is confusing, since in the third line, only the argument of f was under evaluation. 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 [Griffin, 1990]. We shall not need the precise operational semantics of **callcc**, 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 analogous *control operator*)? We can show that the answer is no, by capturing those plays that “respect the call stack discipline”, and refining the whole interpretation to show that the interpretation of terms only yields such plays. This is the goal of the notion of *well-bracketing*.

First we enrich arenas:

Definition 3.2.4. A *Question/Answer labeling* on arena A is a function

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

which satisfies the following conditions:

$$\begin{array}{ll} \text{question-opening:} & \text{if } a \in |A| \text{ is minimal, then } \lambda_A(a) = \mathcal{Q}, \\ \text{answer-closing:} & \text{if } \lambda_A(a) = \mathcal{A}, \text{ then } a \text{ is maximal for } \leq_A, \end{array}$$

Questions intuitively correspond to variable calls, while *Answers* correspond to returns. From now on, we assume that all **arenas** have a Question/Answer labeling. For basic arenas \mathbf{U}, \mathbf{B} and \mathbf{N} , the labelling is as indicated in Figures 3.7, 3.8 and 3.9: the initial Opponent move is a Question, and the different values are Answers. This is extended to the arrow arena with $\lambda_{A \Rightarrow B}(1, a) = \lambda_A(a)$ and $\lambda_{A \Rightarrow B}(2, b) = \lambda_B(b)$.

If $s \in \Downarrow\text{-Plays}(A)$ and s_i is an answer, it cannot be minimal in A by *question-opening*. Its antecedent in A appears 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:

Definition 3.2.5. Consider A an arena, and $s \in \Downarrow\text{-Plays}(A)$ an alternating legal 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 monograph until now are well-bracketed, with the exception of the example at the beginning of Section 3.2.2. We can then define:

Definition 3.2.6. Let $\sigma : A$ be a strategy on A .

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





Figure 3.11: A play for a strategy with references

Thus, a well-bracketed strategy is never the first to break the call stack discipline.

We shall see that all the operations on strategies used in the interpretation of PCF (notably, composition) preserve well-bracketing. In the execution at the start of Section 3.2.2, Player is the first to break well-bracketing; so this execution cannot occur in a well-bracketed strategy, and is therefore – as expected – unrealizable in PCF.

3.2.3 Visibility and Innocence

Likewise, is this play a possible execution of a term?



This also seems unfeasible in PCF. Again, on the rhs 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 of PCF with a primitive \oplus for non-deterministic choice, this play would be realisable by $\lambda f^{B \rightarrow U}. f (\underline{tt} \oplus \underline{ff})$. 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

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

in IA (see Section 2.2), *i.e.* PCF extended with references: **newref** r in M allocates a reference r initialized to 0. We show in Figure 3.11 an operational description.

Yet, again, this play cannot be realised in PCF. To show this, we give a version of *innocence* [Hyland and Ong, 2000], formalizing that without state, independent evaluations of the same subterm must yield the same response. The first step is a mathematical way to state that two Opponent-ending plays “evaluate the same subterm”, like the two prefixes of the play of Figure 3.11 terminating with a q^- on the left.




Figure 3.12: A non P-visible play

The operation computing (a mathematical notion of) “current subterm” is:

Definition 3.2.7. Let $s \in \Downarrow\text{-Plays}(A)$. Its **P-view** is the subsequence defined by:

$$\begin{aligned} \ulcorner \varepsilon \urcorner &= \varepsilon \\ \ulcorner sa^+ \urcorner &= \ulcorner s \urcorner a^+ \\ \ulcorner sa_1^+ s' a_2^- \urcorner &= \ulcorner s \urcorner a_1^+ a_2^- \quad \text{if } a_2 \text{ points to } a_1, \\ \ulcorner sa_2^- \urcorner &= a_2^- \quad \text{if } a_2 \text{ is minimal,} \end{aligned}$$

where pointers are preserved if their target remains in the subsequence.

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 3.11 have the same P-view, capturing that they correspond to the same subterm. This is a fundamental definition – really, one of the cornerstones of HO games [Hyland and Ong, 2000].

But in order to use this definition it better be a well-defined operation on plays, which is not yet the case: for instance, in Figure 3.12 we gray out moves not selected in computing the P-view of $s \in \Downarrow\text{-Plays}(A)$ for $A = [(U \rightarrow U) \rightarrow U]$. The subsequence of $\ulcorner s \urcorner$ in black is an alternating sequence of moves as needed, but it does not contain the justifier of \checkmark^+ . So \checkmark^+ has no justifier in $\ulcorner s \urcorner$, which means that $\ulcorner s \urcorner \notin \Downarrow\text{-Plays}(A)$ as \checkmark^+ is not minimal in A . Accordingly, we focus on those plays where $\ulcorner - \urcorner$ is defined:

Definition 3.2.8. A play $s \in \Downarrow\text{-Plays}(A)$ is **P-visible** if $\forall t \sqsubseteq s, \ulcorner t \urcorner \in \Downarrow\text{-Plays}(A)$. Likewise, a strategy $\sigma : A$ is **P-visible** (or just **visible**) iff all its plays are P-visible.

We write $\mathbf{Vis}(A)$ for the set of P-visible strategies on arena A .

So, “computing P-views never drops pointers”, or “Player always points in the P-view”. On P-visible $s \in \Downarrow\text{-Plays}(A)$, the P-view always yields a (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 3.2.9. A P-visible strategy $\sigma : A$ is **innocent** if it satisfies:

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

We write $\mathbf{Inn}(A)$ for the set of innocent strategies on arena A .

The clause $\ulcorner sa \urcorner = \ulcorner ta \urcorner$ serves to ensure that a has the “same” justifier in s and in t . Again we shall see that innocence is preserved by all operations involved in the interpretation of PCF so that any term of PCF is interpreted by an innocent strategy.

The play at the beginning of Figure 3.2.3 contains a direct failure of innocence – hence, it cannot appear in an innocent strategy and thus cannot be realised in PCF.

O-visibility. Whenever a condition constrains the behaviour of Player, one may consider the dual concerning Opponent. As a short digression we consider *O-visibility*, which will be technically helpful later on in this chapter. First we define:

Definition 3.2.10. Let $s \in \Downarrow\text{-Plays}(A)$. Its **O-view** is the subsequence defined by:

$$\begin{aligned} \ulcorner \varepsilon \urcorner &= \varepsilon \\ \ulcorner sa^- \urcorner &= \ulcorner s \urcorner a^- \\ \ulcorner sa_1^- s' a_2^+ \urcorner &= \ulcorner s \urcorner a_1^- a_2^+ \quad \text{if } a_2 \text{ points to } a_1, \end{aligned}$$

where pointers are preserved if their target remains in the subsequence.

The definition mirrors that of the P-view, with one missing clause for the initial move: this mirrors an asymmetry in the definition of arenas/plays, where initial moves must always be negative. As for P-views, this operation is well-defined only for:

Definition 3.2.11. A play $s \in \Downarrow\text{-Plays}(A)$ is **O-visible** if $\forall t \sqsubseteq s, \ulcorner t \urcorner \in \Downarrow\text{-Plays}(A)$.

O-visibility captures those plays where the external environment acts in a P-visible way, *i.e.* does not have access to effects that break P-visibility, see Section 3.3.3. We say that $s \in \Downarrow\text{-Plays}(A)$ is **visible** when it is both P-visible and O-visible.

One can also apply the same principle to innocence and define O-innocence capturing these plays where Opponent acts innocently, but this turns out less useful.

3.2.4 Finite Definability and Full Abstraction

At this point, via well-bracketing and innocence we have effectively eliminated all PCF non-definable behaviour. The theorem making this formal is called *finite definability*; we show the main steps of its proof as well as how it entails *intensional full abstraction*.

Meager innocent strategies. We have used the P-view operation as a mathematical attempt to capture the “current subterm” of a term reached in a play. Accordingly, it makes sense to attempt recovering syntax by examining those plays obtained by $\ulcorner - \urcorner$. Clearly, for any P-visible $s \in \Downarrow\text{-Plays}(A)$, we have $\ulcorner \ulcorner s \urcorner \urcorner = \ulcorner s \urcorner$, so the plays obtained by $\ulcorner - \urcorner$ are exactly those invariant under $\ulcorner - \urcorner$, easily characterized as those P-visible plays where *Opponent always points to the previous move* – call these the **P-views**.

Following this discussion, it should be no surprise that the P-views are the key to reconstructing a term from a strategy. If $\sigma : A$ is P-visible, we define

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

the set of **P-views** of σ . If σ is innocent, it is a simple fact that $\ulcorner \sigma \urcorner \subseteq \sigma$. Moreover, σ can then be recovered as the set of P-visible plays s such that for all $t \sqsubseteq s$, we have $\ulcorner t \urcorner \in \ulcorner \sigma \urcorner$. Thus, sets of P-views are an equivalent representation of innocent strategies.

It is informative to examine those sets of P-views matching innocent strategies:

Definition 3.2.12. A *meager innocent strategy* on A is a set V of P-views on A , s.t.:

- non-empty: $\varepsilon \in V$,
- prefix-closed: for all $s \in V$ and $t \sqsubseteq s$, then $t \in V$
- receptive: for all $s \in V$, if sa^- is a P-view then $sa \in V$,
- deterministic: for all $sa_1^+, sa_2^+ \in V$ we have $sa_1 = sa_2$.

Finally, V is **well-bracketed** if all its plays are well-bracketed.

We write **Meager**(A) for the set of meager strategies on arena A .

From the above and routine verification, we obtain:

Proposition 3.2.13. Consider A an arena. Then we obtain a bijection:

$$\ulcorner - \urcorner : \mathbf{Inn}(A) \simeq \mathbf{Meager}(A),$$

and moreover σ is well-bracketed iff $\ulcorner \sigma \urcorner$ is.

Proof. The only subtle point is that if $\ulcorner \sigma \urcorner$ is well-bracketed, then so is σ . This follows from the preliminary observation, proved by induction on s , that if s is well-bracketed then its pending question (if any) always appears in the P-view $\ulcorner s \urcorner$. \square

In a certain sense, meager innocent strategies *are* syntax – or at least, each finite meager innocent strategy may be equivalently written in the syntax of PCF:

Finite definability. Let us say that an innocent strategy $\sigma : A$ is **finite** if the corresponding meager strategy $\ulcorner \sigma \urcorner$ has finitely many even-length plays. The **size** of finite σ is then the cardinal of the set of $s \in \ulcorner \sigma \urcorner$ of even length.

Theorem 3.2.14. Let A be a PCF type, and $\sigma : \llbracket A \rrbracket$ be finite well-bracketed innocent. Then, there is a PCF term $\vdash M : A$ s.t. $\llbracket M \rrbracket = \sigma$.

Sketch. For more details than this sketch, see [Hyland and Ong, 2000].

W.l.o.g., the type 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 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} . The situation is

$$A_1 \rightarrow \dots \rightarrow (A_{i,1} \rightarrow \dots \rightarrow A_{i,p_i} \rightarrow \mathbb{X}_i) \rightarrow \dots \rightarrow A_n \rightarrow \mathbb{X}$$

with, in grey, the possible extended P-views. For each, 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

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

a meager innocent strategy on $\llbracket A \rrbracket$ of size strictly lesser than σ . By IH there is $\vdash M_v : A$ with $\llbracket M_v \rrbracket = \sigma_v$. 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 $\mathbf{q}^- \mathbf{q}_{i_0}^+ \mathbf{q}_{i_0,j}^- s \in \llbracket \sigma \rrbracket$ where as a P-view, s answers neither $\mathbf{q}_{i_0}^+$, nor \mathbf{q}^- by well-bracketing. Such a P-view yields

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

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 meager 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. $\llbracket M_{i_0,j} \rrbracket = \sigma_{i_0,j}$. 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 $\llbracket M \rrbracket = \sigma$. \square

This final statement is conceptually straightforward (if one has access to the definition of the interpretation, which we save for the next chapter), but technically challenging as reasoning concretely on the categorical definition of the interpretation and on the (forthcoming) composition of strategies is unwieldy. In the author's opinion, there is still room for improvement in finding a conceptually clean, fully detailed and human-readable proof of this that makes it as easy technically as it is conceptually.

Here, **case** is the syntax introduced in Section 2.2, 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 [Hyland and Ong, 2000] as the historical presentation of PCF does not include a **let** construct. This yields a term that is not quite σ , but is nonetheless observationally equivalent.

We conclude this discussion on definability with the statement:

Conjecture 3.2.15 (Intensional Universality). *Consider a PCF type A . For any computable well-bracketed innocent $\sigma : \llbracket A \rrbracket$, there is $\vdash M : A$ s.t. $\llbracket M \rrbracket = \sigma$.*

Hyland and Ong [Hyland and Ong, 2000] prove universality up to *contextual equivalence* (with a similar independent result in AJM games [Abramsky et al., 2000]). Their

construction fails to achieve intensional universality because their version of PCF lacks a “**let**” construct, making them unable to avoid repeated computation. There is strong confidence that their construction yields intensional universality in the presence of **let** (intensional universality does hold in call-by-value [Murawski and Tzevelekos, 2013]), but as far as we know the details have not been properly checked.

Intensional Full Abstraction. Recall from Section 2.3 that a model is *intensionally fully abstract* for a language when its observational quotient is fully abstract.

The famous *full abstraction* for PCF follows by a fairly typical argument, once one has an adequate model with sufficient definability properties.

Theorem 3.2.16 (Intensional Full Abstraction for PCF). *Innocent well-bracketed strategies form an intensionally fully abstract model for PCF.*

Sketch. Consider $\vdash M, N : A$ observationally equivalent, and assume $\llbracket M \rrbracket$ and $\llbracket N \rrbracket$ are not, *i.e.* there is a test α which distinguishes them. One may easily prove (see [Hyland and Ong, 2000]) that it suffices to consider α finite. By Theorem 3.2.14, α is defined by a PCF term, which we present as a context $C[-]$. By Proposition 3.2.3, it then follows that $C[-]$ distinguishes M and N , contradiction. \square

We have seen that by imposing well-bracketing and innocence, we have successfully banned all non PCF-definable behaviour. But dually, are references and control operators *complete*, are they sufficient to express unrestricted strategies?

3.3 Non-Innocence and Effects

Though Hyland-Ong games were originally designed as a model of plain PCF, what is in hindsight one of their most striking feature is that a crisp description of a few computational effects already hide in the definitions seen so far.

3.3.1 Non-Local State

As before, we start the discussion by looking at an example:

$$\begin{array}{c} \mathbb{B} \\ \mathbf{q}^- \\ \vdots \\ \mathbf{tt}^+ \\ \\ \mathbf{q}^- \\ \vdots \\ \mathbf{ff}^+ \end{array}$$

Clearly, this play is not realizable in PCF since it breaks innocence. But unlike the play of Section 3.2.3, it is not realizable either in IA, the extension of PCF with local state. Indeed, if this is to be realized by a deterministic program, some state must be

shared – or some information must flow – between the two calls. In contrast, a closed term with access to local state may initialize variables only once computation has begun, and thus separate calls cannot enjoy any shared state.

Concretely, an hypothetical $\vdash \mathbf{bad} : \mathbb{B}$ acting as above would break β -equivalence:

$$(\lambda x. \mathbf{if} \ x \ x \ x) \mathbf{bad}$$

should evaluate to \mathbf{ff} , as all calls are routed to \mathbf{bad} . Yet, $\mathbf{if} \ \mathbf{bad} \ \mathbf{bad} \ \mathbf{bad}$ should evaluate to \mathbf{tt} : we now have three separate instances of \mathbf{bad} , so only the first call matters. In this monograph we deem this pathological: we are concerned with call-by-name languages, which must hence enjoy the call-by-name equational theory, including β -equivalence.

We ban \mathbf{bad} via the notion of *single-threadedness*: a single-threaded strategy only depends on the *current thread*, so that it cannot pass information between distinct instances. For the next definition, in a play $s \in \Downarrow\text{-Plays}(A)$, we say that s_i is **hereditarily justified** by s_j (with $j < i$) if following the pointers from s_i eventually hits s_j . If s_j is initial, we say that s_j is the (necessarily unique) **initial hereditary justifier** of s_i .

Definition 3.3.1. Consider A an arena, and $s \in \Downarrow\text{-Plays}(A)$.

The *current thread* of non-empty s , written $[s]$, is the subsequence of s comprising all moves with the same initial hereditary justifier as the last move of s ; and $[\varepsilon] = \varepsilon$.

The definition of single-threaded strategies then mimics that of innocent strategies:

Definition 3.3.2. Consider $\sigma : A$ any strategy. It is *single-threaded* if it satisfies:


- well-threaded: for all $sa^+ \sqsubseteq t \in \sigma$, the justifier of a appears in $[s]$,
- single-threaded: for all $sa^+, t \in \sigma$, if $[s] = [t]$ then $ta^+ \in \sigma$ with $[sa] = [ta]$.

This bans pathological strategies such as \mathbf{bad} . On the other hand, single-threadedness is preserved under all operations involved in the interpretation of \mathbf{IA} so that the interpretation of any term of \mathbf{IA} yields a single-threaded strategy.

3.3.2 Higher-Order State

Finally, for our last example, consider the following play:

$$(((\mathbb{U} \rightarrow A \rightarrow \mathbb{U}) \rightarrow A \rightarrow \mathbb{U}) \rightarrow \mathbb{U})$$



$$\begin{aligned} & \lambda F. F(\lambda xy. M) N \\ & \lambda F. \underline{F}(\lambda xy. M) N \\ & \lambda F. F(\lambda xy. M) N \\ & \lambda F. F(\lambda xy. \underline{x}) N \\ & \lambda F. F(\lambda xy. M) N \\ & \lambda F. F(\lambda xy. M) \underline{y} \end{aligned}$$

As shown in the indicative term on the right hand side, and following the intuitions on the operational content of game semantics so far, the last Player move corresponds to calling variable y . But how can that make sense, since y is not even in the scope!

Yet, again, this play is computationally meaningful – it witnesses *higher-order state*:

$$\lambda F. \mathbf{newref}_A r \text{ in } F (\lambda xy. r := y; x) (!r)$$

Though we borrow the notation from IA, this must be read completely differently. Here A can be an arbitrary type, *i.e.* potentially a function. Whereas in IA, references store *values*, here they store *computations*; so that when the reference is read, the control flow *jumps* accordingly, giving control to arbitrary program phrases.

Higher-order state is a sensible programming feature, found in the wild in call-by-value languages such as ML. In call-by-value game semantics, a games model with morphisms set to comprise all well-bracketed single-threaded strategies is fully abstract model for a version of ML with higher-order store [Abramsky et al., 1998]. Semantic investigations have also been conducted in call-by-name [Laird, 2002, Goyet, 2013], but this is somewhat less natural from a programming language point of view. Either way, we leave higher-order state out of the scope of this monograph.

So how should one get rid of (the behaviours generated by) higher-order state? Without higher-order state, only information can flow through shared state, not control. So Player can only pass control (*i.e.* call variables) within the *scope*, *i.e.* the current syntactic branches. The same sentence, translated in game semantics terminology along the dictionary elaborated so far, reads: *Player can only point within the P-view*. But we have already seen this: this is *P-visibility* (Definition 3.2.8). So P-visibility, originally imposed as a preliminary construction just to make P-views well-defined, is actually also the condition we need to allow ground state but ban higher-order state!

P-visibility is preserved by all constructions used in the interpretation of IA (including ground type references), while the play above clearly is not P-visible.

3.3.3 The “Semantic Cube”

It is striking that almost, if not all, strategies in Hyland-Ong games correspond to existing programming features, for a framework originally designed for plain PCF only! From now on we assume that all strategies are single-threaded and P-visible, the canvas for the celebrated “semantic cube” that we shall present now.

Idealized Algol. Well-bracketed strategies support an adequate interpretation of IA:

Proposition 3.3.3. *There is an interpretation $\llbracket - \rrbracket$ sending types of IA to arenas, and terms $\vdash M : A$ to well-bracketed, single-threaded, P-visible $\llbracket M \rrbracket : A$.*

Moreover, for all $\vdash M : \mathbb{U}$, $M \Downarrow$ iff in $\llbracket M \rrbracket$, the initial move has an answer.

The interpretation of IA will be detailed in Section 4.2.

It turns out that well-bracketed strategies support “only” the interpretation of IA: we have a finite definability property similar to Theorem 3.2.14. We state:

Theorem 3.3.4. *Consider A an IA type.*

For any $\sigma : \llbracket A \rrbracket$ visible, well-bracketed, single-threaded, and finite (i.e. with finitely many even-length plays) there is $\vdash M : A$ a term of IA such that $\llbracket M \rrbracket = \sigma$.

We shall not give the detailed proof of this as their statement relies on the compositional structure, but the key arguments appear in Section 4.3.2.

As for PCF, from those definability properties follow:

Theorem 3.3.5 (Intensional Full Abstraction for IA). *P-visible well-bracketed single-threaded strategies form an intensionally fully abstract model for IA.*

The reasoning for that is the same as for Theorem 3.2.16.

Recall that intensional full abstraction means that the observational quotient of the model is fully abstract; in other words, visible well-bracketed single-threaded strategies have the same distinguishing power as terms of IA. As for PCF this makes the quotiented fully abstract model hard to reason with, but for IA something much more interesting happens: we can characterize concretely observational equivalence.

Complete plays. This is achievable by understanding which parts of a P-visible, well-bracketed strategy can be explored by a P-visible, well-bracketed Opponent – and among those, which lead to an observable result at toplevel. This is done via:

Definition 3.3.6. *Consider A an arena. A play $s \in \Downarrow\text{-Plays}(A)$ is **complete** iff it is well-bracketed, O-visible, P-visible, and every question has an answer.*

If $\sigma : A$, we write $\text{comp}(\sigma)$ for the set of its complete plays. Complete plays indeed characterize observational equivalence for P-visible well-bracketed strategies:

Proposition 3.3.7. *Consider $\sigma, \tau : A$ two P-visible, well-bracketed strategies. Then,*

$$\sigma \simeq \tau \quad \Leftrightarrow \quad \text{comp}(\sigma) = \text{comp}(\tau).$$

Again we omit the (rather simple) details which rely on the compositional structure introduced in the next chapter only, but the idea is the following. If a strategy $\alpha : A \Rightarrow \mathbb{U}$ distinguishes σ and τ , then it is straightforward to read from that a complete play in one but not the other. Reciprocally, if s is a complete play in σ but not in τ , then $\mathbf{q}^- s \checkmark^+$ (with adequate pointers) is a play in $A \Rightarrow \mathbb{U}$. The set of all prefixes of $\mathbf{q}^- s \checkmark^+$, closed under receptivity, is then a valid strategy that distinguishes σ and τ .

From Proposition 3.3.7 and Theorem 3.3.5, it shall be clear that we get:

Theorem 3.3.8. *Consider A a type, and $\vdash M, N : A$ two terms of IA. Then,*

$$M \simeq N \quad \Leftrightarrow \quad \text{comp}(\llbracket M \rrbracket) = \text{comp}(\llbracket N \rrbracket).$$

This entails that the fully abstract model is *effectively presentable*: it can be phrased as a computable function into an effective domain [Plotkin, 1981]; concretely this means that observational equality is decidable between finite strategies (those with finitely

many even-length plays; forming the compact elements of the domain of strategies). This is remarkable: in the theory of programming languages, situations like that where one has a concrete characterization of observational equivalence, are rare!

This comes from the ability to write a context $C[-]$ reproducing one complete play s and *only* s ; so that $C[M] \Downarrow$ is exactly equivalent to $s \in \llbracket M \rrbracket$. This is impossible for an innocent context, because if an innocent strategy includes one play, it also includes all plays with the same P-views. In fact, we know that the observational quotient in the model of PCF cannot be characterized concretely, because of Loader’s result that the (unique) fully abstract model of PCF is not effectively presentable [Loader, 2001].

Full abstraction with control. We now drop well-bracketing. We also consider temporarily the extensions CIA and CPCF respectively of IA and PCF with

$$\mathbf{callcc}_{A,B} : ((A \rightarrow B) \rightarrow A) \rightarrow A,$$

a *control operator* with behaviour described informally in Section 3.2.2. We shall not cover its behaviour or interpretation in detail as it plays no role in this monograph beyond painting the big picture. But we do mention two main results, first for CIA:

Theorem 3.3.9. *Visible single-threaded strategies form a fully abstract model for CIA. Moreover, if $\vdash M, N : A$ are two terms of CIA, then*

$$M \simeq N \quad \Leftrightarrow \quad \mathbf{O-vis}(\llbracket M \rrbracket) = \mathbf{O-vis}(\llbracket N \rrbracket),$$

with $\mathbf{O-vis}(\sigma)$ the set of *O-visible* plays of a strategy σ .

This is proved again via an adequate definability result. This theorem was surely known long before it was published⁵ in [Murawski, 2007]. Here again, the fully abstract model is also effectively presentable. Note that contexts from CIA can observe more than contexts from IA: **callcc** lets us observe plays that cannot be completed, and hence would never have led to an observable behaviour within IA.

Finally, we mention one last fully abstract model:

Theorem 3.3.10. *Visible innocent strategies form a fully abstract model for CPCF.*

This is due to [Laird, 1997], proved again via an adequate definability property. In this case also, the fully abstract model is effectively presentable, but not as concretely as above: instead, one shows that equality between compact elements is decidable, by reducing equivalence checking to a finite number of tests. The fully abstract model can also be described concretely – and in fact *was* described concretely, before even the fully abstract games model for PCF were developed [Cartwright et al., 1994]!

The “semantic cube”. We have, so far, seen the following full abstraction results:

⁵Murawski uses a different primitive for control, but the difference is superficial within IA.




Figure 3.13: The Semantic Cube

Theorem 3.3.11 (Semantic Cube). *We have four intensional full abstraction results:*

G	<i>is fully abstract for</i>	$PCF + \text{state} + \text{control}$
$G + \text{innocence}$	<i>is fully abstract for</i>	$PCF + \text{control}$
$G + \text{well-bracketing}$	<i>is fully abstract for</i>	$PCF + \text{state}$
$G + \text{innocence} + \text{well-bracketing}$	<i>is fully abstract for</i>	PCF

where G denotes the model formed of visible, single-threaded strategies.

This “Semantic Cube”, drawn in Figure 14.2, 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 effects. But independently of purely semantic purposes, this gives us a microscope to study behaviourally interactions between effects in programming languages. We demonstrate this with the following orthogonality property⁶ between state and control which nicely illustrates the strength of game semantics. We only state it as a conjecture, because it relies on intensional universality for PCF.

Conjecture 3.3.12. *Let $\vdash M : A$ a term of CIA with A a PCF type. Assume that*

- (1) $M \simeq N_1$ where N_1 does not use **callec**,
- (2) $M \simeq N_2$ where N_2 does not use **newref**,

then there is $\vdash N : A$ in pure PCF such that $M \simeq N$.

Proof. We have $N_1 \simeq N_2$. By Theorem 3.3.9, they have the same O-visible plays. The P-views of $\llbracket N_2 \rrbracket$ are O-visible, so they must be in $\llbracket N_1 \rrbracket$, so they must be well-bracketed. But $\llbracket N_2 \rrbracket$ is innocent, so if its P-views are well-bracketed, then it must be well-bracketed (Proposition 3.2.13). Summing up, $\llbracket N_2 \rrbracket$ is an innocent well-bracketed strategy. But it is also computable, since the interpretation is computable. Hence, by Conjecture 3.2.15, there is $\vdash N : A$ in pure PCF such that $\llbracket N \rrbracket = \llbracket N_2 \rrbracket$. In particular, it has the same O-visible plays as M , hence $M \simeq N$ by Theorem 3.3.9 again. \square

We finish this section with two concluding notes.

Firstly, game semantics has the reputation that its full abstraction results are unsatisfactory, because they consist in constraining strategies until they are essentially syntax,

⁶We learnt of it from a talk by Paul Levy in 2014 [Levy, 2014].

followed by the same quotient that one could have done already in the syntax. This criticism holds for PCF, and even then seems short-sighted: without the fully abstract model for PCF, we would not have the semantic cube. But furthermore, beyond PCF, *all fully abstract models in the semantic cube can be effectively presented!*

Secondly, unfortunately the semantic cube holds only for sequential deterministic languages. Beyond that we still have models (we shall see one in Chapter 5), but they do not fit in the same unified landscape as the cube above. What can be done about that? This is one of the leading questions behind this monograph.

3.4 An Alternative to Pointers: Copy Indices

As the final section of this chapter, we briefly present an alternative to the use of pointers for thread indexing: *copy indices*, used in particular in the other major family of traditional game semantics models, called AJM games for Abramsky, Jagadeesan and Malacaria [Abramsky et al., 2000]. We do this, on the one hand, to account for the large part of the game semantics literature that depends on it; and on the other hand, to help prepare the reader for concurrent games, which use copy indices.

The idea of copy indices is that rather than explicitly authorizing repetitions (and adding pointers to disambiguate the relationships between copies), we keep the invariant that each move must appear *at most once* in a play. So as to allow players to perform duplications, we create explicit copies of duplicable moves, which are kept apart by adjoining to each a unique identifier, a natural number called a *copy index*.

Copy indices must be added by an explicit construction on games. To that end, games with copy indices typically follow linear logic [Girard, 1987] and decompose the arrow construction in two steps $A \Rightarrow B = !A \multimap B$. Here, \multimap is a *linear* or *affine* arrow that does not allow replication of the argument; and $!A$ is the *bang* construction from linear logic, which creates countably many copies of a game by adjoining copy indices:

Definition 3.4.1. *Consider A a simple game. Then, the simple game $!A$ has:*

$$\begin{aligned} |!A| &= \mathbb{N} \times |A| \\ \text{pol}_{!A}(i, a) &= \text{pol}_A(a) \\ P_{!A} &= \{s \in |!A|^* \mid s \text{ alternating}, \forall i \in \mathbb{N}, s \upharpoonright i \in P_A\} \end{aligned}$$

where $s \upharpoonright i \in |A|^*$ is the obvious operation only keeping moves of the form (i, a) .

From here and throughout this monograph, we adopt the convention that copy indices are typeset in grey: this helps in easily distinguishing them from other indices, and is in line with the intuition that though their presence is needed to give copies an identity, their precise value is in general irrelevant. Now we have, for instance:

$$P_{!B}^I ::= \varepsilon \mid \mathbf{q}_i^- \mid \mathbf{q}_i^- \mathbf{u}_i^+ P_{!B}^{I \uplus \{i\}} \mid \mathbf{q}_i^- \mathbf{f}_i^+ P_{!B}^{I \uplus \{i\}} \quad i \notin I$$

and $P_{!B} = P_{!B}^\emptyset$; so plays in $!B$ is any sequence of plays in B , each with a copy index, in any order but such that each copy index is used at most once.




Figure 3.14: Two plays with copy indices disambiguating Figure 3.3.

As mentioned above, the usual call-by-name arrow is defined in linear logic by “Girard’s translation”. Accordingly, the interpretation of types include an occurrence of $!$ each time we cross the left of an arrow. Now, we saw that each move corresponds to a precise occurrence of a ground type in the type/typing judgment; so in the scope of a fixed number of $!$. Accordingly, each move carries a *sequence* of copy indices, once for each ambient $!$ – by convention, we write first the copy index for the outermost $!$.

With this convention, the two plays corresponding to Figure 3.5 may be drawn⁷ with copy indices in Figure 3.14. This shows how copy indices carry implicitly the hierarchical relationship between moves expressed by pointers. Yet, plays are still straightforward sequences of moves, without unwieldy additional structure.

But if the information expressed by pointers can be recovered on mere plays with such a small technical overhead, why bother with pointers at all? In fact, we have yet to see the main subtlety – and technical hurdle – arising from copy indices.

Equality up to copy indices. The issue is that there are now many more strategies, as when playing a positive move, Player may seemingly arbitrarily choose a copy index – which does not seem to carry any relevant computational information.

In fact, this has the concrete consequence that simple games with replication as in Definition 3.4.1 fail to be a model of the simply-typed λ -calculus. Indeed we have

$$x : \mathbb{B} \vdash (\lambda f. f \ x \ x) (\lambda yz. y) \cong_{\beta} (\lambda f. f \ x \ x) (\lambda yz. z) : \mathbb{B} \quad (3.2)$$

yet, following the usual definition of the interpretation (which we shall not present right now), the left hand side calls x with copy index 0 while the right hand side calls x with copy index 1 . Copy indices trace back which occurrence of x was selected in $\lambda f. f \ x \ x$, but remembering this is incompatible with β -equivalence.

To address this, AJM games add to simple games an equivalence relation on plays:

⁷Again, the play on the right hand side is not accepted by the usual definitions of simple games or AJM games as it is not *locally alternating*. But it is nevertheless illustrative of the real phenomenon.

Definition 3.4.2. An **AJM game** comprises a simple game $A = (|A|, \text{pol}_A, P_A)$, together with an equivalence relation \cong_A on P_A satisfying the following conditions:

- (1) if $s_1 \dots s_n \cong_A t_1 \dots t_p$, then $n = p$ and for all $1 \leq i \leq n$, $\text{pol}_A(s_i) = \text{pol}_A(t_i)$,
- (2) if $s_1 \dots s_n \cong_A t_1 \dots t_n$, then for all $1 \leq i \leq n$, $s_1 \dots s_i \cong_A t_1 \dots t_i$,
- (3) if $s \cong_A t$ and $sa \in P_A$, then there is $tb \in P_A$ such that $sa \cong_A tb$.

We think of the equivalence relation \cong_A as formalizing *reindexings*, i.e. we have $s \cong_A t$ iff s and t are the same play, but with a different choice of copy indices. For ground types such as \mathbb{B} the equivalence relation is restricted to reflexive pairs, but:

Definition 3.4.3. Consider A an AJM game. We set $!A$ following Definition 3.4.1, plus:

$$s_1 \dots s_n \cong_{!A} t_1 \dots t_n \Leftrightarrow \begin{cases} \text{there is a permutation } \pi : \mathbb{N} \simeq \mathbb{N} \text{ such that} \\ \text{for all } 1 \leq i \leq n, \text{ if } s_i = (k, a) \text{ then } t_i = (\pi(k), a'), \\ \text{and for all } k \in \mathbb{N}, s \upharpoonright k \cong_A t \upharpoonright \pi(k). \end{cases}$$

In other words, an equivalence $s \cong_{!A} t$ follows a global reindexing on copy indices introduced by $!$, and local reindexings on each individual copy.

These reindexings on plays induce a partial equivalence relation on strategies:

Definition 3.4.4. Consider A an AJM game.

For $\sigma, \tau : A$ strategies on the underlying simple game, we write $\sigma \approx \tau$ iff:

- \rightarrow -simulation: $\forall sa^+ \in \sigma, t \in \tau, s \cong_A t \Rightarrow \exists b^+, tb^+ \in \tau \ \& \ sa^+ \cong_A tb^+$
- \leftarrow -simulation: $\forall s \in \sigma, tb^+ \in \tau, s \cong_A t \Rightarrow \exists a^+, sa^+ \in \sigma \ \& \ sa^+ \cong_A tb^+$
- \rightarrow -receptive: $\forall sa^- \in \sigma, t \in \tau, sa^- \cong_A tb^- \Rightarrow tb^- \in \tau$
- \leftarrow -receptive: $\forall s \in \sigma, tb^- \in \tau, sa^- \cong_A tb^- \Rightarrow sa^- \in \sigma$

We say that $\sigma : A$ is **uniform** iff $\sigma \approx \sigma$.

This allows us to consider equal strategies behaving the same up to their choice of copy indices – for instance, it identifies the two strategies corresponding to the terms of (3.2). But it does more: it select the *uniform* strategies, which are unable to observe the copy indices chosen by Opponent; i.e. with respect to which \approx is a congruence.

When setting up a game semantics based on copy indices, the idea is that one only considers strategies which are uniform, and those are considered up to \approx .

Pointers or Copy Indices? In designing a game semantics supporting replication, one must use one or the other of these two structures. Which one should be chosen?

Beyond thread indexing, pointers also import into plays relevant semantical structures. In Hyland-Ong games, pointers are key to defining *P-visibility* and *innocence* [Hyland and Ong, 2000], the conditions that eventually gave rise to the semantic cube. On the other hand, pointers are hard to manipulate in a rigorous yet unobtrusive way.

In contrast, the AJM model has simpler plays, and enjoys a clean linear decomposition in the sense of Linear Logic [Girard, 1987]. However, terms are interpreted as

\approx -equivalence classes rather than concrete strategies and all constructions come with a sometimes challenging proof obligation that it preserves \approx . Conditions which rely on pointers such as visibility and innocence are beyond reach without further structure.

Much of the literature in game semantics for programming languages has been developed based on pointers and Hyland-Ong games. However, in this monograph, we shall see that *concurrent games* are based on copy indices. This is for two reasons: firstly, their finer intensional structure seems to require it. Secondly, despite using copy indices and unlike in AJM games, the *pointers* of Hyland-Ong games which have proved so fruitful in studying program behaviour are retained as derived structures.

3.5 Conclusions and Historical Notes

Much of the basic definitions presented in this chapter originate (though with minor variations) in Hyland and Ong’s original paper [Hyland and Ong, 2000], including the definitions of *visibility*, of *innocence* and of *well-bracketing* (on plays). The realization that lifting these constraints yielded fully abstract models for richer languages came very quickly: full abstraction for IA is due to [Abramsky and McCusker, 1996], full abstraction for CPCF is due to Laird in [Laird, 1997]. To our knowledge, the full abstraction result for the missing corner of the cube (CIA) was only developed in detail significantly later by Murawski [Murawski, 2007] – though it was certainly well-known.

It might seem confusing that apart from this latter paper, all the works cited here, which build on Hyland and Ong’s framework, date from *before* the publication of the original Hyland and Ong paper. In fact, though the Hyland and Ong paper was published only in 2000, preliminary versions of it circulated from the earlier 90s⁸, in parallel with versions of the AJM model. These developments, and announcements by the two competing groups, were closely followed by the broader denotational semantics community, feeling that the resolution of the problem of “full abstraction for PCF” was close.

In hindsight, rather than the full abstraction result for PCF *per se*, what had the most impact was definitely the nice surprise that many different effects could be captured in a unified framework; justifying the 2017 *Alonzo Church Award* awarded to the founders of game semantics. Game semantics has since then grown to a broad topic; and this monograph shall cover some of its most recent developments.

⁸The workshop celebrating the 25 years of game semantics was held in 2018 as part of the federated logic conference (FLoC), which places the birth of game semantics – somewhat arbitrarily – in 1993.

Chapter 4

The Category of Alternating Strategies

In this chapter, we continue our introduction to sequential game semantics, filling the gap left open in the previous chapter: the compositional structure. We show how to compose alternating strategies, forming a cartesian closed category $\Downarrow\text{-Strat}$. Then, we define the interpretation of IA into $\Downarrow\text{-Strat}$, and conclude with some complements.

4.1 The Ambient Cartesian Closed Category $\Downarrow\text{-Strat}$

We organize arenas and P-visible strategies into a cartesian closed category $\Downarrow\text{-Strat}$. Its objects will be arenas, already presented in Section 3.1.3. But the cartesian structure requires more arena constructions beyond the well-opened case shown in Chapter 3.

4.1.1 More Arena Constructions


We start with the *product arena*, used for interpreting contexts:

Definition 4.1.1. Consider A_1, A_2 arenas. Their *product* $A_1 \& A_2$ has components:

$$\begin{array}{llll} \text{moves:} & |A_1 \& A_2| & = & |A_1| + |A_2| \\ \text{causality:} & (i, a) \leq_{A_1 \& A_2} (j, a') & \Leftrightarrow & i = j \ \& \ a \leq_{A_i} a' \\ \text{polarities:} & \text{pol}_{A_1 \& A_2}(i, a) & = & \text{pol}_{A_i}(a), \\ \mathcal{Q}/\mathcal{A}\text{-labeling:} & \lambda_{A_1 \& A_2}(i, a) & = & \lambda_{A_i}(a), \end{array}$$

generalizing to an n -ary construction $\&_{i \in I} A_i$ in the obvious way (for I a finite set).

This construction puts arenas A_1 and A_2 side-by-side, without interaction – Opponent may start evaluation in either component. Like a forest is formed of trees, any arena is – up to isomorphism – a product of well-opened arenas:

Figure 4.1: The arena $U \Rightarrow (U \& U)$

Lemma 4.1.2. *Consider A an arena.*

Then there is $(A_i)_{i \in I}$ of well-opened arenas, unique up to iso, s.t. $A \cong \&_{i \in I} A_i$.

Here, an **isomorphism** of arenas is any forest-isomorphism preserving polarities and \mathcal{Q}/\mathcal{A} -labeling – likewise, two families $(A_i)_{i \in I}$ and $(B_j)_{j \in J}$ are **isomorphic** if there is a bijection $\pi : I \simeq J$ and for each $i \in I$, an isomorphism $A_i \cong B_{\pi(i)}$.

Recall that in Definition 3.2.2, we defined the **arrow** $A \Rightarrow B$ of two arenas A and B with B well-opened. Building a cartesian closed category requires us to extend $A \Rightarrow B$ to the case where B may not be well-opened. However, $A \Rightarrow -$ shall be the right adjoint to $A \times -$, and so must preserve products; so we know that we must have:

$$A \Rightarrow (B \times C) \quad \cong \quad (A \Rightarrow B) \times (A \Rightarrow C),$$

and via Lemma 4.1.2 we may take that as a definition:

Definition 4.1.3. *For $A, B = \&_{i \in I} B_i$ any two arenas, we set:*

$$A \Rightarrow B = \&_{i \in I} A \Rightarrow B_i.$$

This has the effect of creating one copy of A for each minimal move in B^1 – in particular, if B is the empty arena, then so is $A \Rightarrow B$. In Figure 4.1 we show an example illustrating the arrow construction on non well-opened arenas.

4.1.2 Composition of Strategies

Restrictions. We first require notations for *restrictions* of plays of an arena to a sub-component. We will actually define the restriction of a play along an *embedding*:

Definition 4.1.4. *Consider A, B arenas.*

*An **embedding** $f : A \hookrightarrow B$ is any injection $f : |A| \rightarrow |B|$.*

¹This definition is equivalent to the original one [Hyland and Ong, 2000]. One may alternatively avoid the duplication of A by considering arenas to be directed acyclic graphs rather than forests, and having minimal moves of A depend on *all* the minimal moves in B : this is done for instance in [Harmer, 2004]. This is fine in the sense that the two constructions generate the same set of plays (see Definition 3.1.4), but arenas then no longer characterize types up to isomorphism, and the connection with thin concurrent games is blurred out.

We use embeddings to address specific components inside arenas: for instance, there is a canonical embedding $A \hookrightarrow A \times B$, etc. The embeddings we use in practice satisfy many additional properties, but it is not useful to require them explicitly.

We can *restrict* a play along an embedding:

Definition 4.1.5. Consider A, B arenas, $f : A \hookrightarrow B$, and s a pointing sequence on B . The **restriction of s along f** , denoted by $s \upharpoonright f$, has moves the sequence defined by

$$\begin{aligned} \varepsilon \upharpoonright f &= \varepsilon \\ sb \upharpoonright f &= (s \upharpoonright f)a && \text{if } b = f(a), \\ sb \upharpoonright f &= s \upharpoonright f && \text{otherwise.} \end{aligned}$$

Moreover, in $s \upharpoonright f$, s_i points to s_j iff there is a sequence of pointers

$$\dots \quad s_j \quad \dots \quad \cdot \quad \dots \quad \cdot \quad \dots \quad \cdot \quad \dots \quad \cdot \quad \dots \quad s_i$$

in s , with all intermediary points not selected in $s \upharpoonright f$.

The restriction $s \upharpoonright f$ is a pointing sequence on A , but there is no guarantee that it is a legal play on A , even if $s \in \Downarrow\text{-Plays}(B)$. However, that will often be the case when using this notion. This overly general definition has two use cases in practice:

Firstly, situations like restrictions along the canonical embedding $A \hookrightarrow A \times B$, where the transitive clause for pointers is not used at all. For $s \in \Downarrow\text{-Plays}(A \times B)$, we write $s \upharpoonright A \in \Downarrow\text{-Plays}(A)$ for the restriction along the (implicit) canonical embedding, which the reader should always be able to disambiguate easily. Secondly, situations like restrictions along the canonical embedding $A \Rightarrow C \hookrightarrow (A \Rightarrow B) \Rightarrow C$. In that case we write $s \upharpoonright A, C$, which may not always be a valid play on $A \Rightarrow C$ (it may not be alternating). As here, in general we never explicitly specify the embedding used in the restriction, but it should hopefully always be clear from the context.

Interactions. Consider $\sigma : A \Rightarrow B$ and $\tau : B \Rightarrow C$ two strategies. Observe that σ and τ have dual perspectives on the polarities in B : when σ plays a Player move in B , it is an Opponent move for τ , and reciprocally. The first step in defining composition, is to *plug* σ and τ , letting them synchronize on B : this is called the *parallel interaction* of σ and τ . Then we *hide* the synchronized events, leaving a strategy on $A \Rightarrow C$.

We first focus on *parallel interaction*. An *interaction* between σ and τ , resulting of the synchronization of plays of σ and τ , is a pointing sequence that spans A, B and C :

Definition 4.1.6. Consider A, B and C arenas.

An **interaction** on A, B, C is a legal pointing sequence u on $(A \Rightarrow B) \Rightarrow C$ s.t.:

$$\begin{aligned} \text{left-legal:} \quad & u \upharpoonright A, B \in \Downarrow\text{-Plays}(A \Rightarrow B), \\ \text{right-legal:} \quad & u \upharpoonright B, C \in \Downarrow\text{-Plays}(B \Rightarrow C), \\ \text{outer-legal:} \quad & u \upharpoonright A, C \in \Downarrow\text{-Plays}(A \Rightarrow C). \end{aligned}$$

We write $I(A, B, C)$ the set of all interactions on A, B, C .




Figure 4.2: State diagram for sequential interactions

In interactions, moves in A, C are said to be **visible** while moves in B are **hidden** or **synchronized**. One may regard an interaction $u \in \mathbf{I}(A, B, C)$ as involving three players: player ℓ plays in A, B ; player r plays on B, C . Finally, the **external Opponent** plays Opponent on A, C , *i.e.* plays moves that have negative polarity on $A \Rightarrow C$. Positive moves in $A \Rightarrow C$ may be due to ℓ or r ; we say that they are due to the **global Player**, which we think of as ℓ and r working together as a team.

The conditions *outer-legal*, *left-legal* and *right-legal* have a strong impact on how interactions may unfold. Indeed, recall from below Definition 3.1.4 that each alternating play has a *state*: O if it is Opponent's turn to play (if the play has even length), and P otherwise. The three conditions of Definition 4.1.6 involve three restrictions

$$\begin{aligned} u \upharpoonright A, B &\in \Downarrow\text{-Plays}(A \Rightarrow B), \\ u \upharpoonright B, C &\in \Downarrow\text{-Plays}(B \Rightarrow C), \\ u \upharpoonright A, C &\in \Downarrow\text{-Plays}(A \Rightarrow C) \end{aligned}$$

for $u \in \mathbf{I}(A, B, C)$. Accordingly, where the polarity state of a play is summarized by a single letter O or P , the polarity state of an interaction is summed up by three letters, keeping track of the state of these three restrictions, in order. There are strong constraints as to how we can jump between these states: for instance, in state OOO no move can be played in B , as no polarity in B can simultaneously satisfy the two constraints that the next move should be an Opponent move in both $A \Rightarrow B$ and $B \Rightarrow C$. Overall, it is fairly easy to show that only three states are reachable from the empty interaction: OOO , OPP , and POP ; and the possible transitions between those are cleanly summarized in the **state diagram for interactions** in Figure 4.2.

We are often interested by interactions ending with a visible move from the global Player (which yield even-length plays of the composition). Those always have the shape

$$e^-(b \dots b)e^+ \dots e^-(b \dots b)e^+,$$

i.e. sequences of synchronized moves in b sandwiched by a pair of visible moves from the external Opponent, and global Player – this immediately follows from the diagram.

Composition. It is now time to define the composition of two strategies:

Definition 4.1.7. Consider $\sigma : A \Rightarrow B$ and $\tau : B \Rightarrow C$ two strategies.

An **interaction** between σ and τ is any $u \in \mathbf{I}(A, B, C)$ such that $u \upharpoonright A, B \in \sigma$ and




Figure 4.3: $u \in \llbracket x : \mathbb{U} \vdash \lambda y. y, x : \mathbb{U} \rightarrow \mathbb{U} \rrbracket \parallel \llbracket f : \mathbb{U} \rightarrow \mathbb{U} \vdash \lambda x. f(f\ x) : \mathbb{U} \rightarrow \mathbb{U} \rrbracket$.

$u \upharpoonright B, C \in \tau$. We write $\sigma \parallel \tau$ for the set of all interactions between σ and τ .

We may then define the **composition** $\tau \odot \sigma = \{u \vdash A, C \mid u \in \sigma \parallel \tau\} : A \Rightarrow C$.

The process that to an interaction $u \in \sigma \parallel \tau$ associates $u \upharpoonright A, C \in \tau \odot \sigma$ is called **hiding**, hence completing the *parallel interaction plus hiding* paradigm. If $s \in \tau \odot \sigma$, we call any $u \in \sigma \parallel \tau$ such that $s = u \upharpoonright A, C$ a **witness** for s .

We give in Figure 4.3 an example interaction between two innocent strategies. In the diagram, moves are annotated with the player responsible: $-$ if it is the external Opponent, $+$ if it is the global Player, and then ℓ or r . This interaction is a witness to the play of the composition keeping the moves in black only.

When composing alternating strategies, the following property is crucial:

Lemma 4.1.8 (Unique witness). *Consider $\sigma : A \Rightarrow B$ and $\tau : B \Rightarrow C$ strategies.*

For all $s \in \tau \odot \sigma$, there is a unique $u \in \sigma \parallel \tau$ such that $u \upharpoonright A, C = s$.

Proof. Direct via the state diagram for interactions and determinism of σ and τ . \square

The next proposition states that composition is well-defined; and that all the conditions on strategies introduced in the previous chapter are stable under composition.

Proposition 4.1.9. *Take $\sigma : A \Rightarrow B, \tau : B \Rightarrow C$ strategies as in Definition 3.1.5.*

Then, $\tau \odot \sigma : A \Rightarrow C$ is a strategy in the sense of Definition 3.1.5. Moreover,

if σ and τ are P -visible, then so is $\tau \odot \sigma$,
 if σ and τ are well-bracketed, then so is $\tau \odot \sigma$,
 if σ and τ are single-threaded, then so is $\tau \odot \sigma$,
 if σ and τ are innocent, then so is $\tau \odot \sigma$.

Proof. For Definition 3.1.5, *non-empty* and *prefix-closed* are clear, while *receptive* and *determinism* are direct via Lemma 4.1.8 and the state diagram of interactions.

For some of the further conditions, the proof of stability under composition can be fairly elaborate – we omit them, but give references. Composition of P -visibility is proved *e.g.* in [Clairambault, 2010, Proposition 2.2.6]; composition of well-bracketing is proved *e.g.* in [Clairambault, 2010, Proposition 2.3.3]; for single-threadedness, it follows from Harmer’s theorem [Harmer, 1999, Proposition 3.5.5] that those are exactly the comonoid morphisms; and finally, a nice detailed proof of stability of innocence under composition appears in [Harmer, 2004, Proposition 3.3.2]. \square

4.1.3 A Category of Arenas and Strategies

We now provide the categorical structure. Though we shall not give details for any of the required proofs, we shall still review the main arguments and provide references.

Associativity. For associativity, one first introduces the set $I(A, B, C, D)$ of **ternary interactions**: those legal pointing sequences u on $((A \Rightarrow B) \Rightarrow C) \Rightarrow D$ such that

$$\begin{aligned} u \upharpoonright A, B &\in \downarrow\uparrow\text{-Plays}(A \Rightarrow B), \\ u \upharpoonright B, C &\in \downarrow\uparrow\text{-Plays}(B \Rightarrow C), \\ u \upharpoonright C, D &\in \downarrow\uparrow\text{-Plays}(C \Rightarrow D), \\ u \upharpoonright A, D &\in \downarrow\uparrow\text{-Plays}(A \Rightarrow D). \end{aligned}$$

The key argument is the so-called “zipping lemma” ([Harmer, 2004, Lemma 2.6.1]):

Lemma 4.1.10. Consider $u \in I(A, C, D)$ and $v \in I(A, B, C)$ s.t. $u \upharpoonright A, C = v \upharpoonright A, C$. There is a unique $w \in I(A, B, C, D)$ s.t. $w \upharpoonright A, C, D = u$, $w \upharpoonright A, B, C = v$.

This is proved by induction on u . As the name suggests, the idea is to “zip” u and v together, using subsequences from v to provide the moves in B missing from u . There is also a mirror lemma zipping $u \in I(A, B, D)$ and $v \in I(B, C, D)$.

Associativity of composition follows:

Proposition 4.1.11. Consider $\sigma : A \Rightarrow B$, $\tau : B \Rightarrow C$ and $\nu : C \Rightarrow D$.

Then, $\nu \odot (\tau \odot \sigma) = (\nu \odot \tau) \odot \sigma$.

Proof. Consider $s \in \nu \odot (\tau \odot \sigma)$. By definition, there is a witness $u \in (\tau \odot \sigma) \parallel \nu$, i.e. $u \in I(A, C, D)$ s.t. $u \upharpoonright A, C \in \tau \odot \sigma$, $u \upharpoonright C, D \in \nu$, and $u \upharpoonright A, D \in \downarrow\uparrow\text{-Plays}(A \Rightarrow D)$. By definition of composition again, there is $v \in \sigma \parallel \tau$ s.t. $v \upharpoonright A, C = u \upharpoonright A, C$. By the

zipping lemma, we can find $w \in I(A, B, C, D)$ s.t. $w \upharpoonright A, C = v$ and $w \upharpoonright A, C, D = u$; and $w \upharpoonright A, B \in \sigma$, $w \upharpoonright B, C \in \tau$ and $w \upharpoonright C, D \in \nu$.

One can then show that $w \upharpoonright B, D$ is alternating (this follows from a state diagram for ternary interactions extending that for interactions in Section 4.1.2). Hence $w \upharpoonright B, C, D \in \tau \parallel \nu$, hence $w \upharpoonright B, D \in \nu \odot \tau$. Since we also have $w \upharpoonright A, B \in \sigma$, we have $w \upharpoonright A, B, D \in \sigma \parallel (\nu \odot \tau)$, so $w \upharpoonright A, D \in (\nu \odot \tau) \odot \sigma$ as required.

The other inclusion is symmetric, using the mirror zipping lemma. \square

Identities. Identities in game semantics are given by *copycat strategies*.

The intuitive behaviour of copycat is very simple: it copies whatever move Opponent plays on either side, to the matching Player move on the other side, maintaining the invariant that we have the same play on A on both sides. More concretely:

Definition 4.1.12. Consider A any arena.

A *copycat play* on A is any P -visible $s \in \Downarrow\text{-Plays}(A_1 \Rightarrow A_2)$ which is:

$$\text{balanced: } \forall ta^+ \sqsubseteq s, ta \upharpoonright A_1 = ta \upharpoonright A_2,$$

and we define $\mathfrak{c}_A : A \Rightarrow A$ as the set of copycat plays on A .

Copycat is a strategy: most conditions are direct, save for determinism that is a bit subtle. The issue is easy to miss as we tend to leave the management of pointers implicit: by *balanced*, an initial move on the left hand side can only immediately follow an initial move on the right hand side. But its *pointer* is not uniquely determined by *balanced*: if there are several initial moves on the right hand side, each can receive the pointer in a way compatible with *balanced*. This is however banned by the P -visible requirement, as in that case only the previous move appears in the P -view.

Copycat is indeed an identity for composition:

Lemma 4.1.13. Consider $\sigma : A \Rightarrow B$ any strategy. Then, $\mathfrak{c}_B \odot \sigma = \sigma \odot \mathfrak{c}_A = \sigma$.

Proof. By *receptive*, it suffices to prove the equality for even-length plays, i.e. Player-ending plays. Consider $s \in \mathfrak{c}_B \odot \sigma$ with even length – write $\mathfrak{c}_B : B_1 \Rightarrow B_2$ – and consider its witness $u \in \sigma \parallel \mathfrak{c}_B$. By the state diagram for interactions, u is in state OOO , therefore $u \upharpoonright B_1, B_2$ is in state OO and is Player-ending, so $u \upharpoonright B_1 = u \upharpoonright B_2$. It follows that $u \upharpoonright A, B_1 = u \upharpoonright A, B_2 = s$, but $u \upharpoonright A, B_1 \in \sigma$, hence $s \in \sigma$.

Reciprocally, from $s \in \sigma$ with even length it is direct to construct by induction on s an interaction $u \in \sigma \parallel \mathfrak{c}_B$ such that $u \upharpoonright A, B_1 = u \upharpoonright A, B_2 = s$, so that $s \in \mathfrak{c}_B \odot \sigma$ as required. Symmetrically, we have $\sigma \odot \mathfrak{c}_A = \sigma$ as well. \square

We conclude from all of the above:

Corollary 4.1.14. There is a category $\Downarrow\text{-Strat}$ (respectively $\Downarrow\text{-WB}$, $\Downarrow\text{-Inn}$, $\Downarrow\text{-InnWB}$) with arenas as objects, and as morphisms from A to B all P -visible single-threaded strategies (respectively all P -visible well-bracketed strategies, all innocent strategies, all innocent well-bracketed strategies) on $A \Rightarrow B$.

4.1.4 Cartesian Closed Structure

Next, we show that $\Downarrow\text{-Strat}$, $\Downarrow\text{-WB}$, $\Downarrow\text{-Inn}$ and $\Downarrow\text{-InnWB}$ are cartesian closed.

Cartesian structure. The empty arena 1 is terminal in all four categories: for all arena A , $A \Rightarrow 1 = 1$, inhabited with a unique, empty, strategy. The *product* of A and B was defined in Definition 4.1.1. For arenas A, B , the two projections

$$\pi_A : A \times B \Rightarrow A, \quad \pi_B : A \times B \Rightarrow B$$

are the copycat strategies defined with the obvious adaptation of Definition 4.1.12.

To define the pairing, we prove:

Lemma 4.1.15. *Consider Γ, A, B arenas, and $\sigma : \Gamma \Rightarrow A, \tau : \Gamma \Rightarrow B$ single-threaded. There exists a unique (necessarily) single-threaded strategy*

$$\langle \sigma, \tau \rangle : \Gamma \Rightarrow A \times B$$

comprising all $s \in \Downarrow\text{-Plays}(\Gamma \Rightarrow A \times B)$ such that for all $t \sqsubseteq s$, $[t] \in \sigma$ or $[t] \in \tau$.

Proof. Analogously to innocent strategies and P-views, single-threaded strategies are uniquely determined by their set of threads, see e.g. [Harmer, 2004, Section 2.8]. \square

Notice, in this definition, the implicit renaming of $[t] \in \Downarrow\text{-Plays}(\Gamma \Rightarrow A \times B)$ to $\Downarrow\text{-Plays}(\Gamma \Rightarrow A)$ or $\Downarrow\text{-Plays}(\Gamma \Rightarrow B)$ depending whether the initial move is in A or B .

This data satisfies the required equations for cartesian products:

Lemma 4.1.16. *Consider $\sigma : \Gamma \Rightarrow A, \tau : \Gamma \Rightarrow B$, and $\nu : \Delta \Rightarrow \Gamma$ single-threaded. Then, the following equations hold:*

$$\begin{aligned} \pi_A \odot \langle \sigma, \tau \rangle &= \sigma \\ \pi_B \odot \langle \sigma, \tau \rangle &= \tau \\ \langle \pi_A, \pi_B \rangle &= \mathbf{c}_{A \times B} \\ \langle \sigma, \tau \rangle \odot \nu &= \langle \sigma \odot \nu, \tau \odot \nu \rangle. \end{aligned}$$

Proof. The first two equations are direct variations of Lemma 4.1.13. The third is direct from the definition, and the fourth immediate on threads, using again that single-threaded strategies are uniquely determined by their sets of threads. \square

From the equations above, the following fact follows immediately.

Corollary 4.1.17. *The categories $\Downarrow\text{-Strat}$, $\Downarrow\text{-WB}$ and $\Downarrow\text{-Inn}$ are cartesian.*

Cartesian closed structure. Now, we describe the cartesian closure.

For any two arenas A and B , we have already defined their arrow $A \Rightarrow B$. We have:

Lemma 4.1.18. *Consider Γ , A , and B three arenas. Then we have*

$$\Xi_{\Gamma,A,B} : (\Gamma \times A) \Rightarrow B \cong \Gamma \Rightarrow (A \Rightarrow B)$$

an isomorphism of arenas, i.e. a bijection preserving and reflecting all structure.

Proof. Straightforward from the definition. \square

Up to renaming, $\Gamma \times A \Rightarrow B$ and $\Gamma \Rightarrow A \Rightarrow B$ are the same. Thus this means that the currying of $\sigma : \Gamma \times A \Rightarrow B$ can be described by applying $\Xi_{\Gamma,A,B}$ move-by-move:

Lemma 4.1.19. *Consider Γ , A and B three arenas. Then, there is a bijection*

$$\Lambda_{\Gamma,A,B} : \Downarrow\text{-Strat}(\Gamma \times A, B) \simeq \Downarrow\text{-Strat}(\Gamma, A \Rightarrow B)$$

preserving and reflecting innocence, and preserving composition in the sense that

$$\Lambda_{\Gamma,A,B}(\sigma) \odot \gamma = \Lambda_{\Delta,A,B}(\sigma \odot (\gamma \times A)) \quad (4.1)$$

for all $\sigma \in \Downarrow\text{-Strat}(\Gamma \times A, B)$ and $\gamma \in \Downarrow\text{-Strat}(\Delta, \Gamma)$.

Proof. A routine manipulation on interactions. \square

In turn, this bijection lets us define the evaluation strategy simply as

$$\text{ev}_{A,B} = \Lambda_{A,A,B}^{-1}(\mathbf{c}_{A \Rightarrow B}) \in \Downarrow\text{-Strat}((A \Rightarrow B) \times A, B)$$

and from the above, it becomes routine to prove cartesian closure:

Proposition 4.1.20. $\Downarrow\text{-Strat}$, $\Downarrow\text{-WB}$, $\Downarrow\text{-Inn}$ and $\Downarrow\text{-InnWB}$ are cartesian closed.

Proof. We need the two equalities for the universal property of cartesian closure, for $\sigma \in \Downarrow\text{-Strat}(\Gamma \times A, B)$, $\gamma \in \Downarrow\text{-Strat}(\Delta, \Gamma)$ and $\tau \in \Downarrow\text{-Strat}(\Gamma, A \Rightarrow B)$:

$$\text{ev}_{A,B} \odot (\Lambda(\sigma) \times A) = \sigma \quad (4.2)$$

$$\Lambda(\text{ev}_{A,B} \odot (\tau \times A)) = \tau, \quad (4.3)$$

but those follow from elementary computation using definition of $\text{ev}_{A,B}$, (4.1), neutrality of copycat for composition and the fact that Λ is a bijection. \square

Since $\Lambda(-)$ preserves and reflects innocence and well-bracketing, the cartesian closed structure automatically holds for our four categories of interest.

Internal language. This cartesian closed structure allows us to use notations from the λ -calculus to manipulate strategies: these may be unfolded to primitive operations of cartesian closed categories in the standard way [Lambek and Scott, 1988]. For instance, the *application* of $\sigma \in \Downarrow\text{-Strat}(\Gamma, A \Rightarrow B)$ to $\tau \in \Downarrow\text{-Strat}(\Gamma, A)$ unfolds to

$$\sigma \tau = \text{ev}_{A,B} \odot \langle \sigma, \tau \rangle,$$

likewise we will sometimes in the sequel use λ -abstraction on expressions for strategies.

4.1.5 Recursion

To handle recursion we must provide a *fixpoint operator*, *i.e.* an innocent strategy

$$\mathcal{Y}_A : (A \Rightarrow A) \Rightarrow A$$

for each arena A such that for every $\sigma \in \Downarrow\text{-Strat}(\Gamma, A \Rightarrow A)$, we have

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

a fixpoint equation. It is solved as usual in denotational semantics, by constructing \mathcal{Y}_A as a least upper bound for an adequate order on strategies – the following is direct:

Proposition 4.1.21. *For arenas A and B , $\Downarrow\text{-Strat}(A, B)$, $\Downarrow\text{-WB}(A, B)$, $\Downarrow\text{-Inn}(A, B)$ and $\Downarrow\text{-InnWB}(A, B)$ are partially ordered by inclusion, yielding pointed dcpos.*

All operations on strategies involved in the cartesian closed structure are continuous, making $\Downarrow\text{-Strat}$, $\Downarrow\text{-WB}$, $\Downarrow\text{-Inn}$ and $\Downarrow\text{-InnWB}$ enriched over pointed dcpos.

Using this, we can obtain the fixpoint operator as the least upper bound of

$$\begin{array}{lll} F & : & \Downarrow\text{-Strat}(A \Rightarrow A, A) \rightarrow \Downarrow\text{-Strat}(A \Rightarrow A, A) \\ & & \sigma \mapsto \lambda f^{A \Rightarrow A}. f(\sigma f) \end{array}$$

which unfolds to the following, for \perp the minimal strategy just closed under receptivity:

$$\mathcal{Y}_A = \bigcup_{n \in \mathbb{N}} F^n(\perp) \in \Downarrow\text{-Strat}(A \Rightarrow A, A) \quad (4.4)$$

satisfying $\mathcal{Y}_A = F \mathcal{Y}_A$ by construction – we may then prove the fixpoint equation:

$$\begin{aligned} \mathcal{Y}_A \sigma &= F(\mathcal{Y}_A) \sigma \\ &= (\lambda f^{A \Rightarrow A}. f(\mathcal{Y}_A f)) \sigma \\ &= \sigma (\mathcal{Y}_A \sigma) \end{aligned}$$

using β -equivalence on the λ -calculus notation [Lambek and Scott, 1988].

4.2 Interpretation of IA

Following [Lambek and Scott, 1988], the cartesian closed structure developed above yields an interpretation of the simply-typed λ -calculus; we also showed the interpretation of the fixpoint combinator. In order to obtain the interpretation of PCF and IA, it remains to provide strategies matching the primitives of our programming languages.

4.2.1 Interpretation of PCF

Let us focus first on the purely functional primitives, *i.e.* on the interpretation of PCF.

Interpretation of types. First of all, to every type A we associate an arena: this is done exactly as in Section 3.2.1, by setting $\llbracket \mathbb{B} \rrbracket = \mathbf{B}$, $\llbracket \mathbb{N} \rrbracket = \mathbf{N}$ and $\llbracket \mathbb{U} \rrbracket = \mathbf{U}$ the arenas of 3.8, 3.9 and 3.7 respectively. Arrow types are interpreted as $\llbracket A \rightarrow B \rrbracket = \llbracket A \rrbracket \Rightarrow \llbracket B \rrbracket$ using the arrow arena construction of Definition 4.1.3, yielding

$$\llbracket A \rrbracket \in \Downarrow\text{-}\mathbf{Strat}$$

an arena for every type A of PCF. This is extended to contexts by

$$\llbracket x_1 : A_1, \dots, x_n : A_n \rrbracket = \big\&_{1 \leq i \leq n} \llbracket A_i \rrbracket.$$

Interpretation of terms. The general methodology of the interpretation of terms is that of the interpretation of a simply-typed λ -calculus in a cartesian closed category. This means that a term² $\Gamma \vdash M : A$ gets interpreted as a morphism in $\Downarrow\text{-}\mathbf{Strat}$:

$$\llbracket M \rrbracket \in \Downarrow\text{-}\mathbf{Strat}(\llbracket \Gamma \rrbracket, \llbracket A \rrbracket).$$

The primitives for the simply-typed λ -calculus are covered by the cartesian closed structure of $\Downarrow\text{-}\mathbf{Strat}$ (see [Lambek and Scott, 1988] for details). For recursion, we set

$$\llbracket \mathcal{Y} M \rrbracket = \mathcal{Y}_{\llbracket A \rrbracket} \odot \llbracket M \rrbracket$$

for $\Gamma \vdash M : A \rightarrow A$, using the strategy \mathcal{Y}_A of (5.2). It remains to provide interpretations for all the additional combinators that PCF adds to the simply-typed λ -calculus.

First of all, all constants are interpreted by the obvious strategy directly returning the corresponding value. In Figure 4.4, we show the typical plays of strategies for sequential composition, conditionals, successor, and zero test; while Figure 4.6 shows typical plays for the predecessor, and Figure 4.7 shows the typical plays for the let binding strategy. Using these strategies, the strategy is obtained by the clauses in Figure 4.5.

4.2.2 Interpretation of State

We complete the interpretation to provide semantics for full IA.

Interpretation of types. Before we describe the interpretation of stateful primitives, we must provide arenas for the two additional types \mathbb{V} and \mathbb{S} that IA adds to PCF.

Intuitively, these arenas describe the interface through which a program may interact with a reference or a semaphore. For instance, a reference may be written to; or read from – a semaphore may be grabbed and released. It is convenient to split the arena

²Or really, a typing derivation – but type annotations in PCF ensure that each term has a unique derivation.




Figure 4.4: Typical plays of basic strategies for PCF








Figure 4.5: Basic interpretation clauses

 Figure 4.6: Strategy for **pred**

 Figure 4.7: Typical plays for $\text{let}_{X,Y}$

Figure 4.8: V_w Figure 4.9: V_r Figure 4.10: S

for references in two: V_w describes the protocol for *writing* to the memory, and V_r for *reading* from the memory. Formally, these and S for semaphores are defined by:

$$V_w = \bigotimes_{n \in \mathbb{N}} U, \quad V_r = N, \quad S = U \& U \quad (4.5)$$

with moves renamed as illustrated in Figures 4.8, 4.9 and 4.10; hopefully making the intention behind the moves self-explanatory. Finally, we pair writing and reading with

$$V = V_w \& V_r,$$

and $\llbracket V \rrbracket = V$, $\llbracket S \rrbracket = S$ completes the interpretation of all IA types.

Queries. Next, we must provide strategies for the primitives corresponding to state: manipulation and declaration of variables and semaphores. The usual methodology in game semantics is to separate these primitives in two categories. On the one hand, we have the primitives performing queries to the state: assignment, dereferenciation, grab and release. Those are viewed as not inherently stateful, as they simply propagate the operations made by the program to the central memory. On the other hand, reference and semaphore declaration actually introduce the stateful behaviour.

We start with the former, by providing the interpretation of assignment, dereferenciation, grab and release. This is done via the four following clauses:

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

using the strategies of Figures 4.11, 4.12, 4.13, and 4.14. For the most part those are simply copycat strategies – or more precisely, following (4.5), projections. In particular we insist that they are *innocent strategies*: they are not responsible for the side effect, as their role is simply to forward queries to the actual memory.

Declaration of references. Now, we must provide the interpretation for the declaration of new references, performing the actual side effect.

The *memory cell* implements the stateful behaviour: it acknowledges the write requests; when queried with a read request it outputs the value provided with the latest




Figure 4.11: assign Figure 4.12: deref Figure 4.13: grab Figure 4.14: release

write. At first ignoring pointers, it seems natural to first capture this behaviour as a language (as in, a set of words) called *cell*, on moves of \mathbb{V} . Hence we define:

$$\text{cell}_n ::= \varepsilon \mid \mathbf{wk}^- \checkmark^+ \text{cell}_k \mid \mathbf{r}^- n^+ \text{cell}_n$$

and we define *cell* as the prefix language of cell_0 (matching the fact that new references are initialized to 0). This definition omits pointers, but they are easily reinstated: Opponent moves are initial, while Player moves point to the previous move.

It is tempting to take $\text{cell} : \mathbb{V}$, and for $x : \mathbb{V} \vdash M : A$ a term of \mathbf{IA} , to define

$$\llbracket \text{newref } x \text{ in } M \rrbracket = \llbracket M \rrbracket \odot \text{cell},$$

however we run into an issue: *cell* is not a strategy, since it is not single-threaded. But actually, of course it cannot be single-threaded: single-threadedness ensures that each call is independent, whereas the very nature of *cell* forces it to link the different queries! To make the definition above single-threaded, we seek to define instead

$$\text{newr}_{\mathbf{X}} : (\mathbb{V} \Rightarrow \mathbf{X}) \Rightarrow \mathbf{X},$$

ensuring that all variable queries belong to the same thread. Fixing \mathbf{X} for now, we set

$$\text{newrt}_n ::= \varepsilon \mid \mathbf{wk}^- \checkmark^+ \text{newrt}_k \mid \mathbf{r}^- n^+ \text{newrt}_n \mid v_1^- v_2^+ \text{newrt}_n$$

and set *newrt* as the prefix language of $q_2^- q_1^+ \text{newrt}_0$, where the annotations 1 and 2 help distinguish components, writing $(\mathbb{V} \Rightarrow \mathbf{X}_1) \Rightarrow \mathbf{X}_2$. To *cell*, *newrt* adds two initial questions q_2^- and q_1^+ and is prepared at any point to propagate an answer v_1^- for q_1^+ to the corresponding answer v_2^+ for q_2^- – this does not necessarily end the play, and it may happen several times as Opponent may not be well-bracketed. Finally, these plays lack *pointers* – but they are uniquely determined so as to make the plays *P*-visible.

This defines *newrt*, a set of *threads*; the corresponding single-threaded strategy is:

Proposition 4.2.1. *Consider \mathbf{X} a basic arena, i.e. one of \mathbf{U} , \mathbf{B} , and \mathbf{N} .*

Then, there is a unique single-threaded strategy

$$\text{newr}_{\mathbf{X}} : (\mathbb{V} \Rightarrow \mathbf{X}) \Rightarrow \mathbf{X}$$

such that for all $s \in \text{newr}_{\mathbf{X}}$, we have $\lceil s \rceil \in \text{newrt}_{\mathbf{X}}$.

$$\begin{aligned}
\llbracket \mathbf{newref} \ x := n \text{ in } M \rrbracket &= \mathbf{newr}_X \odot \Lambda(\llbracket M \rrbracket) \in \Downarrow\text{-Strat}(\llbracket \Gamma \rrbracket, \llbracket A \rrbracket) \\
\llbracket \mathbf{newsem} \ x := n \text{ in } M \rrbracket &= \mathbf{news}_X \odot \Lambda(\llbracket M \rrbracket) \in \Downarrow\text{-Strat}(\llbracket \Gamma \rrbracket, \llbracket A \rrbracket).
\end{aligned}$$

Figure 4.15: Interpretation of reference and semaphore declaration

After each Opponent move, there is exactly one active thread: namely, that where Opponent just moved. The single-threaded strategy \mathbf{newr}_X then simply plays as \mathbf{newrt}_X in the active thread, ignoring other threads and never switching between them. Operationally, this corresponds to the fact that state in **IA** is *local*: each execution of the new reference operation is linked to a separate, dynamically allocated memory cell.

Finally, the interpretation of **newref** is given in Figure 4.15 for $\Gamma, x : \mathbf{V} \vdash M : \mathbf{X}$.

Declaration of semaphores. Now, we perform the same construction for semaphores. The definition is analogous and semaphores are not particularly interesting in a sequential language, so we offer fewer details. The *semaphore cell* is defined first via:

$$\begin{aligned}
\mathbf{lock}_0 &::= \varepsilon \mid \mathbf{g}^- \checkmark^+ \mathbf{lock}_1 \mid \mathbf{d}^- \\
\mathbf{lock}_{n+1} &::= \varepsilon \mid \mathbf{d}^- \checkmark^+ \mathbf{lock}_0 \mid \mathbf{g}^-,
\end{aligned}$$

setting \mathbf{lock} as the prefix language of \mathbf{lock}_0 – pointers are redundant. Observe that in a sequential language, attempting to grab a semaphore that is not free results in divergence; likewise attempting to release a semaphore that is already free yields divergence.

Again, that is not single-threaded. For the actual single-threaded strategy, we set

$$\begin{aligned}
\mathbf{newst}_0 &::= \varepsilon \mid \mathbf{g}^- \checkmark^+ \mathbf{newst}_1 \mid v_1^- v_2^+ \mathbf{newst}_0 \mid \mathbf{d}^- \\
\mathbf{newst}_{n+1} &::= \varepsilon \mid \mathbf{d}^- \checkmark^+ \mathbf{newst}_0 \mid v_1^- v_2^+ \mathbf{newst}_{n+1} \mid \mathbf{g}^-
\end{aligned}$$

and set \mathbf{newst}_X as the prefix language of \mathbf{newst}_0 , where the subscripts 1 and 2 help distinguish components on $(\mathbf{S} \Rightarrow \mathbf{X}_1) \Rightarrow \mathbf{X}_2$. Finally, as for references, we set:

Proposition 4.2.2. *Consider \mathbf{X} a basic arena, i.e. one of \mathbf{U} , \mathbf{B} , and \mathbf{N} .*

Then, there is a unique single-threaded strategy

$$\mathbf{news}_X : (\mathbf{V} \Rightarrow \mathbf{X}) \Rightarrow \mathbf{X}$$

such that for all $s \in \mathbf{news}_X$, we have $\lceil s \rceil \in \mathbf{newst}_X$.

Finally, the interpretation of **newsem** is given in Figure 4.15 for $\Gamma, x : \mathbf{S} \vdash M : \mathbf{X}$. This almost concludes the interpretation of **IA**, but two final constructions remain.

Bad variables and semaphores. Indeed, recall from Section 2.2.1 that **IA** includes

$$\frac{\Gamma \vdash M : \mathbf{N} \rightarrow \mathbf{U} \quad \Gamma \vdash N : \mathbf{N}}{\Gamma \vdash \mathbf{mkvar} \ M \ N : \mathbf{V}} \qquad \frac{\Gamma \vdash M : \mathbf{U} \quad \Gamma \vdash N : \mathbf{U}}{\Gamma \vdash \mathbf{mksem} \ M \ N : \mathbf{S}}$$

the so-called *bad variable* and *bad semaphore* constructs.

Bad variables and semaphores are not sensible programming primitives, they are here by necessity for definability. Indeed, up to isomorphism \mathbf{V} is defined as $(\&_{n \in \mathbb{N}} \mathbf{U}) \& \mathbf{N}$, and strategies inhabiting it may not be linked with actual memory cells: they may be formed by pairing arbitrary “write” and “read” methods. Thus, to obtain a finite definability result, one must include such pairings should be formable in the syntax³.

Through the isomorphisms $\mathbf{V} \cong (\&_{n \in \mathbb{N}} \mathbf{U}) \& \mathbf{N}$ and $\mathbf{S} \cong \mathbf{U} \& \mathbf{U}$, the interpretation is:

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

which are the final two clauses of the interpretation of \mathbf{IA} .

4.3 Complements on Alternating Strategies

In this section, we introduce a few complements on the results and constructions of this chapter. Those may be skipped on first reading; they are intended either as a reference, or because they are referred to further along in this monograph.

4.3.1 Complements on Single-Threadedness

First of all, we have introduced in Section 3.3.1 the notion of single-threaded strategies, as a way to enforce that distinct initial moves generate separated copies of the strategy, which may not communicate or share any state. Here we give some complements.

Here we focus on the concept of *threads*:

Definition 4.3.1. An *alternating thread* on arena A is $t \in \Downarrow\text{-Plays}(A)$ with at most one initial move; write $\Downarrow\text{-Plays}_\bullet(A)$ for the set of alternating threads on A .

We mean to make formal that if $\sigma : A$ is single-threaded, then it is determined by its set of threads – so it should be equivalently described by a *thread-strategy*:

Definition 4.3.2. For A an arena, an *alternating thread-strategy* $\sigma : A$ is a set $\sigma \subseteq \Downarrow\text{-Plays}_\bullet(A)$ satisfying the following conditions:

- non-empty: $\epsilon \in \sigma$,
- prefix-closed: $\forall s \sqsubseteq t \in \sigma, s \in \sigma$,
- receptive: $\forall s \in \sigma, sa^- \in \Downarrow\text{-Plays}_\bullet(A) \Rightarrow sa^- \in \sigma$,
- deterministic: $\forall sa^-b_1^+, sa^-b_2^+ \in \sigma, sab_1 = sab_2$,

and it is additionally *well-bracketed* if it satisfies

well-bracketing: for all $ta^+ \in \sigma$, if t is well-bracketed then so is ta .

³A definability result without bad variables can be obtained by switching to *nominal games* [Tzevelekos, 2009], where nominal techniques are employed to attach reference cells to actual memory locations, but the setting gets more complicated. It must also be mentioned that though finite definability fails for \mathbf{IA} without bad variables, full abstraction still holds as proved by McCusker [McCusker, 2003]

This is, of course, the same conditions as for Definition 3.1.5, except than they are taken within $\Downarrow\text{-Plays}_\bullet(A)$ rather than $\Downarrow\text{-Plays}(A)$. Clearly, if $\sigma : A$ is an alternating strategy (single-threaded or not), then $\sigma \cap \Downarrow\text{-Plays}_\bullet(A)$ is a thread-strategy.

Reciprocally, we construct a single-threaded strategy from a thread-strategy on A :

Proposition 4.3.3. *Consider σ a thread-strategy on arena A .*

Then, constructing the set $\sigma^!$ comprising all plays $s \in \Downarrow\text{-Plays}_\bullet(A)$ which satisfy

threaded-correct: for all $ta^+ \sqsubseteq s$, we have $\lceil ta^+ \rceil \in \sigma$,

it follows that $\sigma^!$ is a single-threaded strategy on A .

Proof. First, the conditions *non-empty*, *prefix-closed* and *receptive* are clear from the definition. For the rest, we first observe that any $s \in \sigma^!$ is *well-threaded*, in the sense that for all $ta^-b^+ \sqsubseteq s$, b points within $\lceil sa^- \rceil$. Indeed, for each initial move s_i in s , we have $s \upharpoonright i$, comprising the moves of s hereditarily justified by s_i , a thread which by condition *thread-correct* must be in σ thus, in particular, alternating. For each $t \sqsubseteq s$ we may define $t \upharpoonright i$ a prefix of $s \upharpoonright i$. Since each $t \upharpoonright i$ is alternating, it is immediate by induction that for $t \sqsubseteq s$ of even length, all $t \upharpoonright i$ have even length, while if t has odd length, exactly one i_0 is such that $t \upharpoonright i_0$ has odd length. But then, considering $ta^-b^+ \sqsubseteq s$, at ta^- the only thread with odd length must be $\lceil ta^- \rceil$, which b^+ must extend so as to satisfy our invariant; thus b^+ must indeed point within $\lceil sa^- \rceil$.

We prove further conditions. For *deterministic*, if $sa^-b_1^+, sa^-b_2^+ \in \sigma^!$, then we have that $\lceil sa^-b_1^+ \rceil, \lceil sa^-b_2^+ \rceil \in \sigma$, but by our observation above those have the form $ta^-b_1^+, ta^-b_2^+ \in \sigma$ for $ta^- = \lceil sa^- \rceil$. Hence $ta^-b_1^+ = ta^-b_2^+$ by *determinism* of σ , from which follows $sa^-b_1^+ = sa^-b_2^+$. Finally, conditions *well-threaded* and *single-threaded* of Definition 3.3.2 follow directly from our observation above. \square

Any thread-strategy yields a single-threaded strategy, but reciprocally all single-threaded strategies have this form, as we show now:

Proposition 4.3.4. *Consider $\sigma : A$ a single-threaded strategy. Then, the set*

$$\llbracket \sigma \rrbracket = \{ \lceil s \rceil \mid s \in \sigma \} = \sigma \cap \Downarrow\text{-Plays}_\bullet(\sigma)$$

is a thread-strategy on A , such that $\sigma = \llbracket \sigma \rrbracket^!$.

Proof. It is direct that if $\sigma : A$ is single-threaded, then $\llbracket \sigma \rrbracket = \sigma \cap \Downarrow\text{-Plays}_\bullet(A)$; from which it is easy that $\llbracket \sigma \rrbracket$ satisfies the conditions of Definition 4.3.2.

Now, we prove that $\sigma = \llbracket \sigma \rrbracket^!$. For \subseteq , consider $s \in \sigma$ and $ta^+ \sqsubseteq s$. Since $ta^+ \in \sigma$, we have $\lceil ta^+ \rceil \in \llbracket \sigma \rrbracket$ by definition, hence $s \in \llbracket \sigma \rrbracket^!$. For \supseteq , consider $s \in \llbracket \sigma \rrbracket^!$. It is then direct to prove by induction on s that $s \in \sigma$: for negative extensions, this follows from *receptivity*, for positive extensions from *single-threaded*. \square

Thus altogether, this gives single-threaded alternating strategies two representations: as a single-threaded set of plays with many initial moves, or as a set of threads. This exactly reflects the two presentations of innocent strategies, as sets of plays or as sets of P-views, that was described in Section 3.2.4. We observe a useful consequence:

Corollary 4.3.5. *Consider $\sigma, \tau : A$ single-threaded strategies. If σ and τ have the same threads, then they are equal.*

Proof. Straightforward from Proposition 4.3.4. \square

Further conditions. Finally, we show how this correspondence preserves the further conditions we considered on alternating strategies. Observe first that the notions of *well-bracketing* and *P-visibility*, respectively introduced in Definitions 3.2.6 and 3.2.8, are conditions on plays and as such transparently apply to thread-strategies. Then:

Lemma 4.3.6. *Consider $\sigma : A$ a thread-strategy. Then:*

- (1) $\sigma^!$ is well-bracketed iff σ is well-bracketed,
- (2) $\sigma^!$ is P-visible iff σ is P-visible.

Proof. (1) *Only if.* Straightforward since $\sigma \subseteq \sigma^!$. *If.* We show by induction on s that for all $s \in \sigma^!$, the pending question is always in the current thread – the result follows.

(2) *Only if.* As above, follows from $\sigma \subseteq \sigma^!$. *If.* Follows from the easy observation that the P-view is always contained within the current thread. \square

4.3.2 Factorization and Definability

With the details of the interpretation in place, we revisit some of the properties of the semantics presented in Chapter 3, providing key steps of the definability proof.

Innocent definability. First of all, recall that all finite innocent well-bracketed strategies are definable within PCF (Theorem 3.2.14). But this result only concerns PCF types – we must first generalize it to all IA types, including references and semaphores:

Theorem 4.3.7. *Let A be an IA type, and $\sigma : \llbracket A \rrbracket$ be finite well-bracketed innocent. There is an IA term $\vdash M : A$, without **newref** or **newsem**, s.t. $\llbracket M \rrbracket = \sigma$.*

Proof. Direct variation of Theorem 3.2.14, crucially using **mkvar** and **mksem**. \square

Factorization. Thus, extending types from PCF to IA does not affect the innocent definability result. But how can it be generalized to *non-innocent* strategies?

For IA, the main argument is a *factorization* result [Abramsky and McCusker, 1996]:

Proposition 4.3.8. *Consider an arena $A = A_1 \Rightarrow \dots \Rightarrow A_n \Rightarrow \mathbf{X}$, and $\sigma : A$ visible, well-bracketed, single-threaded. Then there is an innocent $\mathbf{Inn}(\sigma) : \mathbf{V} \Rightarrow A$ such that*

$$\mathbf{newr}_A \odot \mathbf{Inn}(\sigma) = \sigma$$

for \mathbf{newr}_A obtained from $\mathbf{newr}_{\mathbf{X}}$ via the cartesian closed structure in the obvious way.

Moreover, if σ is finite (resp. computable), then so is $\mathbf{Inn}(\sigma)$.

Sketch. As an innocent strategy, $\mathbf{Inn}(\sigma)$ can only depend on the P-view and hence cannot act as σ , which depends on the full play. But the idea is that $\mathbf{Inn}(\sigma)$ can access the full play indirectly, by maintaining in the reference (an encoding of) the full play. More precisely, each time it receives an Opponent move a^- , $\mathbf{Inn}(\sigma)$ first reads the reference and obtains (the encoding of) a play $s \in \sigma$. Assuming $sa^-b^+ \in \sigma$, Player stores (the encoding of) sa^-b^+ in the reference, then plays b^+ . \square

This provides the missing argument for finite definability (and hence intensional full abstraction) for IA: if $\llbracket A \rrbracket = A_1 \Rightarrow \dots \Rightarrow A_n \Rightarrow \mathbf{X}$ and $\sigma : \llbracket A \rrbracket$ is a finite visible, well-bracketed and single-threaded strategy, then one first obtains a finite innocent and well-bracketed $\mathbf{Inn}(\sigma) : \mathbf{V} \Rightarrow \llbracket A \rrbracket$ by Proposition 4.3.8. By Theorem 4.3.7, there is $\vdash M : \mathbf{V} \rightarrow A$ in IA such that $\llbracket M \rrbracket = \mathbf{Inn}(\sigma)$, but then, as required,

$$\llbracket \lambda x_1 \dots x_n. \mathbf{newref} \, r \text{ in } M \, x_1 \dots x_n \, r \rrbracket = \sigma : \llbracket A \rrbracket$$

concluding the proof of finite definability for IA – note that if Theorem 4.3.7 is extended to a *universality result* as is expected following Conjecture 3.2.15, we get that every *computable* P-visible well-bracketed single-threaded strategy in IA is definable.

4.3.3 Summary of Results

Now that we have introduced the proper categories supporting the interpretations underlying Chapter 3, we restate for future reference the main theorems of that chapter with proper reference to the target categories of the interpretations:

Proposition 4.3.9. *There is a computationally adequate interpretation of IA in $\Downarrow\text{-Strat}$.*

Proof. First stated as Proposition 4.3.9. \square

Theorem 4.3.10. *The interpretation of Proposition 4.3.9 refines into an interpretation of PCF in $\Downarrow\text{-InnWB}$, which is intensionally fully abstract.*

Proof. First stated as Theorem 3.2.16. \square

Theorem 4.3.11. *The interpretation of Proposition 4.3.9 refines into an interpretation of IA in $\Downarrow\text{-WB}$, which is intensionally fully abstract.*

Proof. First stated as Theorem 3.3.5. \square

4.4 Conclusions and Historical Notes

The compositional setting presented here is robust to a number of variations: for instance, *determinism* can be relaxed, obtaining an intensionally fully abstract model for an extension of IA with non-deterministic choice⁴ [Harmer and McCusker, 1999]. Non-deterministic strategies can additionally be weighted by *probabilities*, yielding a model intensionally fully abstract for an extension of IA with probabilistic choice

⁴That is not the case with AJM games – determinism is crucial in the way AJM games handle uniformity.

[Danos and Harmer, 2000]. However, neither of these two extensions support a notion of innocence banning state: they do not restrict to intensionally fully abstract models for either non-deterministic or probabilistic PCF. Among other things, this is what this monograph aims to solve, as motivated in the introduction.

Beyond those non-deterministic and probabilistic extensions, this alternating compositional setting presented here has seen so many extensions and applications that it would be a daunting task to attempt to survey them exhaustively, and we leave this beyond the scope of this short conclusion. Many extensions keep the general structure of plays and strategies, changing mostly the arenas: a crucial example of that is game semantics for call-by-value [Abramsky and McCusker, 1997, Honda and Yoshida, 1999], or for polarized versions of linear logic [Laurent, 2002]. Other extensions build on the present compositional mechanism, but bring additional structure to moves and/or plays, such as *names* [Tzevelekos, 2009] or *exception pointers* [Laird, 2001a].

A more drastic change is to abandon *alternation* in order to model concurrent computation [Laird, 2001b, Ghica and Murawski, 2008], but this does affect rather deeply the compositional machinery presented in this chapter: we present this next.

Chapter 5

Non-Alternating Game Semantics

In this chapter, we continue our introduction to game semantics by presenting the non-alternating version of the games model seen so far. The chapter roughly follows the outline of the previous two chapters: first, we introduce non-alternating plays and strategies appealing to operational intuitions in the style of Chapter 3. Then, we introduce the category $\mathcal{O}\text{-Strat}$, describe the interpretation of \mathbf{IA}_{\parallel} , and prove its properties.

The constructions of this chapter mostly follow [Ghica and Murawski, 2008].

5.1 Concurrency and Non-Alternation

Thus, let us start again from the intuitions described in Section 3.1, representing interactive executions as plays, and ask ourselves the following question: how should plays be amended if execution is no longer sequential, but becomes concurrent?

5.1.1 Non-Alternating Plays

Non-alternation. As a first, example, consider the following program M from \mathbf{PCF}_{\parallel}

$$x : \mathbb{N}, y : \mathbb{N} \vdash \text{let } \begin{pmatrix} x_1 & = & x \\ x_2 & = & y \end{pmatrix} \text{ in } x_1 + x_2 : \mathbb{N},$$

which evaluates x and y in parallel, then returns their sum. Or more precisely, the intention of the program is to evaluate x and y in parallel, but really the execution order depends on the scheduler. It may still be the case that x is evaluated before y , as in the first example of Section 3.1.1. Or, perhaps, the scheduler will evaluate y before x . Or



O : “What is the return value?”
P : “What is the value of x ?”
P : “What is the value of y ?”
O : “The value of x is 12.”
O : “The value of y is 30.”
P : “The value of $x + y$ is 42.”

This naturally brings us to the following variation over Definition 3.1.4:

This is simply Definition 3.1.4 without the condition *alternating*. In other words, *non-alternating pre-play* is a synonym for *legal pointing string*, though we prefer to keep the names separate for conceptual clarity.

Plays and executions. Non-alternating pre-plays are represented with the same conventions as alternating plays. As for alternating plays, the non-alternating pre-plays of a term may be thought of operationally, as illustrated in Figure 5.1.

This example also shows that even for programs which do not use parallel evaluation as a primitive, the switch from alternating plays to non-alternating plays is impactful: plays may also describe executions where it is Opponent who first plays concurrently. In Figure 5.1, Opponent, playing for the argument $f : \mathbb{U} \rightarrow \mathbb{U} \rightarrow \mathbb{U}$, simultaneously calls its two arguments. This triggers a *race*, as the two write requests are being evaluated simultaneously, yielding a non-deterministic result. So in non-alternating games, terms from IA – a deterministic language – may still have non-deterministic behaviour.

Figure 5.1 is misleading in one respect: it makes it seem like the Player move \checkmark^+ *coincides* with the memory update. In reality, it is an acknowledgment for which we can only know for sure that it comes *after* the memory update: the message may stay stuck





Figure 5.2: Alternative result

in a buffer for a while. The behaviour in Figure 5.2, with exactly the same play until the final result, is just as possible. The consequence of this is that the order in which consecutive acknowledgments arrive is *not* representative of the order in which memory operations are handled: it is irrelevant and cannot influence subsequent behaviour.

Well-bracketing. As in sequential games, constructing the games model will require us to introduce a notion of *well-bracketing*. In our presentation of sequential games, *plays* were not well-bracketed; instead well-bracketing was a property of *strategies*¹. Here we make the opposite choice and impose it directly on plays: defining cleanly non-alternating strategies for the primitives of IA_{\parallel} without well-bracketing becomes very challenging, as we are overwhelmed by the number of plays.

As in the alternating case, non well-bracketed plays allow behaviours typical of control operators such as **callcc** – with the new addition of parallel versions, such as the **fork** operator on \mathbb{B} that answers **tt** and **ff** *in parallel*, see [Castellan, 2015]. As in the alternating case, these behaviours must be banned via an adequate condition. By analogy we still call it *well-bracketing*, though this is a bit of a misnomer: for instance,



is a perfectly acceptable play for the constant **tt**, though the third move does not answer the pending question. Intuitively, this play is an interleaving of two evaluations of **tt** progressing independently. While each, taken independently, is indeed well-bracketed in the sense of alternating game semantics, this is blurred out in their interleaving – the notion of “pending question” does not make sense anymore.

Thus we must constrain the interplay between questions and answers differently:

Definition 5.1.2. Consider A an arena, and s a non-alternating pre-play.

¹For modeling IA , both choices are possible, but our discussion on the semantic cube in Section 3.3.3 requires non well-bracketed plays as doing otherwise is incompatible with the interpretation of **callcc**.




Figure 5.3: Logical well-bracketing is weaker than well-bracketing

We say that s is **logically well-bracketed** iff it satisfies the conditions:

- fork: for $t \quad q \quad t' \quad m \sqsubseteq s$, q is unanswered in t' .
 join: for $t \quad q \quad t' \quad a^A \sqsubseteq s$, all questions justified by q are answered in t' .

By *answer-closing* of Definition 3.2.4, moves labelled q are necessarily questions in this definition. Rather than *chronological*, the constraint is *logical*; it is about the hierarchical relationship between functions and arguments. *Fork* means that if a call returns, it can neither return again nor investigate its arguments anymore; *join* means that a call cannot return if there are un-terminated argument calls currently running.

For $s \in \Downarrow\text{-Plays}(A)$ alternating, if s is well-bracketed (in the sense of Definition 3.2.5) then it is logically well-bracketed (in the sense of Definition 5.1.2. However, the converse does not hold as illustrated by Figure 5.3. We may then define:

Definition 5.1.3. A **non-alternating play** on A is a (logically) well-bracketed non-alternating pre-play on A . We write $\Downarrow\text{-Plays}(A)$ for alternating legal plays on A .

Recall that the set of alternating legal plays on A was written $\Downarrow\text{-Plays}(A)$, where \Downarrow evokes the transition between states O and P , depending on whose turn it is. Here, \Downarrow evokes the fact that there is only one state, where both players can potentially play.

It is a puzzling fact that logical well-bracketing may replace chronological well-bracketing in alternating strategies also, and still get intensional full abstraction for IA (but not definability), despite the example in Figure 5.3. This boils down to the fact that logically well-bracketed *complete* plays (*i.e.* P - and O -visible, and all questions have an answer) are automatically well-bracketed – so *e.g.* Figure 5.3 cannot be completed.

5.1.2 Non-Alternating Strategies

Now, for strategies, which conditions of Definition 3.1.5 survive? *Non-empty*, *prefix-closed* and *receptive* still make sense. However *deterministic* is unreasonable, firstly because of the non-determinism introduced by the scheduler's choices, and secondly because a non-deterministic choice may be defined in IA_{\parallel} (see **choice** in Section 2.2.3).



Figure 5.4: Two consecutive P -moves

Figure 5.5: The other order




Figure 5.6: Two exchangeable moves




Figure 5.7: The other order

As a first approximation of non-alternating strategies, we set:

Definition 5.1.4. Consider A an arena.

A *non-alternating prestrategy* $\sigma : A$ is a set $\sigma \subseteq \mathcal{O}\text{-Plays}(A)$ satisfying:

- non-empty: $\epsilon \in \sigma$,
- prefix-closed: $\forall s \sqsubseteq t \in \sigma, s \in \sigma$.

Though *receptive* still makes sense, it is more useful technically to let prestrategies not be receptive, and impose receptivity only later on in Definition 5.1.5.

Saturation. Perhaps less obvious is the fact that we need a new condition.

We already mentioned that in a non-alternating play, part of the order in which moves appear is irrelevant. Imagine a program features consecutive Player moves:

$$\dots m_1^+ m_2^+ \dots \in \mathcal{O}\text{-Plays}(A).$$

Since the play is non-alternating, according to our operational intuition this means that two threads are currently under execution. So we are as in Figure 5.4, which illustrates that: (1) Player has no control over the order in which the two positive moves appear, which is the scheduler's prerogative; and (2) if the positive moves are exchanged, we reach nevertheless the same program state. So the order of two contiguous Player moves is irrelevant, they can be exchanged without affecting the future of the execution.

A symmetric situation is that of two consecutive Opponent moves, in which case, likewise their order is irrelevant and they can be exchanged without affecting the future of execution. Finally, consider a program in IA_{\parallel} with a play of the form

$$\dots m_1^+ m_2^- \dots \in \mathcal{O}\text{-Plays}(A)$$

where m_2 is *not* justified by m_1 . In that case we are as in Figure 5.6 and the positive m_1 can always be postponed until after the negative m_2 , reaching the same program state, and thus not affecting the future computation. But finally, one cannot always postpone an Opponent move after a Player move: the Player move may actually depend on the Opponent move. Overall, the relevant information in $s \in \mathcal{O}\text{-Plays}(A)$ is only that which survives permutations as described above; and accordingly we shall require that non-alternating strategies should be *saturated* under those.

Integrating this condition, the definition of non-alternating strategies becomes:

Definition 5.1.5. *Consider A an arena.*

*A **non-alternating strategy** $\sigma : A$, or \mathcal{O} -strategy for short, is a prestrategy s.t.:*

$$\begin{aligned} \text{receptive: } & \forall s \in \sigma, sa^- \in \mathcal{O}\text{-Plays}(A) \Rightarrow sa^- \in \sigma, \\ \text{courteous: } & \text{for } sabt \in \sigma \text{ with } \text{pol}(a) = + \text{ or } \text{pol}(b) = -, \\ & \text{then if } sbat \in \mathcal{O}\text{-Plays}(A), sbat \in \sigma \text{ as well.} \end{aligned}$$

where it is understood that $sabt$ and $sbat$ have the same pointers.

The earliest appearance of such a saturation that we are aware of is by Selinger [Selinger, 1997]. It seemed to have first appeared in game semantics in [Laird, 2001b], and was called *saturation* in [Ghica and Murawski, 2008]. We refer to it as *courtesy*, a terminology imported from [Melliès and Mimram, 2007]: this is for compatibility with the terminology in the forthcoming concurrent games, where the same behaviour is imposed causally rather than by a saturation condition.

It is notable that non-alternating game semantics is obtained “only” by changing the nature of plays: the arenas, representing types, will remain the same.

5.1.3 Single-Threadedness

Are all non-alternating strategies sensible w.r.t. IA_{\parallel} ?

It is clear that it is not the case: the same phenomenon as in Section 3.3.1 occurs. As in the alternating case, we must impose a condition ensuring that the state is local, in that no information flows between distinct instances of the program. By analogy with the alternating case, we call this *single-threadedness*. But we must warn the reader that this by no means implies that our strategies can only have a single “thread” as in concurrent programming! *Single-threadedness* means that each initial move spawns an independent copy of this program, even if that program is multi-threaded.

But single-threadedness cannot be imported transparently from alternating strategies. Though the *current thread* (Definition 3.3.1) seemingly still makes sense, Player may

have control over several threads simultaneously, as in the play on **B** displayed below

$$\mathbf{q}^- \cdots \mathbf{q}^- \cdots \mathbf{tt}^+ \cdots \mathbf{tt}^+$$

which will belong to the strategy for the constant **tt**. The last move fails condition *well-threaded* of Definition 3.3.2 even though **tt** does behave independently on all copies.

So we have to approach the issue a bit differently. First, define:

Definition 5.1.6. A *non-alternating thread* on arena A is $t \in \mathcal{O}\text{-Plays}(A)$ with at most one initial move; write $\mathcal{O}\text{-Plays}_\bullet(A)$ for the set of non-alternating threads on A .

Single-threaded behaviour is defined first via single threads, which invites:

Definition 5.1.7. Consider A an arena.

A *thread-strategy* $\sigma : A$ is a set $\sigma \subseteq \mathcal{O}\text{-Plays}_\bullet(A)$ satisfying:

$$\text{receptive: } \forall s \in \sigma, sa^- \in \mathcal{O}\text{-Plays}_\bullet(A) \Rightarrow sa^- \in \sigma,$$

as well as non-empty, prefix-closed and courteous as in Definitions 5.1.4 and 5.1.5.

As in Definition 5.1.4, a *thread-prestrategy* is a prestrategy $\sigma \subseteq \mathcal{O}\text{-Plays}_\bullet(A)$.

Note that unlike in Definition 5.1.4, *receptive* must be restricted to threads.

Given a thread-strategy $\sigma : A$, a full non-alternating strategy may be defined by interleaving copies of σ . Given $s, t \in \mathcal{O}\text{-Plays}(A)$, $s \sqcup t \subseteq \mathcal{O}\text{-Plays}(A)$ denotes the set of **interleavings** of s and t . If $X, Y \subseteq \mathcal{O}\text{-Plays}(A)$, $X \sqcup Y$ is defined as the union of all $s \sqcup t$, for $s \in X$ and $t \in Y$. Finally, if $X \subseteq \mathcal{O}\text{-Plays}(A)$, we set

$$\begin{aligned} X^{(0)} &= \{\varepsilon\} \\ X^{(n+1)} &= X \sqcup X^{(n)} \end{aligned}$$

and $X^! \subseteq \mathcal{O}\text{-Plays}(A)$, the **iterated interleaving** of X , is $X^! = \bigcup_{n \in \mathbb{N}} X^{(n)}$. We have:

Definition 5.1.8. Consider A an arena, and $\sigma : A$ a non-alternating strategy on A .

We say that σ is *single-threaded* if there is $\tau : A$ a thread-strategy such that $\sigma = \tau^!$.

In that case τ is uniquely determined, as it must be the set of threads in σ . If $\sigma : A$ is a non-alternating strategy on A , write σ^\bullet for its set of threads – σ is single-threaded iff $\sigma = (\sigma^\bullet)^!$. It also directly follows from Definition 5.1.8 that we have:

Lemma 5.1.9. Consider $\sigma, \tau : A$ single-threaded non-alternating strategies.

Then, $\sigma = \tau$ if and only if $\sigma^\bullet = \tau^\bullet$.

5.1.4 Full Abstraction for \mathbf{IA}_\parallel

As we shall spell out in Section 5.2, arenas and single-threaded non-alternating strategies form a cartesian closed category $\mathcal{O}\text{-Strat}$ which supports the interpretation of \mathbf{IA}_\parallel . More precisely, types and contexts are interpreted as arenas exactly as in Chapter 4 (\mathbf{IA}_\parallel brings no new type), and a term $\Gamma \vdash M : A$ is interpreted as a morphism

$$\llbracket M \rrbracket \in \mathcal{O}\text{-Strat}(\llbracket \Gamma \rrbracket, \llbracket A \rrbracket).$$





Figure 5.8: A realizable non P-visible play

Figure 5.9: Example of definability of plays in $\mathbf{IA}_{//}$

Realizing plays. But do we have more than $\mathbf{IA}_{//}$? It is natural to expect that the answer is yes: in the alternating setting, *P-visibility* is required to ban higher-order references (see Section 3.3.2), and we have not imposed any such condition here.

But P-visibility is not adequate: strategies from $\mathbf{IA}_{//}$ will not satisfy it, even though we only have ground type references. Figure 5.8 illustrates why: several threads in parallel may cause executions which are undistinguishable, viewed from the outside, from a non-local control flow due to a higher-order reference. In Figure 5.8, one cannot observe that the occurrence of the last move after the next-to-last is entirely coincidental, in that the two moves really belong to completely independent threads.

This phenomenon is general: given a play s , it is fairly easy to produce a term that can realize s , though typically along with many other plays. We can do this via a purely parallel term that simply mimics the hierarchical dependencies imposed by justification pointers. For instance, for s on the left hand side of Figure 5.9, one obtains the term

$$\lambda x^{U_1} f^{U_2 \rightarrow U_3} y^{U_4}. (x \parallel f \text{ skip} \parallel y); \perp. \quad (5.1)$$

Definability. Now, can we produce a term that realizes *exactly* s , and not more? Not in the strictest sense, since strategies interpreting terms will always be sets of plays satisfying Definition 5.1.5. But there is always a *smallest* strategy containing s :

Lemma 5.1.10. *Consider $\sigma : A$ a thread-prestrategy.*

Then, there is a unique least (for inclusion) thread-strategy $\hat{\sigma} : A$ s.t. $\sigma \subseteq \hat{\sigma}$. Its plays are deduced from σ by applying receptive and courteous from Definition 5.1.5.

Regarding $s \in \mathcal{O}\text{-Plays}_s(A)$ as a thread-prestrategy implicitly by prefix-closure, we have $\hat{s} : A$ its **saturation**, a thread-strategy, and $\hat{s}^! : A$ a single-threaded strategy.

The saturation process generates many plays, which are however succinctly summarized via a causal dependency between moves: a move depends on another when they *cannot be permuted by courtesy*. This happens if one points to the other; or for a Player move appearing *after* an Opponent move in s . It is helpful conceptually to represent these constraints via a diagram on moves, and this is done on the rhs of Figure 5.9: we draw the dependency due to pointers as dotted lines (read from top to bottom), and the dependency from negative moves to positive moves as arrows \rightarrow . The resulting object is close to a concurrent strategy in the sense of Part II – by design, \hat{s} comprises exactly the plays that are linearizations of this diagram (with pointers following dotted lines).

Now, we observe that in our example, the purely parallel term of (5.1) captures *exactly* the static/dotted dependencies on the rhs of Figure 5.9. To capture \hat{s} , we must also account for the arrows \rightarrow , and this is done by exploiting the memory. For that we define two helper functions. If $M : \mathbb{V}$, we write $\mathbf{set}(M) : \mathbb{U}$ for $M := 1$, and $\mathbf{test}(M) : \mathbb{U}$ for $\mathbf{if}(\mathbf{iszero} \ !M) \perp \mathbf{skip}$ which converges iff $\!M$ is non-zero. Then:

$$\lambda x^{\mathbb{U}_1} f^{\mathbb{U}_2 \rightarrow \mathbb{U}_3} y^{\mathbb{U}_4}. \left(\begin{pmatrix} x; \\ \mathbf{set}(\checkmark_1^-) \end{pmatrix} \parallel \begin{pmatrix} f(\mathbf{set}(\mathbf{q}_2^-); \\ \mathbf{test}(\checkmark_4^-); \\ \mathbf{grab}(\checkmark_2^+); \\ \mathbf{skip}) \end{pmatrix} \parallel \begin{pmatrix} \mathbf{test}(\checkmark_1^-); \\ \mathbf{test}(\mathbf{q}_2^-); \\ y; \\ \mathbf{set}(\checkmark_4^-) \end{pmatrix} \right); \perp$$

borrowes the shape of (5.1), additionally signaling the \rightarrow -links through memory (we omit reference and semaphore creations wrapping the term). We use one fresh reference (initialized to 0) for each Opponent move, which gets set to 1 when the Opponent move occurs. Finally, we must 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 5.9. For that we add semaphores which are grabbed, but never released.

Done systematically for arbitrary plays [Ghica and Murawski, 2008], we get:

Proposition 5.1.11. *Consider A an arena, and $s \in \mathcal{O}\text{-Plays}_s(A)$.*

Then, there is a term $\vdash M : A$ of \mathbf{IA}_{\parallel} such that $\llbracket M \rrbracket = \hat{s}^!$.

From this, it is deduced in [Ghica and Murawski, 2008] that we have:

Theorem 5.1.12. *Arenas and single-threaded strategies form an intensionally fully abstract model for \mathbf{IA}_{\parallel} . Moreover, given $\Gamma \vdash M, N : A$, we have*

$$M \simeq N \quad \Leftrightarrow \quad \mathbf{comp}(\llbracket M \rrbracket) = \mathbf{comp}(\llbracket N \rrbracket),$$

where $\text{comp}(\sigma)$ is the set of **non-alternating complete plays**, i.e. non-alternating plays such that all questions have an answer.

5.2 The Ambient Cartesian Closed Category $\mathcal{U}\text{-Strat}$

As in the alternating case, we now focus on the compositional construction of non-alternating strategies from terms, starting with the categorical structure.

5.2.1 Constructing $\mathcal{U}\text{-Strat}$

As for alternating strategies, a **non-alternating strategy from A to B** is a non-alternating strategy on $A \Rightarrow B$. The composition of non-alternating strategies largely follows the same route as for alternating strategies in Section 4.1.2, with a few minor changes.

Interactions. Consider $\sigma : A \Rightarrow B$ and $\tau : B \Rightarrow C$ non-alternating strategies. As in Section 4.1.2, we start with an appropriate notion of *interactions*:

Definition 5.2.1. Consider A, B and C arenas.

A **non-alternating interaction** on A, B, C is a legal pointing sequence on $(A \Rightarrow B) \Rightarrow C$ satisfying the conditions fork and join of Definition 5.1.2.

We write $\mathcal{U}\text{-I}(A, B, C)$ the set of all non-alternating interactions on A, B, C .

It is useful to compare with Definition 4.1.6 in the alternating case: the conditions *outer-legal*, *left-legal* and *right-legal*, which amount to all restrictions being alternating, become unnecessary: all restrictions are automatically valid non-alternating plays.

We use the same terminology and conventions as in the alternating case (notably regarding the polarities of moves, see below Definition 4.1.6) for non-alternating interactions. However, the state diagram of Figure 4.2 disappears. Just as plays, non-alternating interactions have no well-defined state, and both players can play anywhere at any time provided the move they play is justified and compatible with well-bracketing.

Composition. We now define composition, first for *prestrategies*:

Definition 5.2.2. Consider $\sigma : A \Rightarrow B$ and $\tau : B \Rightarrow C$ non-alternating prestrategies.

An **interaction** between σ and τ is any $u \in \mathcal{U}\text{-I}(A, B, C)$ such that $u \upharpoonright A, B \in \sigma$ and $u \upharpoonright B, C \in \tau$. We write $\sigma \parallel \tau$ for the set of all interactions between σ and τ .

We may then define the **composition** $\tau \odot \sigma = \{u \upharpoonright A, C \mid u \in \sigma \parallel \tau\} : A \Rightarrow C$.

As in the alternating case, if $s \in \tau \odot \sigma$ then by definition there must be some $u \in \sigma \parallel \tau$ such that $s = u \upharpoonright A, C$, called a **witness** for s . However, unlike in the alternating case, the witness is no longer unique. Composition is associative:

Proposition 5.2.3. Consider $\sigma : A \Rightarrow B, \tau : B \Rightarrow C$ and $\nu : C \Rightarrow D$ prestrategies.

Then, $(\nu \odot \tau) \odot \sigma = \nu \odot (\tau \odot \sigma)$.

Proof. Same proof as in the alternating case: one defines the set $\text{IA}(A, B, C, D)$ of ternary interactions as comprising well-bracketed legal pointing sequence on $((A \Rightarrow$

$B) \Rightarrow C) \Rightarrow D$. Then, one proves a “zipping lemma” with almost the same statement as Lemma 4.1.10 – with only the uniqueness clause missing – and the proof follows transparently as in Proposition 4.1.11 (uniqueness of witness is not needed). \square

Note that associativity holds already for composition of prestrategies, without relying on *receptivity* or *courtesy*. However, we will need these conditions shortly.

We mention the following preservation proposition [Ghica and Murawski, 2008]:

Proposition 5.2.4. *Consider $\sigma : A \Rightarrow B$ and $\tau : B \Rightarrow C$ prestrategies. Then:*

if σ and τ are strategies, then so is $\tau \odot \sigma$,
if σ and τ are single-threaded, then so is $\tau \odot \sigma$.

Copypcat. Composition of non-alternating strategies is as direct as in the alternating case – perhaps simpler, as interactions are “wilder”, not having to follow the rigid structure imposed by alternation and the state diagram. But non-alternation also means that identities cannot be as simple as in Definition 4.1.12: copypcat must handle Opponent moves arriving in bulk, possibly before earlier Opponent moves are propagated.

Nevertheless, the non-alternating copypcat can be succinctly defined by:

Definition 5.2.5. *Consider A any arena, and \mathfrak{C}_A the set of threads of \mathfrak{C}_A .*

The non-alternating copypcat on A , written $\mathfrak{C}\text{-}\mathfrak{C}_A$, is simply defined as $\widehat{\mathfrak{C}_A}^!$.

So $\mathfrak{C}\text{-}\mathfrak{C}_A$ is the smallest non-alternating strategy including the alternating copypcat. Despite the overload of notations, we will often still refer to the non-alternating copypcat simply with \mathfrak{C}_A – hopefully, the context should always be sufficient to disambiguate. We shall reserve $\mathfrak{C}\text{-}\mathfrak{C}_A$ for when we need to be extra explicit. Likewise, we shall use $\Downarrow\text{-}\mathfrak{C}_A$ for the alternating copypcat, when we feel it is useful to emphasize that this is not $\mathfrak{C}\text{-}\mathfrak{C}_A$.

This definition is succinct but not so explicit, so we illustrate it with a play of the non-alternating copypcat in Figure 5.10. As long as Opponent “plays in turn”, copypcat remains alternating. But if Opponent plays several moves in a row – before the previous ones can be propagated – then copypcat maintains a buffer of moves to be forwarded and plays them (not necessarily in the same order as they arrived), attempting to restore the same state between the right hand side and the left hand side.

As required, the non-alternating copypcat is neutral for composition with respect to non-alternating strategies. In fact, we have the following proposition:

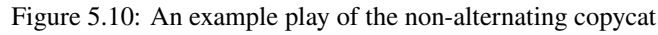
Proposition 5.2.6. *Consider A, B arenas, and $\sigma : A \Rightarrow B$ a prestrategy.*

Then, σ is a non-alternating strategy iff $\mathfrak{C}\text{-}\mathfrak{C}_B \odot \sigma = \sigma \odot \mathfrak{C}\text{-}\mathfrak{C}_A = \sigma$.

Proof. Direct adaptation of [Ghica and Murawski, 2008, Lemma 16]. \square

It follows that we have:

Corollary 5.2.7. *There is a category $\mathfrak{C}\text{-Strat}$ with arenas as objects, and as morphisms from A to B all single-threaded strategies on $A \Rightarrow B$.*



The cartesian closed structure essentially follows that of \Downarrow -**Strat**.

$$\circlearrowleft\text{-}\pi_A : A \times B \Rightarrow A, \quad \circlearrowleft\text{-}\pi_B : A \times B \Rightarrow B$$

To define the pairing, we prove:

$$\langle \sigma, \tau \rangle : \Gamma \Rightarrow A \times B,$$

Proof. Straightforward by definition of single-threaded non-alternating strategies. \square

Lemma 5.2.9. *Consider $\sigma : \Gamma \Rightarrow A$, $\tau : \Gamma \Rightarrow B$, and $\nu : \Delta \Rightarrow \Gamma$ single-threaded. Then, the following equations hold:*

$$\begin{aligned}\mathfrak{U}\text{-}\pi_A \odot \langle \sigma, \tau \rangle &= \sigma \\ \mathfrak{U}\text{-}\pi_B \odot \langle \sigma, \tau \rangle &= \tau \\ \langle \mathfrak{U}\text{-}\pi_A, \mathfrak{U}\text{-}\pi_B \rangle &= \mathfrak{U}\text{-}\mathfrak{C}_{A \times B} \\ \langle \sigma, \tau \rangle \odot \nu &= \langle \sigma \odot \nu, \tau \odot \nu \rangle.\end{aligned}$$

Proof. The first two equalities are a variation of Proposition 5.2.6, first proved for threads and extended to strategies via Lemma 5.1.9. The last two are direct on threads and generalized by Lemma 5.1.9. \square

It follows from these equations that we have the desired structure:

Corollary 5.2.10. *The category $\mathcal{O}\text{-Strat}$ is cartesian.*

Cartesian closed structure. Now, we describe the cartesian closure.

On arenas, $A \Rightarrow B$ is of course the same construction as for alternating strategies. We have seen in Lemma 4.1.18 that there is an isomorphism of arenas

$$\Xi_{\Gamma, A, B} : (\Gamma \times A) \Rightarrow B \cong \Gamma \Rightarrow (A \Rightarrow B)$$

which, as in the alternating case, immediately yields a bijection:

Lemma 5.2.11. *Consider Γ, A and B three arenas. Then, there is a bijection*

$$\Lambda_{\Gamma, A, B} : \mathcal{O}\text{-Strat}(\Gamma \times A, B) \simeq \mathcal{O}\text{-Strat}(\Gamma, A \Rightarrow B)$$

preserving composition in the sense that

$$\Lambda_{\Gamma, A, B}(\sigma) \odot \gamma = \Lambda_{\Delta, A, B}(\sigma \odot (\gamma \times A))$$

for all $\sigma \in \mathcal{O}\text{-Strat}(\Gamma \times A, B)$ and $\gamma \in \mathcal{O}\text{-Strat}(\Delta, \Gamma)$.

Proof. A routine manipulation on interactions. \square

In turn, this bijection lets us define the evaluation strategy simply as

$$\mathcal{O}\text{-ev}_{A, B} = \Lambda_{A, A, B}^{-1}(\mathcal{O}\text{-}\mathfrak{c}_{A \Rightarrow B}) \in \mathcal{O}\text{-Strat}((A \Rightarrow B) \times A, B)$$

and from the above, it becomes routine to prove cartesian closure:

Proposition 5.2.12. *The category $\mathcal{O}\text{-Strat}$ is cartesian closed.*

Proof. As for Proposition 4.1.20. \square

5.2.3 Recursion

As for alternating strategies, non-alternating strategies are partially ordered by inclusion. This partial order is preserved by all operations, yielding:

Proposition 5.2.13. *For any two arenas A and B , $\mathcal{O}\text{-Strat}(A, B)$ is partially ordered by inclusion, yielding a pointed dcpo. All operations on strategies involved in the cartesian closed structure are continuous, making $\mathcal{O}\text{-Strat}$ enriched over pointed dcpos.*

Proof. Direct verifications. \square

As in the alternating case, we define

$$\begin{aligned} F & : \mathcal{O}\text{-}\mathbf{Strat}(A \Rightarrow A, A) \rightarrow \mathcal{O}\text{-}\mathbf{Strat}(A \Rightarrow A, A) \\ \sigma & \mapsto \lambda f^{A \Rightarrow A}. f(\sigma f) \end{aligned}$$

a continuous operation whose least fixed point

$$\mathcal{O}\text{-}\mathcal{Y}_A = \bigcup_{n \in \mathbb{N}} F^n(\perp) \in \mathcal{O}\text{-}\mathbf{Strat}(A \Rightarrow A, A) \quad (5.2)$$

satisfies $\mathcal{O}\text{-}\mathcal{Y}_A = F \mathcal{O}\text{-}\mathcal{Y}_A$ by construction and provides the interpretation of recursion.

5.3 Interpretation of \mathbf{IA}_{\parallel}

5.3.1 Interpretation of \mathbf{PCF}_{\parallel}

First of all, *types* and *contexts* of \mathbf{PCF} are interpreted exactly as in $\Downarrow\text{-}\mathbf{Strat}$.

For *terms*, any $\Gamma \vdash M : A$ gets interpreted as a morphism

$$\llbracket M \rrbracket \in \mathcal{O}\text{-}\mathbf{Strat}(\llbracket \Gamma \rrbracket, \llbracket A \rrbracket).$$

As for $\Downarrow\text{-}\mathbf{Strat}$, the primitives for the simply-typed λ -calculus and for recursion are interpreted using the structure above. It remains to provide interpretations for all the additional combinators that \mathbf{PCF} adds to the simply-typed λ -calculus.

PCF. Constants are interpreted by single-threaded strategies returning directly the corresponding value – due to non-alternation, this means they have plays such as

$$\mathbf{q}^- \cdots \checkmark^+ \quad \mathbf{q}^- \cdots \mathbf{q}^- \cdots \mathbf{q}^- \cdots \checkmark^+ \quad \mathbf{q}^- \cdots \checkmark^+ \cdots \checkmark^+ \cdots \checkmark^+$$


for $\text{skip} : \mathbf{U}$, interleaving various calls in an asynchronous fashion – the generating thread-strategy is simpler to describe, with only one maximal play $\mathbf{q}^- \checkmark^+$.

The interpretation of \mathbf{PCF} is obtained via the same clauses as in alternating games, in Figure 4.5. This uses single-threaded non-alternating strategies seq , if , succ , iszero and pred whose thread-strategies are exhaustively presented as comprising prefixes of plays of the form displayed in Figures 4.4 and 4.6. Non-alternation does not induce additional plays, because we required that non-alternating plays should be well-bracketed thus Opponent can answer at most once. The strategy let comprises more plays than its alternating version in Figure 4.7. We omit however its definition, obtained as an obvious simplification of the new operation that \mathbf{PCF}_{\parallel} adds to \mathbf{PCF} : the parallel let .

\mathbf{PCF}_{\parallel} . To complete the interpretation of \mathbf{PCF}_{\parallel} , we need a single-threaded strategy

$$\text{plet}_{X,Y} : X \Rightarrow X \Rightarrow (X \Rightarrow X \Rightarrow Y) \Rightarrow Y$$

for the parallel let . Intuitively, its behaviour is simple enough: upon being called, plet will interrogate its first two arguments in parallel. Upon receiving two values, it will

Figure 5.11: Typical play of $\text{plet}_{X,Y}$.

interrogate its function argument and provide it the two values. We show one of its typical plays, in Figure 5.11. It acts in two subsequent stages: first the parallel evaluation of arguments, and then the call to the functional argument with the memoized values.

To define $\text{plet}_{X,Y}$ formally, we define sets of pointing sequences for these two stages. Annotating types as $X_1 \Rightarrow X_2 \Rightarrow (X \Rightarrow X \Rightarrow Y) \Rightarrow Y$, for v and w values of X , we set

$$\text{force}_{v,w} ::= \begin{array}{l} \mathbf{q}_1^+ v_1^- \mathbf{q}_2^+ w_2^- \mid \mathbf{q}_2^+ w_2^- \mathbf{q}_1^+ v_1^- \mid \mathbf{q}_1^+ \mathbf{q}_2^+ v_1^- w_2^- \\ \mid \mathbf{q}_1^+ \mathbf{q}_2^+ w_2^- v_1^- \mid \mathbf{q}_2^+ \mathbf{q}_1^+ v_1^- w_2^- \mid \mathbf{q}_2^+ \mathbf{q}_1^+ w_2^- v_1^- \end{array}$$

where it is understood that v_1^- points to \mathbf{q}_1^+ and w_2^- points to \mathbf{q}_2^+ .

Next, for the second stage we define another set of pointing sequences, written $\text{eval}_{v,w}$. Its elements are obtained by first considering the threads in the strategy

$$\lambda x^X. \lambda y^X. \lambda f^{X \Rightarrow X \Rightarrow Y}. f \ v \ w : X \Rightarrow X \Rightarrow (X \Rightarrow X \Rightarrow Y) \Rightarrow Y,$$

defined exploiting the cartesian closed structure of $\mathcal{O}\text{-Strat}$ and the corresponding internal language. From such threads, elements of $\text{eval}_{v,w}$ are obtained by removing the initial question. Finally, we form the set of pointing sequences

$$\text{prelet}_{X,Y} = \bigcup_{v,w} \mathbf{q}^- \cdot \text{force}_{v,w} \cdot \text{eval}_{v,w}$$

completed into a set of plays by reinstating the missing pointers in the unique possible way, with missing pointers assigned to the initial move \mathbf{q}^- . Finally, the non-alternating thread-strategy for $\text{plet}_{\mathbf{X}, \mathbf{Y}}$ is obtained as interleavings of all prefixes of $\text{prelet}_{\mathbf{X}, \mathbf{Y}}$.

Altogether, we obtain $\text{plet}_{\mathbf{X}, \mathbf{Y}} : \mathbf{X} \Rightarrow \mathbf{X} \Rightarrow (\mathbf{X} \Rightarrow \mathbf{X} \Rightarrow \mathbf{Y}) \Rightarrow \mathbf{Y}$ and we complete the interpretation of PCF_{\parallel} by adding to Figure 4.4 the additional clause

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

5.3.2 Interpretation of State

The types \mathbb{V} and \mathbb{S} are interpreted by the same arenas as in Section 4.2.2.

Queries. As in the alternating case, we set the interpretation of queries as

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

using the single-threaded non-alternating strategies with threads as in Figures 4.11, 4.12, 4.13, and 4.14 – the same basic plays as in the alternating case.

Declaration. Now, we must provide the interpretation for the declaration of new references and semaphores. For that, we need non-alternating versions of the strategies for reference and semaphore declaration introduced in Section 4.2.2.

Rather than going again through the concrete definition of these strategies, we set

$$\mathcal{O}\text{-newr}_{\mathbf{X}} = \widehat{\text{newr}_{\mathbf{X}}}^!, \quad \mathcal{O}\text{-news}_{\mathbf{X}} = \widehat{\text{news}_{\mathbf{X}}}^!$$

via the saturation process of Lemma 5.1.10, and use those in:

$$\begin{aligned} \llbracket \text{newref } x := n \text{ in } M \rrbracket &= \mathcal{O}\text{-newr}_{\mathbf{X}} \odot \Lambda(\llbracket M \rrbracket) \in \mathcal{O}\text{-Strat}(\llbracket \Gamma \rrbracket, \llbracket A \rrbracket) \\ \llbracket \text{newsem } x := n \text{ in } M \rrbracket &= \mathcal{O}\text{-news}_{\mathbf{X}} \odot \Lambda(\llbracket M \rrbracket) \in \mathcal{O}\text{-Strat}(\llbracket \Gamma \rrbracket, \llbracket A \rrbracket). \end{aligned}$$

Bad variables and semaphores. As in the alternating case, through the isomorphisms $\mathbf{V} \cong (\&_{n \in \mathbb{N}} \mathbf{U}) \& \mathbf{N}$ and $\mathbf{S} \cong \mathbf{U} \& \mathbf{U}$, the interpretation is:

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

which concludes the interpretation of IA_{\parallel} .

5.3.3 Summary of Results

This concludes the construction of the cartesian closed category $\mathcal{O}\text{-Strat}$, as well as the interpretation of \mathbf{IA}_{\parallel} . As in the alternating case, we conclude the technical content of this chapter by summing up a few key results for future reference.

Proposition 5.3.1. *There is an adequate interpretation of \mathbf{IA}_{\parallel} in $\mathcal{O}\text{-Strat}$.*

Proof. The interpretation is as described in its chapter. Computational adequacy corresponds to the particular case of Theorem 5.1.12 instantiated on ground type. \square

Theorem 5.3.2. *The interpretation of Proposition 5.3.1 refines into an interpretation of \mathbf{IA} in $\mathcal{O}\text{-Strat}$, which is intensionally fully abstract.*

Proof. First stated as Theorem 5.1.12. \square

5.4 Conclusions and Historical Notes

The first to consider a non-alternating version of play-based game semantics was Laird, who proposed a fully abstract model for a message-passing language called Idealized CSP [Laird, 2001b]. This model had much in common with the one presented here, however plays are equipped with additional “concurrency pointers”, which explicitly mark the different threads acting in a non-alternating plays. Ghica and Murawski’s model was first published in 2004 [Ghica and Murawski, 2004], in conference format. It is strongly related to the non-alternating games of Melliès and Mimram published soon after [Melliès and Mimram, 2007], which contains the definition of a category of games and linear non-alternating strategies, along with an adequate notion of innocence.

Just like in the alternating case in the presence of state, non-alternating games form an effectively presentable fully abstract model, which lead to decision algorithms for observational equivalence – though for very restricted languages: first bounding the number of threads via a type system [Ghica et al., 2006] so that strategies are representable as finite automata, and more recently via other finitary restrictions that allow a representation of non-alternating strategies via more elaborate notions of data automata [Dixon et al., 2021a, Dixon et al., 2021b].

It is noteworthy, in the definability process outlined in Proposition 5.1.11, that it seems we really *need* semaphores in order to capture a single play. Murawski explored what happened without semaphores [Murawski, 2010]; he proved that then strategies are closed under a stuttering behaviour that affect observational equivalence. This shows that, perhaps surprisingly, semaphores are *not* definable just via references up to observational equivalence – another of those very results which illustrates well the strength of game semantics.

This concludes the preliminary part. One of the purposes of the rest of this monograph will be to construct and explore $\rightarrow\text{-Strat}$, a *causal* version of $\mathcal{O}\text{-Strat}$.

Part II

Thin Concurrent Games

Introduction to Part II




Figure 5.12: To survive pointer showers, remember to bring your parallel-or umbrella!

Now that we have spent significant time building up a background on traditional play-based game semantics, let us forget it all and start from scratch!

Concurrent games reject the premise that there exists a global clock; instead relying on a *causal* presentation of program behaviour. This idea was pioneered by Abramsky and Melliès [Abramsky and Melliès, 1999], and its consequences were mainly explored by Melliès in his sequence of insightful papers on *asynchronous games* [Melliès, 2003, Melliès, 2004a, Melliès, 2004b, Melliès, 2005, Melliès and Mimram, 2007]. Inheriting from this, the definitions on which we build were first proposed by Rideau and Winskel [Rideau and Winskel, 2011], also building on other work [Faggian and Hyland, 2002, Curien and Faggian, 2005, Faggian and Piccolo, 2009] inspired by *ludics* – we redirect to Section 6.5 for a longer discussion on the history of these notions.

While thin concurrent games share much conceptually with the game semantics of Part I, they build on a completely different – and more sophisticated – technical basis: the theory of event structures. This enhances their expressiveness, but with the cost of an involved technical development, for two reasons.

Firstly, event structures are more elaborate than the usual objects used in game semantics, such as sequences of moves. Manipulating them rigorously is more complex than for plays, involving a demanding – but (perhaps subjectively) rather elegant – infrastructure involving significant mathematical engineering and ingenuity.

Secondly, we cannot get away with some of the usual informal practice of traditional game semantics. Indeed, researchers in game semantics have sometimes adopted a somewhat informal mathematical style – the strategy does this, if Opponent plays this then the strategy reacts by playing that, *etc.* While this has been criticized, it is true that such phrasings in the literature can usually be unambiguously made formal if needed (though in practice this is rarely done); but this is thanks to the fact that the mathematical objects involved are rather simple and closely follow their natural language description. So it is both the curse and the blessing of concurrent games that when strategies are event structures, such informal descriptions are no longer reasonable.

So, we must prepare the reader for a long mathematical development, in part because of the complexity of our strategies, because the development is carried out with more mathematical precision than is the standard in many game semantics texts, but also because we have taken the space to make the development as pedagogical as possible.

The purpose of Part II is the construction of so-called “thin concurrent games” along with the categorical structure needed to interpret the simply-typed λ -calculus, leaving for Part III the interpretation of $\mathbf{IA}_{//}$ and its fragments. We consider that the key contribution of Part II is the bicategory **TCG**, a setting for strategies-as-event-structures able to handle symmetry and the interchangeability of moves as required for copy indices and uniformity (in the sense of Section 3.4). We arrive there in two stages: first, in Chapter 6 we develop the bicategory **CG** of concurrent games *without symmetry*; then Chapter 7 introduces **TCG**, a conservative extension of **CG** equipping it to handle symmetry. Finally, **TCG** is only a bicategory and as such does not yet have the further categorical structure required to handle the interpretation of programming language; for that, in Chapter 8 we construct **NTCG**, essentially a model of intuitionistic linear logic with an exponential ! so that the Kleisli category **NTCG**_! is cartesian closed.

Chapter 6

Basic Concurrent Games

In this chapter, relying on the intuitions introduced previously for traditional game semantics, we introduce the bicategory **CG**, the foundation for concurrent games on event structures. We do this for pedagogical reasons: **CG** is only a stepping stone, as it does not yet support the exponential (!) – this will require the addition of *symmetry* in Chapter 7. But it seems better to first construct **CG** separately, so as to give the proper spotlight to quite a few new conceptual ideas, instead of introducing everything simultaneously.

CG is a non-deterministic extension of Melliès and Mimram’s *non-alternating asynchronous games* [Melliès and Mimram, 2007]; it was first introduced by Rideau and Winskel in [Rideau and Winskel, 2011] (see also [Castellan et al., 2017a]). The presentation given here is new, exploiting insights obtained over the years, from our experience of working with concurrent games.

6.1 Games and Strategies as Event Structures


6.1.1 Basic Intuitions on Concurrent Strategies

Parallelism. In the beginning of Section 5.1.1, we motivated non-alternating strategies by considering the plays in a parallel implementation of the arithmetic sum

$$x : \mathbb{N}, y : \mathbb{N} \vdash \text{let } \begin{pmatrix} x_1 & = & x \\ x_2 & = & y \end{pmatrix} \text{ in } x_1 + x_2 : \mathbb{N},$$

which we argued required to account for all the orders in which the evaluations of x and y could be performed by the scheduler. Rather than writing down all these interleavings,

an alternative could be to write – naively for now – a diagram of the form:




the Hasse diagram of a *partial order* in which the two sub-computations occur in independent branches. As for sequential plays in game semantics, this diagram is read from top to bottom: first, Opponent initiates computation with q^- . But then, there are *two* follow-ups: the two occurrences of q^+ , prompting the evaluation of the two arguments. Despite them being in the same row, they are not thought of as simultaneous. Instead they are *independent*: there is no well-defined ordering between them. Concurrent strategies display explicitly this independence, whereas the non-alternating game semantics of Chapter 5 would convey it implicitly by allowing all interleavings.

In a diagram as above, the vertical axis no longer represents time. Instead, the partial order is thought of as expressing *causality*: in order to observe the two occurrences of q^+ , we *must* have seen q^- . In order to observe (this occurrence of) the answer 42, we *must* have seen the two arguments converge with values 12 and 30.

Non-determinism. In such diagrams, the branching structure indicates parallelism, not non-determinism. But then, how is non-determinism going to be represented?


Our answer is to mark branchings that correspond to non-deterministic choice:



In this diagram, which represents the non-deterministic choice of a boolean, tt^+ and ff^+ are linked with a wiggly line indicating *conflict*: only one of them can be seen in a given execution. This marks that the branching is a *non-deterministic* rather than a *parallel* branching. Such situations with these two branching structures are captured by the mathematical notion of an *event structure*, introduced formally in Section 6.1.2.


Concurrent strategies explicitly display all non-deterministic choices, even when

they do not yet yield visibly different behaviours. For instance, in the diagram



corresponding *e.g.* to the term $\lambda f^{\mathbb{B} \rightarrow \mathbb{B}}. (f \mathbf{tt}) \odot (f \mathbf{ff})$ for \odot a non-deterministic choice operator, the two conflicting occurrences of \mathbf{q}^+ correspond to the *same* observable behaviour, but they should be kept separate because they have distinct futures. In general, concurrent strategies remember all non-deterministic choices even when they lead to no observable difference: $\mathbf{tt} \odot \mathbf{tt}$ will *not* yield the same concurrent strategy as \mathbf{tt} .

Causality. Though we have motivated the use of partial order models for representing strategies by examples from parallel and non-deterministic computation, we shall see throughout this monograph that it has a deep impact even for pure sequential deterministic computation. As an example, the term $\lambda f. f \mathbf{tt} \mathbf{ff}$ admits the causal behaviour



where it is apparent that the two branches where Opponent explores the two arguments are *causally independent* from each other. In fact, in that case the tree structure coincides with the tree of *P-views* as described in Section 3.2.4, exposing its causal nature.

6.1.2 Formalizing Concurrent Strategies

Such diagrams as above, mixing causal dependency and conflict, are certain *event structures* – the partial order model at the core of concurrent games.

Event structures. We first define *event structures*¹:

Definition 6.1.1. 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

¹In the literature, those event structures are called *prime event structures with binary conflict*.

irreflexive symmetric binary relation on $|E|$ called **conflict**, satisfying:

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

A given event has only finitely many causes: this fundamental axiom means that any event can appear in a finite execution. *Conflict inheritance* is sometimes called the *vendetta axiom*: conflicts transfer to children, and to their children, and so on for ever. It ensures that two non-conflicting events can always appear together in a finite execution (as the causal dependencies of non-conflict events must be non-conflicting).

The diagrams appearing earlier in this section may be read as event structures. For E an event structure, the **immediate causal dependency** \rightarrow_E is defined as $e \rightarrow_E e'$ if $e <_E e'$ with no event in between: for all $e'' \in E$, $e \leq_E e'' \leq_E e'$ we have $e = e''$ or $e'' = e'$. Likewise, the **immediate conflict** $e \rightsquigarrow_E e'$ means that $e \#_E e'$ and the conflict is not inherited: if $e'' <_E e$, $\neg(e'' \#_E e')$ and if $e'' <_E e'$, $\neg(e \#_E e'')$. These two notations are intensively used, in particular when drawing event structures. For instance, in the diagram (6.1), \mathbf{tt}^+ and \mathbf{ff}^+ are in conflict (as they depend on conflicting events), but not in immediate conflict – so they are not linked with a wiggly line.

The most fundamental definition on an event structure is *configurations* – one may regard an event structure as a concrete representation of its set of configurations:

Definition 6.1.2. A (finite) **configuration** of event structure E is a finite $x \subseteq |E|$ s.t.:

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

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

Configurations are the *states* of an event structure: conflict-free finite sets of events with all causal dependencies satisfied. Though a configuration is merely a set, it inherits a partial order which is often drawn in pictures from that of the ambient event structure.

The set $\mathcal{C}(E)$ is naturally ordered by inclusion; it is the *domain of configurations*. Configurations are typically ranged over by variables x, y, z . For $x, y \in \mathcal{C}(E)$, we write $x \subset y$ if x is immediately below y in the inclusion order, i.e. there is $e \in E$ such that $e \notin x$ and $y = x \cup \{e\}$ – we also write $x \vdash_E e$ and say that x **enables** e .

Games. Before we formalize strategies, let us look at *games*.

As for arenas, our *games* present observable computational events, along with their dependencies and conflict. In spirit, the way this is done is very close to the arenas of Chapter 3. For instance, the games representing respectively \mathbb{B} and \mathbb{U} are:

$$\mathbb{B} = \begin{array}{c} \mathbf{q}^- \\ \swarrow \quad \searrow \\ \mathbf{tt}^+ \rightsquigarrow \mathbf{ff}^+ \end{array} \qquad \mathbb{U} = \begin{array}{c} \mathbf{q}^- \\ \vdots \\ \checkmark^+ \end{array} \quad (6.2)$$

where the first is to be compared with the arena for booleans in traditional arena games in Figure 3.8. It is the same, except that the two answers are conflicting: unlike in

traditional game semantics, concurrent games are inherently linear (or rather, affine) and a question may be by default answered at most once. This hints at a deeper difference: in Figure 3.8, the initial question may appear arbitrarily many times in a play. In contrast, the intention here is that a given move can be played at most once: an execution such as at the beginning of Section 3.3.1 would make sense not on B but on its “bang”

$$!B = \begin{array}{c} \mathbf{q}_0^- \quad \mathbf{q}_1^- \quad \mathbf{q}_2^- \quad \dots \\ \vdots \quad \vdots \quad \vdots \quad \vdots \\ \mathbf{tt}_0^+ \sim \mathbf{ff}_0^+ \quad \mathbf{tt}_1^+ \sim \mathbf{ff}_1^+ \quad \mathbf{tt}_2^+ \sim \mathbf{ff}_2^+ \quad \dots \end{array} \quad (6.3)$$

involving *copy indices* as in AJM games (see Section 3.4) – but we push this operation aside until Chapter 7, as developing it properly also requires a model in which one can express that all copies are interchangeable, which plain event structures cannot do.

Games feature causal dependency and conflict, hence are *also* event structures:

Definition 6.1.3. A *game* is an *es* $A = (|A|, \leq_A, \#_A)$ with a *polarity function*

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

The polarity function simply indicates whether a given move is due to Player (+) or Opponent (−). Though games are event structures, we draw them with specific conventions: in particular, as in (6.2) and following our earlier convention for arenas, immediate causal dependency in games is drawn with dotted lines rather than with the usual symbol \rightarrow . Having distinct notations for immediate causal dependency in games and strategies will allow us later on to have diagrams with two immediate dependencies superimposed, from a game and from a strategy.

In practice, games in the sense of Definition 6.1.3 are way too wild for our purposes. The games arising from the interpretation of types will have a very specific shape: they are negative (Opponent starts), alternating (immediate causal links are always between different players), forestial (there is no causal merge) and conflict is local (between two moves with the same predecessor, if any). But for the sake of economy, we shall impose such conditions only when they become technically required.

Prestrategies. If both strategies and games are to be event structures, what does it mean that an event structure “plays” on another event structure? It is natural to model such a situation as an event structure (the strategy) *labeled* by an event structure (the game), where the labelling function should be a *map of event structures*:

Definition 6.1.4. Consider E, F two *es*. A *map of event structures* $f : E \rightarrow F$ is

$$f : |E| \rightarrow |F|$$

a function on events satisfying the further two conditions below:

- configuration-preserving: for all $x \in \mathcal{C}(E)$, $fx \in \mathcal{C}(F)$,
- locally injective: for all $e_1, e_2 \in x \in \mathcal{C}(E)$, if $fe_1 = fe_2$ then $e_1 = e_2$.

This is a kind of simulation map. It transports valid states into valid states while preserving atomicity: within a configuration $x \in \mathcal{C}(E)$, *local injectivity* ensures that f induces a bijection $f : x \simeq fx$; adding an event to x must add an event to fx . We shall review a few basic properties of maps of event structures in Section 6.1.3, but for now let us introduce a first approximation of strategies-as-event-structures:

Definition 6.1.5. A *prestrategy* on game A is an event structure $\sigma = (|\sigma|, \leq_\sigma, \#_\sigma)$ with


$$\partial_\sigma : \sigma \rightarrow A$$

a map of event structures.

Observe that here σ is a plain event structures, not assumed equipped with a polarity function. Nevertheless, there is an induced notion of polarity on σ , obtained as $\text{pol}_\sigma(s) = \text{pol}_A(\partial_\sigma(s))$ – we shall use this implicitly from now on.

In contrast with Definition 3.1.2, a prestrategy is not a collection of standalone executions (plays), but one *global* object, an event structure σ , aggregating all possible executions. The events in $|\sigma|$ are not events of the game: they are internal to the prestrategy, only linked to the game via the *display map* ∂_σ .

For instance, we may have a prestrategy (in fact, a strategy) on U :



with the obvious mapping $\partial_\sigma(\bullet_1) = q^-$ and $\partial_\sigma(\bullet_2) = \partial_\sigma(\bullet_3) = \checkmark^+$. We generally prefer drawing (pre)strategies as on the right hand side, which retains all the relevant information, only omitting the *names* of events in σ (i.e. $\bullet_1, \bullet_2, \bullet_3$)².

Strategies. The definition of prestrategies ignores polarities, thus not all prestrategies are computationally sensible – we need conditions analogous to Definition 5.1.5:

Definition 6.1.6. A prestrategy σ on game A is a *strategy* if it satisfies:

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

We write $\sigma : A$ to mean that σ is a strategy on game A .

Receptive forces strategies to acknowledge uniquely any possible move by Opponent. Finally, *courtesy* gets two birds with one stone. Firstly, it forbids causal links of the form $s_1 \rightarrow_\sigma s_2^-$ which are not in the game – a strategy cannot force Opponent to wait before playing a move available in the game. Secondly, it also forbids causal links of the form $s_1^+ \rightarrow_\sigma s_2^+$ not in the game. This is to be understood as an *asynchrony* requirement:

²This diagram exploits our convention to write immediate causal links differently in games and strategies.

while a program may play two positive moves in a row, it cannot control in which order those will reach the Opponent; the moves are thought of as propagating independently in an asynchronous medium (say, the internet), so that their order may change.

This completes our first notion of *concurrent strategy*.

Mediating between strategies. Hopefully, it is apparent to the reader why strategies are formalized as certain maps $\partial_\sigma : \sigma \rightarrow A$, with σ built on a set of events separated from A . However, this also means that equations between strategies should not be expected to hold up to *equality* – otherwise the two maps of event structures

$$\begin{array}{ccccc}
 \sigma & & U & & \tau \\
 \begin{array}{c} \bullet 1 \\ \swarrow \quad \searrow \\ \bullet 2 \quad \bullet 3 \\ \text{---} \end{array} & \xrightarrow{\partial_\sigma} & \begin{array}{c} q^- \\ \vdots \\ \checkmark + \end{array} & \xleftarrow{\partial_\tau} & \begin{array}{c} \bullet a \\ \swarrow \quad \searrow \\ \bullet b \quad \bullet c \\ \text{---} \end{array}
 \end{array} \quad (6.4)$$

would denote different strategies, though they only differ via the *name* they assign events. To account for that, we shall compare strategies via the following *morphisms*:

Definition 6.1.7. Consider σ and τ two (pre)strategies on a game A .

A *morphism* from σ to τ , written $f : \sigma \Rightarrow \tau$, is a map of es $f : \sigma \rightarrow \tau$ satisfying:

- compatibility with display maps: $\partial_\tau \circ f = \partial_\sigma$,
- rigidity: for all $s \leq_\sigma t$, we have $f s \leq_\tau f t$.

An *isomorphism of strategies* $f : \sigma \cong \tau$ is simply an invertible morphism. We sometimes write $\sigma \cong \tau$ for the mere existence of an isomorphism between σ and τ .

We assume rigidity here as it simplifies some proofs, notably later on in the presence of symmetry. Our next goal is to organize concurrent strategies into a bicategory **CG**; but for that we will need a bit of the basic theory on event structures.

6.1.3 Basic Properties of Event Structures and Maps

First, event structures and their maps (as in Definition 6.1.4) form a category **ES**. For E an event structure, we often write $e \in E$ as an alias for $e \in |E|$. If $e \in E$,

$$[e]_E = \{e' \in E \mid e' \leq_E e\}$$

is the set of **causal dependencies** of e . It is an easy exercise to prove that $[e]_E \in \mathcal{C}(E)$ – we shall use that fact without mentioning it. We say that $x \in \mathcal{C}(E)$ is a **prime configuration** iff it has a top element, or equivalently if it is $x = [e]_E$ for some $e \in E$. It is clear that events of E are in one-to-one correspondence with prime configurations.

Compatibility with the order. Event structures are posets, so one may expect their maps to preserve the order. But this is not at all the intention, maps of event structures are not meant to be monotone. The following is a perfectly acceptable map:

$$1 \longrightarrow 2 \quad \rightarrow \quad 1 \quad 2$$

sending two events in causal dependency to causally independent moves. Thus causality is not preserved, but we do have a property in the other direction:

Lemma 6.1.8. *Consider $f : E \rightarrow F$ a map of event structures.*

Then, for all $e, e' \in x \in \mathcal{C}(E)$, if $fe \leq_F fe'$ then $e \leq_E e'$.

Proof. Since $e' \in E$, we have $[e']_E \in \mathcal{C}(E)$, hence $f[e']_E \in \mathcal{C}(F)$ since f preserves configurations. But by construction, $fe' \in f[e']_E$. Since $f[e']_E \in \mathcal{C}(F)$, it is in particular down-closed. Hence, since $fe \leq_F fe'$, we must have $fe \in f[e']_E$. But this means that there is some $e'' \in [e']_E$ such that $fe'' = fe$ – in other words, there is $e'' \leq_E e'$ such that $fe'' = fe$. Now, we remark that we have $e' \in x$, so as x is down-closed, we must have $e'' \in x$ as well. Thus $e, e'' \in x$ have the same image by f , thus $e = e''$ by local injectivity. Hence, we have proved that $e \leq_E e'$ as required. \square

Recall that we have defined prestrategies as maps $\partial_\sigma : \sigma \rightarrow A$. For $\sigma : A$, the lemma above means that if $s, s' \in x \in \mathcal{C}(\sigma)$, if $\partial_\sigma s \leq_A \partial_\sigma s'$ then $s \leq_\sigma s'$ – in other words, any “static” causal dependency imposed by the game must be respected by the strategy as well. But causality is reflected only within configurations: the map

$$\begin{array}{ccc} 1 & \sim & 1' \\ \downarrow & & \downarrow \\ 2 & & 2 \end{array} \quad \rightarrow \quad \begin{array}{c} 1 \\ \downarrow \\ 2 \end{array}$$

is valid, though we have $f1' \leq f2$ but $\neg(1' \leq 2)$ as they are in conflict.

The maps of event structures that are monotone are called **rigid**. In that case:

Lemma 6.1.9. *Consider $f : E \rightarrow F$ a rigid map of event structures.*

Then, f preserves immediate causality.

Proof. Left as an exercise to the reader. \square

Rigid maps $f : E \rightarrow F$ are too constrained to be interesting as strategies, but they are useful as a sort of inclusions of E within F , collapsing non-deterministic branches. The *morphisms between strategies*, as defined in Definition 6.1.7, are rigid maps.

In Definition 6.1.7 we encountered for the first time the *isomorphisms* in the category of event structures, as induced by Definition 6.1.4. It is worth pointing out explicitly that such isomorphisms of event structures (and thus, isomorphisms of strategies) are simply *renamings* of the events, preserving and reflecting all structure:

Lemma 6.1.10. *Consider E, F event structures, $f : E \rightarrow F$ an isomorphism of es.*

Then, f is a bijection between events, satisfying the following two properties:

- (1) *for all $e, e' \in E$, $e \leq_E e'$ iff $f(e) \leq_F f(e')$,*
- (2) *for all $e, e' \in E$, $e \#_E e'$ iff $f(e) \#_F f(e')$.*

Proof. Left as an exercise to the reader. \square

Reciprocally, bijections satisfying properties (1) and (2) above are obviously isomorphisms of event structures – clearly, the two strategies of (6.4) are isomorphic.

Mapification. The above shows that a map of es $f : E \rightarrow F$ should not be viewed as a map of posets; instead it should be seen as a concrete presentation of

$$f : \mathcal{C}(E) \rightarrow \mathcal{C}(F)$$

a monotone function between the generated domains of configurations (the sets of configurations, ordered by inclusion). In fact, we will often reason on event structures via their domains of configurations; likewise we shall build maps of event structures via their action on configurations – we call the lemma below the “mapification” lemma, letting us build maps of event structures from functions on configurations:

Lemma 6.1.11. *Consider $f : \mathcal{C}(E) \rightarrow \mathcal{C}(F)$ preserving unions and cardinality.*

Then, there is a unique map of es $\hat{f} : E \rightarrow F$ s.t. for all $x \in \mathcal{C}(E)$, $\hat{f}(x) = f(x)$.

Proof. Existence. As f preserves unions it is monotone. As it preserves cardinality, it preserves \emptyset and $\neg\subset$. The key point is that it preserves *covering squares*:

$$\begin{array}{ccc} x & \xrightarrow{a} & y \\ b \downarrow & & \downarrow b \\ u & \xrightarrow{a} & v \end{array}$$

for $a \neq b$. Indeed as f preserves $\neg\subset$, $f(x) \neg\subset f(y)$ and $f(x) \neg\subset f(u)$. But also $f(v) = f(y \cup u) = f(y) \cup f(u)$ since f preserves unions, thus $f(v) = f(x) \cup \{a', b'\}$. Hence,

$$\begin{array}{ccc} f(x) & \xrightarrow{a'} & f(y) \\ b' \downarrow & & \downarrow b' \\ f(u) & \xrightarrow{a'} & f(v) \end{array}$$

Now, we are in position to define \hat{f} . Given $e \in E$, write

$$[e]_E = \{e' \in E \mid e' <_E e\}$$

its set of **strict causal dependencies**, which is automatically a configuration. Of course,

$$[e]_E \xrightarrow{e} [e]_E$$

so that $f[e]_E \xrightarrow{e'} f[e]_E$ for some $e' \in F$ – set $\hat{f}(e) = e'$. Now, for any $x \xrightarrow{e} y$, we have

$$\begin{array}{ccccccccccc} [e]_E & \xrightarrow{e_1} & x_1 & \xrightarrow{e_2} & x_2 & \xrightarrow{e_3} & x_3 & \dots & x_{n-1} & \xrightarrow{e_n} & x \\ e \downarrow & & \downarrow & & \downarrow & & \downarrow & & \downarrow & & \downarrow e \\ [e]_E & \xrightarrow{e_1} & y_1 & \xrightarrow{e_2} & y_2 & \xrightarrow{e_3} & y_3 & \dots & y_{n-1} & \xrightarrow{e_n} & y \end{array}$$

a tiling of covering squares, which are all preserved by f , implying $f x \hat{\sim}^{\hat{f}(e)} f y$. But now, for any $x \in \mathcal{C}(E)$, one may reach it in $\mathcal{C}(E)$ adding events one by one, *i.e.* there is

$$\emptyset = x_0 \xrightarrow{e_1} x_1 \xrightarrow{e_2} \dots x_{n-1} \xrightarrow{e_n} x_n = x$$

a **covering chain** in E obtained as any linearization of x according to \leq_E . But then

$$\emptyset = f(x_0) \xrightarrow{\hat{f}(e_1)} f(x_1) \xrightarrow{\hat{f}(e_2)} \dots f(x_{n-1}) \xrightarrow{\hat{f}(e_n)} f(x_n) = f(x)$$

from the observation above. But this means that $f(x) = \hat{f}(x)$, the latter defined by the direct image. It immediately follows that \hat{f} is a map of event structures: as it coincides with f of configuration it must preserve configurations, and a failure of local injectivity would immediately contradict the fact that f preserves cardinality.

Uniqueness. For $h : E \rightarrow F$ s.t. $h(x) = f(x)$ for all $x \in \mathcal{C}(E)$, then if $e \in E$,

$$f([e]_E) \xrightarrow{h(e)} f([e]_F)$$

so that $h(e) = \hat{f}(e)$ as required, concluding the proof. \square

Of course, reciprocally, if $f : E \rightarrow F$ is a map of event structures then $f : \mathcal{C}(E) \rightarrow \mathcal{C}(F)$ must preserve cardinality and unions as it is defined as a direct image. It follows that to test equality of maps of es, it suffices to evaluate them on configurations:

Lemma 6.1.12. *Consider $f, g : E \rightarrow F$ maps of event structures.*

If for all $x \in \mathcal{C}(E)$ we have $f(x) = g(x)$, then $f = g$.

Proof. Obvious by the *uniqueness* clause of Lemma 6.1.11. \square

Another particularly useful consequence of the mapification lemma is:

Proposition 6.1.13. *Consider E, F two es, and $f : \mathcal{C}(E) \cong \mathcal{C}(F)$ an order-iso.*

*Then, there is a unique $\hat{f} : E \cong F$ an iso in **ES** s.t. for all $x \in \mathcal{C}(E)$, $\hat{f}(x) = f(x)$.*

Proof. Firstly, f preserves unions as it is an order-isomorphisms and unions are least upper bounds in $\mathcal{C}(E)$. Secondly, as f is a bijection preserving and reflecting inclusion, it must preserve $\neg\subset$ and thus preserves cardinality. Hence, by Lemma 6.1.11 there is $\hat{f} : E \rightarrow F$ a map of event structures such that for all $x \in \mathcal{C}(E)$, $\hat{f}(x) = f(x)$. The same argument applies to f^{-1} , yielding an inverse to f by Lemma 6.1.12. \square

Paired with Lemma 6.1.10, we see that given two event structures E and F , an order-iso between the domains of configurations $\mathcal{C}(E)$ and $\mathcal{C}(F)$ induces a concrete *renaming* between E and F , a bijection between events preserving and reflecting all structure. In the future, this will provide our main tool to establish isomorphisms between event structures: by exhibiting an order-isomorphism between the domains of configurations.

Parallel composition. Next we introduce the *parallel composition* of event structures, an operation which shall play a central role throughout the development.

Definition 6.1.14. Consider E_1 and E_2 two event structures.

Their *simple parallel composition*, written $E_1 \parallel E_2$, has components:

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

This simply puts the two event structures side by side, with no interaction. We shall also apply this construction to configurations: if $x_1 \in \mathcal{C}(E_1)$ and $x_2 \in \mathcal{C}(E_2)$, then $x_1 \parallel x_2 \in \mathcal{C}(E_1 \parallel E_2)$ is defined as the same tagged disjoint union as above, i.e. $x_1 \parallel x_2 = x_1 + x_2$. Moreover, all configurations on $E_1 \parallel E_2$ have this form:

Lemma 6.1.15. Consider E_1 and E_2 two event structures. Then, we have

$$(- \parallel -) : \mathcal{C}(E_1) \times \mathcal{C}(E_2) \cong \mathcal{C}(E_1 \parallel E_2)$$

an order-isomorphism, with the order on $\mathcal{C}(E_1) \times \mathcal{C}(E_2)$ set as the product order.

Proof. Straightforward. \square

Along this isomorphism, all configurations of $E_1 \parallel E_2$ are uniquely written as pairs – accordingly, we shall often write $x_1 \parallel x_2 \in \mathcal{C}(E_1 \parallel E_2)$ for arbitrary configurations on a parallel composition. Note that up to isomorphism, $E_1 \parallel E_2$ is the *unique* event structure whose domain of configurations is order-isomorphic to $\mathcal{C}(E_1) \times \mathcal{C}(E_2)$ – this immediately follows from Lemma 6.1.15 and Proposition 6.1.13.

Basic properties of strategies. To put our preliminaries into play, we mention a few facts on strategies. The proofs are relatively elementary with the ingredients introduced so far. We omit them, but invite the reader to work them out as exercises.

Recall that if $\sigma : A$ is a prestrategy on game A , then events of σ inherit a polarity from A via $\text{pol}_\sigma(s) = \text{pol}_A(\partial_\sigma(s))$. Our first lemma states two properties of strategies in the particular case where the game A is forestial and alternating (which is almost always the case when interpreting types of programming languages).

Lemma 6.1.16. Consider a game A forestial and alternating as in Definition 3.1.3.

Then, we have the following properties:

- (1) σ is alternating: if $s_1 \rightarrow_\sigma s_2$, then $\text{pol}_\sigma(s_1) \neq \text{pol}_\sigma(s_2)$.
- (2) consider $s_1, s_2 \in x \in \mathcal{C}(\sigma)$ such that $\text{pol}_\sigma(s_1) = +$ or $\text{pol}_\sigma(s_2) = -$.
Then $s_1 \rightarrow_\sigma s_2$ iff $\partial_\sigma s_1 \rightarrow_A \partial_\sigma s_2$.

Thus for such games, the causal structure of strategies is very much constrained: it is alternating, and only immediate causal links $a_1^- \rightarrow_A a_2^+$ in the game may be postponed by the strategy. In general and even without alternation, a strategy may only introduce new immediate conflicts (with respect to the game) between positive events:




Figure 6.1: Example of an interaction and composition

Lemma 6.1.17. Consider A a game, and $\sigma : A$ a strategy.

If $s_1 \sim_{\sigma} s_2$ with $\text{pol}_{\sigma}(s_1) = -$ or $\text{pol}_{\sigma}(s_2) = -$, then $\partial_{\sigma}(s_1) \sim_A \partial_{\sigma}(s_2)$ as well.

Now, we move to the *composition* of (pre)strategies.

6.2 Composition of Prestrategies

Before we can phrase composition, we must define *strategies between games*. As in sequential games, a strategy *from a game A to a game B* is a strategy that plays on A and B in parallel, except that the polarities in A are reversed.

Strategies between games. If A is a game, its **dual** A^{\perp} has the same components as A , except for $\text{pol}_{A^{\perp}} = -\text{pol}_A$. The **simple parallel composition** of Definition 6.1.14 extends to games with $\text{pol}_{A \parallel B}(1, a) = \text{pol}_A(a)$ and $\text{pol}_{A \parallel B}(2, b) = \text{pol}_B(b)$. Finally, the **hom-game** $A \vdash B$ is defined as $A^{\perp} \parallel B$ – by Lemma 6.1.15 we have an order-iso

$$(- \vdash -) : \mathcal{G}(A) \times \mathcal{G}(B) \cong \mathcal{G}(A \vdash B),$$

and we often write configurations of $A \vdash B$ as $x_A \vdash x_B$ for $x_A \in \mathcal{G}(A)$ and $x_B \in \mathcal{G}(B)$.

Now if A and B are games, a **strategy from A to B** is simply defined as a strategy on $A \vdash B$ – the same definition applies to prestrategies.

Towards composition. Now, consider two strategies $\sigma : A \vdash B$ and $\tau : B \vdash C$ that we wish to compose, forming a composite strategy $\tau \odot \sigma : A \vdash C$.

Intuitively, the computation of $\tau \odot \sigma$ proceeds in two steps. First, we form the *interaction* $\tau \otimes \sigma$; an event structure displayed to $A \parallel B \parallel C$. The interaction has three kinds of events: those displayed to A are directly imported from σ , those displayed to C come from τ , while those in B are thought of *synchronizations* between events of σ and events of τ whose display on B match. The *causality* of $\tau \otimes \sigma$ is obtained as (the transitive closure) of the causal dependencies imposed by σ and by τ . Figure 6.1

displays the interaction of two strategies³, and again the conflict is imported from the two strategies. In Figure 6.1 on the right hand side, the synchronized events are grayed out while those in the outer interface remain in black – those are called *visible*. The composition $\tau \odot \sigma$ is obtained from $\tau \otimes \sigma$ by keeping the visible events only.

This intuition is quite natural, but rather hard to capture formally in a way that is both mathematically rigorous and a tractable basis for further developments – much more so than the composition of sequential strategies as in Part I, which is already non-trivial. And while our forthcoming development of composition will indeed follow the route above, it has been crucial in working with concurrent games that the composition of strategies may be characterized almost relationally, via the proposition:

Proposition 6.2.1. *Consider A, B, C games, and $\sigma : A \vdash B$ and $\tau : B \vdash C$ strategies. Then there is a strategy $\tau \odot \sigma : A \vdash C$, unique up to iso, s.t. there are order-isos:*

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

such that for $x^\sigma \in \mathcal{C}^+(\sigma)$ and $x^\tau \in \mathcal{C}^+(\tau)$ causally compatible,

$$\partial_{\tau \odot \sigma}(x^\tau \odot x^\sigma) = x_A^\sigma \vdash x_C^\tau$$

where we write $\partial_\sigma x^\sigma = x_A^\sigma \vdash x_B^\sigma$ and $\partial_\tau x^\tau = x_B^\tau \vdash x_C^\tau$.

Here, $\mathcal{C}^+(-)$ denotes the restriction of configurations to those that are *+covered*, i.e. whose maximal events are positive – without trailing Opponent moves. Likewise, *causally compatible* refers to a notion that we shall introduce soon: intuitively, $x^\sigma \in \mathcal{C}(\sigma)$ and $x^\tau \in \mathcal{C}(\tau)$ are causally compatible if they reach the same state in B (i.e. $x_B^\sigma = x_B^\tau$), and *their synchronization induces no deadlocks*. These two notions will be covered in depth soon, but the message for now is that for most purposes, we can ignore the concrete definition of $\tau \odot \sigma$ and treat its (+covered) configurations simply as certain pairs of (+covered) configurations of σ and τ . In addition, this completely captures the composition of σ and τ – there is *one* strategy (up to iso) whose +covered configurations are such causally compatible pairs, and it is $\tau \odot \sigma$.

Although noticed somewhat late (say around 2020), this proposition has become central in recent developments in concurrent games, as it lets us reason on the composition of strategies while leaving encapsulated the rather elaborate concrete definition. For most purposes, this is all you need to know about composition!

However, this still leaves us with the task of *proving* it, and thus to construct the composition concretely – first the *interaction*, and then its *hiding*. In the present section we focus on *prestrategies*, while Section 6.3 contains the extension to *strategies*.

6.2.1 Interaction of prestrategies

The next goal is to define *interactions* of prestrategies. But before that, we define formally the notion of *causal compatibility* appearing in Proposition 6.2.1.

³The game $B \multimap U$ is defined only in Section 7.1.4, and likewise for other forthcoming examples; hopefully this is not an obstacle to understanding. Though do note that in such examples, the game can be read from the diagram by only keeping dotted lines and ignoring immediate causal links \rightarrow .




Figure 6.2: Matching, secured




Figure 6.3: Matching, non-secured

Causally compatible pairs. Let us fix $\sigma : A \vdash B$ and $\tau : B \vdash C$ prestrategies. For configurations $x^\sigma \in \mathcal{C}(\sigma)$, $x^\tau \in \mathcal{C}(\tau)$, as a convention we write their display as

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

this convention will be used silently from now on.

First, we capture when $x^\sigma \in \mathcal{C}(\sigma)$ and $x^\tau \in \mathcal{C}(\tau)$ may successfully *synchronise*:

Definition 6.2.2. Two configurations $x^\sigma \in \mathcal{C}(\sigma)$ and $x^\tau \in \mathcal{C}(\tau)$ 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

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

defined on (the graph of) $\varphi[x^\sigma, x^\tau]$ from the 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 ; and this induces a synchronization between the events matching 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 6.2 and 6.3 two attempted synchronizations between configurations of the strategy on the left hand side of Figure 7.6 and that for the identity $\lambda x^\cup. x$. In both cases, the configurations are matching. In Figure 6.2, the synchronization is successful and yields a causally compatible pair. However, in Figure 6.3 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.

This disambiguates the statement of Proposition 6.2.1 – we want an *event structure* $\tau \odot \sigma$ whose (+-covered) configurations correspond to (+-covered) causally compatible pairs. Its construction involves several steps, starting with the *interaction*.

Interactions. Given $\sigma : A \vdash B$ and $\tau : B \vdash C$ we mean to construct their interaction $\tau \otimes \sigma$. An *interaction* is not quite a strategy, because it involves events in B that have no clear polarity. Accordingly, it is convenient to define a variation:

Definition 6.2.3. An *interaction* on A, B, C is an event structure $\mu = (|\mu|, \leq_\mu, \#_\mu)$ with

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

a display map subject to the following conditions:

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

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

In other words an interaction on A, B, C is a prestrategy on $A \parallel B \parallel C$, but it is conceptually clearer to keep the notions separate. Now we must define the interaction, written $\tau \otimes \sigma$, of σ and τ ; but what should be its *events*? Intuitively, each event of $\tau \otimes \sigma$ arises as the synchronization of an event in $\sigma \parallel C$ and one in $A \parallel \tau$ – note the padding out so that both $\sigma \parallel C$ and $A \parallel \tau$ cover $A \parallel B \parallel C$.

But it is not that simple, as can be observed in Figure 6.1, showing an interaction $\tau \otimes \sigma$ of $\sigma : B \multimap U$ and $\tau : B \multimap U \vdash N$. The two final events 0^+ of $\tau \otimes \sigma$ are two unsynchronized versions of the *same* event in τ . Yet they are different: they have different causal histories, following the earlier non-deterministic choice in τ . This has to do with an important property of the model we already mentioned: it retains the *non-deterministic branching point* – any event carries its entire causal history.

A pair of synchronized events with a causal history is a *prime secured bijection*:

Prime secured bijections. If $x^\sigma \in \mathcal{E}(\sigma)$ and $x^\tau \in \mathcal{E}(\tau)$ are causally compatible, from Definition 6.2.2 we know that the relation on the graph of $\varphi[x^\sigma, x^\tau]$ defined as

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

is acyclic: its reflexive transitive closure yields a partial order $\leq_{\varphi[x^\sigma, x^\tau]}$. From now on, for x^σ, x^τ causally compatible, we consider $\varphi[x^\sigma, x^\tau]$ partially ordered with $\leq_{\varphi[x^\sigma, x^\tau]}$.

Definition 6.2.4. A *secured bijection* between $\sigma : A \vdash B$ and $\tau : B \vdash C$ is any partial order of the form $\varphi[x^\sigma, x^\tau]$, for $x^\sigma \in \mathcal{E}(\sigma)$ and $x^\tau \in \mathcal{E}(\tau)$ causally compatible.

We write $\mathcal{B}(\sigma, \tau)$ for the set of secured bijections between σ and τ .

We shall be particularly interested in those secured bijections with a top element:

Definition 6.2.5. We say $\varphi \in \mathcal{B}(\sigma, \tau)$ is *prime* if \leq_φ has a top $\text{top}(\varphi) = (m, n)$.

We write $\mathcal{B}^\top(\sigma, \tau)$ for the set of prime secured bijections between σ and τ .

So we regard $\varphi \in \mathcal{B}^\top(\sigma, \tau)$ with $\text{top}(\varphi) = (m, n)$ as the synchronized pair (m, n) together with a causal explanation for that synchronization. Those will provide the *events* of the interaction. To proceed, we will need the lemma:

Lemma 6.2.6. Consider $x^\sigma, y^\sigma \in \mathcal{E}(\sigma)$, $x^\tau, y^\tau \in \mathcal{E}(\tau)$. Then, we have:

- (1) If x^σ, x^τ caus. comp., y^σ, y^τ caus. comp., $x^\sigma \cup y^\sigma \in \mathcal{E}(\sigma)$ and $x^\tau \cup y^\tau \in \mathcal{E}(\tau)$, then $x^\sigma \cup y^\sigma$ and $x^\tau \cup y^\tau$ are causally compatible.
- (2) If $x^\sigma \cup y^\sigma \in \mathcal{E}(\sigma)$ and $x^\tau \cup y^\tau \in \mathcal{E}(\tau)$ are causally compatible, if x^σ, x^τ matching and y^σ, y^τ matching, then they are causally compatible,

and in both cases, we have $\varphi[x^\sigma \cup y^\sigma, x^\tau \cup y^\tau] = \varphi[x^\sigma, x^\tau] \cup \varphi[y^\sigma, y^\tau]$.

Proof. Straightforward by local injectivity of ∂_σ and ∂_τ . \square

We are now in position to form an event structure:

Proposition 6.2.7. *The tuple $(|\tau \otimes \sigma|, \leq_{\tau \otimes \sigma}, \#_{\tau \otimes \sigma})$ defined with*

$$\begin{aligned} \text{events:} \quad & |\tau \otimes \sigma| = \{\varphi[x^\sigma, x^\tau] \mid \text{prime, for } x^\sigma, x^\tau \text{ causally compatible}\} \\ \text{causality:} \quad & \phi \leq_{\tau \otimes \sigma} \psi \Leftrightarrow \phi \subseteq \psi \\ \text{conflict:} \quad & \neg(\phi \#_{\tau \otimes \sigma} \psi) \Leftrightarrow \exists x^\sigma, x^\tau \text{ causally compatible, } \phi, \psi \subseteq \varphi[x^\sigma, x^\tau]. \end{aligned}$$

forms an event structure. Moreover, there are monotone functions

$$\begin{aligned} \chi_{\sigma, \tau} : \mathcal{C}(\tau \otimes \sigma) &\rightarrow \mathcal{B}(\sigma, \tau) \\ x &\mapsto \cup x \\ \bar{\chi}_{\sigma, \tau} : \mathcal{B}(\sigma, \tau) &\rightarrow \mathcal{C}(\tau \otimes \sigma) \\ \varphi &\mapsto \{\psi \in \mathcal{B}^\top(\sigma, \tau) \mid \psi \subseteq \varphi\} \end{aligned}$$

forming an order-isomorphism, with both sets ordered by inclusion.

Proof. The conditions for an event structure are easily checked. It is immediate that $\bar{\chi}_{\sigma, \tau}$ is well-formed, for $\chi_{\sigma, \tau}$ it follows from Lemma 6.2.6; both maps are monotone with respect to inclusion. Next we observe that for all $\varphi \in \mathcal{B}(\sigma, \tau)$, we have

$$\varphi = \cup \{\psi \in \mathcal{B}^\top(\sigma, \tau) \mid \psi \subseteq \varphi\};$$

indeed if $(m, n) \in \varphi$ we set $\psi = [(m, n)]_\varphi$ the set of all pairs $(m', n') \in \varphi$ such that $(m', n') \leq_\varphi (m, n)$ (recall from above Definition 6.2.4 that causal compatibility endows the graph of any secured bijection with a partial order); by construction, $\psi \in \mathcal{B}^\top(\sigma, \tau)$ and $\psi \subseteq \varphi$, hence $(m, n) \in \cup \{\psi \in \mathcal{B}^\top(\sigma, \tau) \mid \psi \subseteq \varphi\}$ – the other inclusion is direct. Altogether, this shows that $\chi_{\sigma, \tau} \circ \bar{\chi}_{\sigma, \tau} \varphi = \varphi$ for all $\varphi \in \mathcal{B}(\sigma, \tau)$ a secured bijection.

Now consider $X \in \mathcal{C}(\tau \otimes \sigma)$ along with $\varphi[x^\sigma, x^\tau] \in \bar{\chi}_{\sigma, \tau} \circ \chi_{\sigma, \tau} X$, and take $(m, n) = \text{top}(\varphi[x^\sigma, x^\tau])$. By definition, $(m, n) \in \cup X$, but this means that there is $\psi \in X$ such that $(m, n) \in \psi$. But then it follows from local injectivity that $\varphi[x^\sigma, x^\tau] \subseteq \psi$, hence $\varphi[x^\sigma, x^\tau] \in X$ as X is down-closed. The other inclusion is straightforward. \square

We have now constructed an event structure with configurations order-isomorphic to the secured bijections between σ and τ . But furthermore:

Proposition 6.2.8. *There is an order-isomorphism*

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

with causally compatible pairs ordered by pairwise inclusion.

Proof. By definition, $\mathcal{B}(\sigma, \tau)$ is in order-isomorphism with the set of causally compatible pairs (x^σ, x^τ) . We conclude by composition with that of Proposition 6.2.7. \square

This is starting to look like Proposition 6.2.1. In the sequel, we will almost always reason on $\mathcal{C}(\tau \otimes \sigma)$ via this correspondence, and routinely write any configuration $x \in \mathcal{C}(\tau \otimes \sigma)$ as $x^\tau \otimes x^\sigma \in \mathcal{C}(\tau \otimes \sigma)$, without mentioning the use of this proposition.

In passing, let us mention a few lemmas. Firstly, using the mapification lemma will require us to have information on the cardinal of configurations of the interaction:

Lemma 6.2.9. *Consider $x^\tau \otimes x^\sigma \in \mathcal{C}(\tau \otimes \sigma)$.*

Then, $\text{card}(x^\tau \otimes x^\sigma) = \text{card}(x_A^\sigma) + \text{card}(x_B^\sigma) + \text{card}(x_C^\tau)$.

Proof. Straightforward from the Propositions 6.2.7 and 6.2.8. \square

Similarly, we need compatibility of configurations of the interaction with unions:

Lemma 6.2.10. *Consider $x^\sigma, y^\sigma \in \mathcal{C}(\sigma)$ and $x^\tau, y^\tau \in \mathcal{C}(\tau)$ satisfying either of the conditions of Lemma 6.2.6. Then, we have the equality:*

$$(x^\tau \cup y^\tau) \otimes (x^\sigma \cup y^\sigma) = (x^\tau \otimes x^\sigma) \cup (y^\tau \otimes y^\sigma).$$

Proof. By Lemma 6.2.6, across the order-isomorphism of Proposition 6.2.7. \square

The interaction pullback. As an important aside, we have:

Lemma 6.2.11. *There are maps of event structures Π_σ and Π_τ making the diagram*

$$\begin{array}{ccc} & \tau \otimes \sigma & \\ \Pi_\sigma \swarrow & \downarrow \vee & \searrow \Pi_\tau \\ \sigma \parallel C & & A \parallel \tau \\ \downarrow \partial_\sigma \parallel C & & \downarrow A \parallel \partial_\tau \\ & A \parallel B \parallel C & \end{array}$$

a pullback in the category of event structures and their maps. Moreover, for any $x^\tau \otimes x^\sigma \in \mathcal{C}(\tau \otimes \sigma)$, we have $\Pi_\sigma(x^\tau \otimes x^\sigma) = x^\sigma \parallel x_C^\tau$ and $\Pi_\tau(x^\tau \otimes x^\sigma) = x_A^\sigma \parallel x^\tau$.

Proof. We first define the projection maps, as

$$\begin{array}{ll} \Pi_\sigma : \tau \otimes \sigma \rightarrow \sigma \parallel C & \Pi_\tau : \tau \otimes \sigma \rightarrow A \parallel \tau \\ \varphi \mapsto \pi_1(\text{top}(\varphi)) & \varphi \mapsto \pi_2(\text{top}(\varphi)), \end{array}$$

and check the two conditions for these to be maps of event structures. For *rule-abiding*, consider $x \in \mathcal{C}(\tau \otimes \sigma)$. By Proposition 6.2.7, $\cup x \in \mathcal{B}(\sigma, \tau)$. But furthermore,

$$\cup x : \Pi_\sigma x \simeq \Pi_\tau x. \quad (6.5)$$

Indeed, if $m \in \Pi_\sigma x$ there is some $\varphi \in x$ such that $m = \pi_1(\text{top}(\varphi))$, but then clearly $m \in \text{dom}(\cup x)$. Reciprocally, if $m \in \text{dom}(\cup x)$, this means there is $\varphi \in x$ such that $(m, n) \in \varphi$ for some n . But then, considering the set defined as

$$[(m, n)]_\varphi = \{(m', n') \in \varphi \mid (m', n') \leq_\varphi (m, n)\},$$

one can check that $[(m, n)]_\varphi \in \mathcal{B}^\top(\sigma, \tau)$, and $[(m, n)]_\varphi \subseteq \varphi$ by construction, so $[(m, n)]_\varphi \in x$ as x is *down-closed*. But so, $m = \Pi_\sigma[(m, n)]_\varphi$ so $m \in \Pi_\sigma x$ as required. It also follows that $\Pi_\sigma x \in \mathcal{C}(\sigma \parallel C)$, and using (6.5) and Proposition 6.2.7,

$$\begin{aligned} \Pi_\sigma(x^\tau \otimes x^\sigma) &= \Pi_\sigma(\bar{\chi}_{\sigma, \tau} \varphi[x^\sigma, x^\tau]) \\ &= \text{dom}(\cup(\bar{\chi}_{\sigma, \tau} \varphi[x^\sigma, x^\tau])) \\ &= \text{dom}(\varphi[x^\sigma, x^\tau]) \end{aligned}$$

which is $x^\sigma \parallel x_C^\tau$ – and likewise, $\Pi_\tau(x^\tau \otimes x^\sigma) = x_A^\sigma \parallel x^\tau \in \mathcal{C}(A \parallel \tau)$. For local injectivity, assume that $\varphi, \psi \in x$ are such that $\Pi_\sigma \varphi = \Pi_\sigma \psi$, write it m . But $\varphi, \psi \subseteq \cup x \in \mathcal{B}(\sigma, \tau)$, and there is a unique n such that $(m, n) \in \cup x$ as it is a bijection. But we must have $\varphi = [(m, n)]_{\cup x}$ and likewise for ψ , so $\varphi = \psi$.

That the diagram commutes and satisfies the universal property of the diagram is routine, using the above along with Lemma 6.1.11 (the conditions for mapification come from Lemmas 6.2.9 and 6.2.10) and Proposition 6.2.7. \square

In fact, the category **ES** of event structures and their maps has all pullbacks. The construction of $\tau \otimes \sigma$ given here closely follows the construction of pullbacks in **ES**, which in turn is a restriction of the usual synchronizing product of event structures. So this is old technology here, essentially the same used by Winskel in the early 80s to give semantics to parallel composition of CCS processes [Winskel, 1982].

In the sequel we shall not exploit the universal property of this pullback, but we will often refer to its projections: after this section we shall often treat the interaction construction as a black box. The *projections* will be our main device to read back an event $p \in \tau \otimes \sigma$ as a synchronization of $\Pi_\sigma p \in \sigma \parallel C$ and $\Pi_\tau p \in A \parallel \tau$.

Finally, we have seen that configurations of the interaction correspond to secured bijections. But the connection is even tighter: both configurations and secured bijections are partially ordered sets, and it will be important for later on to keep in mind that this correspondence is, locally, an order-isomorphism. This is expressed by:

Lemma 6.2.12. *Consider $x^\sigma \in \mathcal{C}(\sigma)$ and $x^\tau \in \mathcal{C}(\tau)$ causally compatible.*

Then, there is an order-isomorphism $v[x^\sigma, x^\tau] : x^\tau \otimes x^\sigma \cong \varphi[x^\sigma, x^\tau]$ such that

$$\begin{array}{ccccc} & x^\tau \otimes x^\sigma & & & \\ \Pi_\sigma \swarrow & & \downarrow v[x^\sigma, x^\tau] & & \searrow \Pi_\tau \\ x^\sigma \parallel x_C^\tau & & & & x_A^\sigma \parallel x^\tau \\ \swarrow \pi_1 & & \downarrow & & \searrow \pi_2 \\ & \varphi[x^\sigma, x^\tau] & & & \end{array}$$

commutes, in the category of finite sets and bijections. Moreover, this assignment is monotone, in the sense that if $x^\sigma \subseteq y^\sigma$ and $x^\tau \subseteq y^\tau$, then $v[x^\sigma, x^\tau] \subseteq v[y^\sigma, y^\tau]$.

Proof. We set the order-isomorphism $v[x^\sigma, x^\tau]$ as follows:

$$\begin{aligned} v[x^\sigma, x^\tau] : x^\tau \otimes x^\sigma &\rightarrow \varphi[x^\sigma, x^\tau] \\ \psi &\mapsto \text{top}(\psi), \end{aligned}$$

it is a direct that this defines an order-iso, and that this assignment is monotone. \square

The interaction. At last, this provides us with the ingredients for the interaction.

Proposition 6.2.13. *There is an interaction $\tau \otimes \sigma$, unique up to iso, s.t. there is*

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

an order-iso s.t. $\partial_{\tau \otimes \sigma}(x^\tau \otimes x^\sigma) = x_A^\sigma \parallel x_B \parallel x_C^\tau$ for all x^σ, x^τ causally compatible.

Proof. Existence. The event structure $\tau \otimes \sigma$ is defined in Proposition 6.2.7, with display map $\partial_{\tau \otimes \sigma} : \tau \otimes \sigma \rightarrow A \parallel B \parallel C$ defined as either way around the square in Lemma 6.2.11. The order-iso for configurations comes from Proposition 6.2.8.

Uniqueness. Immediate from the order-isos, by Lemma 6.1.12. \square

Here, an isomorphism of interactions simply means an isomorphism of event structures commuting with the display maps. The construction of $\tau \otimes \sigma$ is elaborate – there is no going around it; but fortunately, Proposition 6.2.13 packages most of what we need about the interaction, and may often be used as a black box. It is suspiciously close to Proposition 6.2.1, which may be surprising since this is not composition but interaction, without hiding. We shall see that the restriction to $+$ -covered configurations in Proposition 6.2.1 does a lot to span the distance between interaction and composition.

Structure of causality. Proposition 6.2.13 gives a neat *state-based* description of the interaction, putting the focus on configurations. This is sufficient for many purposes, but it is sometimes necessary to reason on individual events of the interaction.

As explained before, any event $p \in \tau \otimes \sigma$ may be regarded as a synchronization between its projections $\Pi_\sigma p \in \sigma \parallel C$ and $\Pi_\tau p \in A \parallel \tau$, along with a causal history for this synchronization. These projections help us classify every $p \in \tau \otimes \sigma$ into:

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

In case (1), the only relevant projection is $\Pi_\sigma(p) = (1, s)$ as $\Pi_\tau(p) = \partial_\sigma(s)$. We write $p_\sigma = s$ and p_τ is undefined, and we say that p **occurs in** A . In case (3), the relevant projection is $\Pi_\tau(p) = (2, t)$ as $\Pi_\sigma(p) = \partial_\tau(t)$. We write $p_\tau = t$ and p_σ is undefined, and p **occurs in** C . Finally, in case (2) the two projections $\Pi_\sigma(p) = (1, s)$ and $\Pi_\tau(p) = (2, t)$ are relevant, but we must have $\partial_\sigma(s) = (2, b)$ and $\partial_\tau(t) = (1, b)$ for some $b \in B$. We write $p_\sigma = s$, $p_\tau = t$, we say that p **occurs in** B . If p occurs in A or C we also say that it is **visible**, while if it occurs in B it is **synchronized** or **invisible**.

Immediate causal links between events of $\tau \otimes \sigma$ must originate from σ or τ :

Lemma 6.2.14. *Consider $p, p' \in \tau \otimes \sigma$, and assume $p \rightarrow_{\tau \otimes \sigma} p'$.*

Then $p_\sigma \rightarrow_\sigma p'_\sigma$ or $p_\tau \rightarrow_\tau p'_\tau$, where p_σ, p_τ are defined whenever used.

Proof. Consider $x^\tau \otimes x^\sigma \in \mathcal{C}(\tau \otimes \sigma)$ such that $p, p' \in x^\tau \otimes x^\sigma$. Write $\text{top}(p) = (m, n)$ and $\text{top}(p') = (m', n')$. By Lemma 6.2.12, we have an immediate causal link

$$(m, n) \rightarrow_{\varphi[x^\sigma, x^\tau]} (m', n'),$$




Figure 6.4: Example of an interaction with polarity information

which easily implies, by definition of $\leq_{\varphi[x^\sigma, x^\tau]}$, that $m \rightarrow_{\sigma \parallel C} m'$ or $n \rightarrow_{A \parallel \tau} n'$.

If $m \rightarrow_{\sigma \parallel C} m'$, then either $m = (1, s)$, $m' = (1, s')$ and $s \rightarrow_\sigma s'$ with $s = p_\sigma$ and $s' = p'_\sigma$ defined as required; or, $m = (2, c)$ and $m' = (2, c')$ with $c \rightarrow_C c'$. In that case, $n = (2, t)$ and $n' = (2, t')$ and $t = p_\tau$ and $t' = p'_\tau$ defined. Moreover, $c \rightarrow_C c'$ implies $\partial_\tau(t) \rightarrow_{B \vdash C} \partial_\tau(t')$, so $t <_\tau t'$ by Lemma 6.1.8. If this causal link was not immediate, this would directly contradict $(m, n) \rightarrow_{\varphi[x^\sigma, x^\tau]} (m', n')$ – so $p_\tau \rightarrow_\tau p'_\tau$ as required. \square

For σ and τ strategies, we can track down the responsible of a move via a polarity analysis. Events of $\tau \otimes \sigma$ cannot sensibly be assigned a polarity in $\{-, +\}$, as σ and τ disagree on B . A more useful polarity is $\text{pol}_{\tau \otimes \sigma} : |\tau \otimes \sigma| \rightarrow \{-, \ell, \nu\}$ set as:

$$\begin{aligned} \text{pol}_{\tau \otimes \sigma}(p) &= \ell && \text{if } p_\sigma \text{ is defined and } \text{pol}_\sigma(p_\sigma) = +, \\ \text{pol}_{\tau \otimes \sigma}(p) &= \nu && \text{if } p_\tau \text{ is defined and } \text{pol}_\tau(p_\tau) = +, \\ \text{pol}_{\tau \otimes \sigma}(p) &= - && \text{otherwise.} \end{aligned}$$

We show in Figure 6.4 a version of Figure 6.1 with those polarities. Then:

Lemma 6.2.15. *If $\sigma : A \vdash B$ and $\tau : B \vdash C$ are strategies and $p \rightarrow_{\tau \otimes \sigma} p'$, then*

- (1) *if $\text{pol}_{\tau \otimes \sigma}(p') = \ell$, then $p_\sigma \rightarrow_\sigma p'_\sigma$,*
- (2) *if $\text{pol}_{\tau \otimes \sigma}(p') = \nu$, then $p_\tau \rightarrow_\tau p'_\tau$,*
- (3) *if $\text{pol}_{\tau \otimes \sigma}(p') = -$, then $\partial_{\tau \otimes \sigma}(p) \rightarrow_{A \parallel B \parallel C} \partial_{\tau \otimes \sigma}(p')$*

where again, $p_\sigma, p'_\sigma, p_\tau$ and p'_τ are defined whenever used.

Proof. (1) By Lemma 6.2.14, p_σ, p'_σ defined and $p_\sigma \rightarrow_\sigma p'_\sigma$ and we are done; or p_τ, p'_τ defined and $p_\tau \rightarrow_\tau p'_\tau$. Since $\text{pol}_{\tau \otimes \sigma}(p') = \ell$, $\text{pol}_\tau(p'_\tau) = -$. By courtesy, $\partial_\tau(p_\tau) \rightarrow_{B \vdash C} \partial_\tau(p'_\tau)$; hence m occurs in B and $\partial_\sigma(p_\sigma) \rightarrow_{A \vdash B} \partial_\sigma(p'_\sigma)$. By Lemma 6.1.8, $p_\sigma <_\sigma p'_\sigma$, and it is direct that the causality must be immediate. (2) is symmetric.

(3) Assume p' occurs in A , the other case is symmetric. In that case only p'_σ is defined, so Lemma 6.2.14 entails that p_σ is defined and $p_\sigma \rightarrow_\sigma p'_\sigma$. But $\text{pol}_\sigma(p'_\sigma) = -$, so by courtesy $\partial_\sigma(p_\sigma) \rightarrow_{A \vdash B} \partial_\sigma(p'_\sigma)$, from which the conclusion follows. \square

Next, we perform the hiding and introduce the *composition* of prestrategies.

6.2.2 Composition of prestrategies

Hiding. Composition consists simply in removing all synchronized events:

Definition 6.2.16. The *composition* of $\sigma : A \vdash B$ and $\tau : B \vdash C$ comprises:

$$\begin{aligned} |\tau \odot \sigma| &= \{p \in \tau \otimes \sigma \mid p \text{ occurs in } A \text{ or } C\}, \\ p_1 \leq_{\tau \odot \sigma} p_2 &\Leftrightarrow p_1 \leq_{\tau \otimes \sigma} p_2, \\ p_1 \#_{\tau \odot \sigma} p_2 &\Leftrightarrow p_1 \#_{\tau \otimes \sigma} p_2. \end{aligned}$$

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

It is straightforward that $\tau \odot \sigma$ is an event structure. In Figure 6.4, the composition simply keeps the events in black. This means that it has two conflicting positive events, both corresponding to 0^+ : 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 $\tau \otimes \sigma$. Thus an event of the composition always carries a unique causal explanation: *itself*, regarded as an event of the interaction $\tau \otimes \sigma$. Accordingly, we set:

Definition 6.2.17. Consider $x \in \mathcal{C}(\tau \odot \sigma)$. Its *interaction witness* is defined as

$$[x]_{\tau \otimes \sigma} = \{p \in \tau \otimes \sigma \mid \exists p' \in x, p \leq_{\tau \otimes \sigma} p'\} \in \mathcal{C}(\tau \otimes \sigma)$$

its down-closure in $\tau \otimes \sigma$.

This entails – and it shall play a crucial role in the development – that configurations of the composition are in one-to-one correspondence with those configurations of the interaction whose maximal events are visible. From this idea, we get:

Proposition 6.2.18. There is an order-isomorphism

$$(- \odot -) : \{(x^\tau, x^\sigma) \in \mathcal{C}(\tau) \times \mathcal{C}(\sigma) \mid x^\sigma, x^\tau \text{ minimal caus. comp.}\} \simeq \mathcal{C}(\tau \odot \sigma)$$

where x^σ, x^τ *minimal* means that the maximal elements of $\phi[x^\sigma, x^\tau]$ occur in A or C .

Proof. We decompose the order-isomorphism in two steps.

As a first step, we note that the isomorphism of Proposition 6.2.13 refines to:

$$(- \otimes -) : \{(x^\tau, x^\sigma) \in \mathcal{C}(\tau) \times \mathcal{C}(\sigma) \mid x^\sigma, x^\tau \text{ min. caus. comp.}\} \simeq \mathcal{C}^v(\tau \otimes \sigma)$$

where $\mathcal{C}^v(\tau \otimes \sigma)$ refers to configurations whose maximal events are *visible*, i.e. occur in A or C – this is clear by Lemma 6.2.12. Next, we have the order-iso formed with:

$$\begin{array}{ccc} \mathcal{C}^v(\tau \otimes \sigma) & \rightarrow & \mathcal{C}(\tau \odot \sigma) \\ x & \mapsto & x \cap V \end{array} \qquad \begin{array}{ccc} \mathcal{C}(\tau \odot \sigma) & \rightarrow & \mathcal{C}^v(\tau \otimes \sigma) \\ x & \mapsto & [x]_{\tau \otimes \sigma} \end{array}$$

where V denotes visible events – it is straightforward that this defines an order-iso. \square

This shows how even through the hiding operation, configurations of the composition can still be regarded as certain pairs of configurations of the compound strategies. We are certainly starting to get close to our target statement of Proposition 6.2.1.

Note that the operation from left to right is well-defined even if $x^\sigma \in \mathcal{C}(\sigma)$ and $x^\tau \in \mathcal{C}(\tau)$ are causally compatible but not *minimal* causally compatible. In that case we shall still write $x^\tau \odot x^\sigma = (x^\tau \otimes x^\sigma) \cap V \in \mathcal{C}(\tau \odot \sigma)$, even though x^σ and x^τ cannot necessarily be recovered if the interaction has some maximal synchronized events.

The next lemma states a simple property of this construction, useful for reference:

Lemma 6.2.19. *For $x^\sigma \in \mathcal{C}(\sigma)$ and $x^\tau \in \mathcal{C}(\tau)$ minimal causally compatible, we have*

$$x^\tau \odot x^\sigma = (x^\tau \otimes x^\sigma) \cap V \quad x^\tau \otimes x^\sigma = [x^\tau \odot x^\sigma]_{\tau \otimes \sigma},$$

where V is the set of visible events.

Proof. By def. of $x^\tau \odot x^\sigma$ through the order-iso $\mathcal{C}^v(\tau \otimes \sigma) \simeq \mathcal{C}(\tau \odot \sigma)$ above. \square

It will also be useful later on that this is compatible with unions:

Lemma 6.2.20. *Consider $x^\sigma, y^\sigma \in \mathcal{C}(\sigma)$ and $x^\tau, y^\tau \in \mathcal{C}(\tau)$ satisfying either of the conditions of Lemma 6.2.6. Then, we have the equality:*

$$(x^\tau \odot x^\sigma) \cup (y^\tau \odot y^\sigma) = (x^\tau \cup y^\tau) \odot (x^\sigma \cup y^\sigma).$$

Proof. By Lemma 6.2.10, we have $(x^\tau \cup y^\tau) \otimes (x^\sigma \cup y^\sigma) = (x^\tau \otimes x^\sigma) \cup (y^\tau \otimes y^\sigma)$, thus $(x^\tau \cup y^\tau) \odot (x^\sigma \cup y^\sigma) = (x^\tau \odot x^\sigma) \cup (y^\tau \odot y^\sigma)$ by restricting to visible events. \square

Finally, we relate immediate causality in the interaction and composition:

Lemma 6.2.21. *Consider $p, p' \in \tau \odot \sigma$. If $p \rightarrow_{\tau \odot \sigma} p'$, then we have*

$$p \rightarrow_{\tau \otimes \sigma} q_1 \rightarrow_{\tau \otimes \sigma} \dots \rightarrow_{\tau \otimes \sigma} q_n \rightarrow_{\tau \otimes \sigma} p'$$

for some $q_1, \dots, q_n \in \tau \otimes \sigma$ synchronized.

Proof. Straightforward by definition. \square

We wrap up with the main proposition for composition of prestrategies:

Proposition 6.2.22. *Consider $\sigma : A \vdash B$ and $\tau : B \vdash C$ prestrategies.*

Then there is a prestrategy $\tau \odot \sigma$, unique up to iso, s.t. there is an order-iso:

$$(- \odot -) : \{(x^\tau, x^\sigma) \in \mathcal{C}(\tau) \times \mathcal{C}(\sigma) \mid x^\sigma, x^\tau \text{ min. caus. comp.}\} \simeq \mathcal{C}(\tau \odot \sigma)$$

s.t. for all $x^\sigma \in \mathcal{C}(\sigma), x^\tau \in \mathcal{C}(\tau)$ min. caus. comp., $\partial_{\tau \odot \sigma}(x^\tau \odot x^\sigma) = x_A^\sigma \vdash x_C^\tau$.

$$\begin{array}{lll}
(\partial_\sigma \parallel C \parallel D)(m) \triangleleft_\sigma (\partial_\sigma \parallel C \parallel D)(m') & \text{where } m <_{\sigma \parallel C \parallel D} m' & (m, m' \in x^\sigma \parallel x_C \parallel x_D), \\
(A \parallel \partial_\tau \parallel D)(m) \triangleleft_\tau (A \parallel \partial_\tau \parallel D)(m') & \text{where } m <_{A \parallel \tau \parallel D} m' & (m, m' \in x_A \parallel x^\tau \parallel x_D), \\
(A \parallel B \parallel \partial_\delta)(m) \triangleleft_\delta (A \parallel B \parallel \partial_\delta)(m') & \text{where } m <_{A \parallel B \parallel \delta} m' & (m, m' \in x_A \parallel x_B \parallel x^\delta),
\end{array}$$

Figure 6.5: Ternary causal compatibility

Proof. Existence. The event structure, along with the display map on events, are defined in Definition 6.2.16, with order-isomorphism defined in Proposition 6.2.18. We must check that $\partial_{\tau \odot \sigma}$ is a map of event structures. By Proposition 6.2.13, $\partial_{\tau \odot \sigma}(x^\tau \otimes x^\sigma) = x_A^\sigma \parallel x_B \parallel x_C^\tau$, hence by definition of $\partial_{\tau \odot \sigma}$, $\partial_{\tau \odot \sigma}(x^\tau \odot x^\sigma) = x_A^\sigma \vdash x_C^\tau$. Now as $\partial_{\tau \odot \sigma}$ is a map of event structures, it induces a bijection between $x^\tau \otimes x^\sigma$ and $x_A^\sigma \parallel x_B \parallel x_C^\tau$, which by definition of $\tau \odot \sigma$ restricts to a bijection between $x^\tau \odot x^\sigma$ and $x_A^\sigma \vdash x_C^\tau$, so

$$\text{card}(x^\tau \odot x^\sigma) = \text{card}(x_A^\sigma) + \text{card}(x_C^\tau)$$

ensuring that $\partial_{\tau \odot \sigma}$ preserves cardinal. By Lemma 6.2.20 it follows that it also preserves unions, and hence forms a map of event structures by Lemma 6.1.11.

Uniqueness up to strong iso is straightforward by Lemma 6.1.12. \square

6.2.3 The associator

We have defined the composition of prestrategies; not yet that of strategies, as we have not yet established that composition preserves receptivity and courtesy. Before doing that we shall study *associativity* of composition, as it turns out that it holds independently of the additional conditions for strategies. Thus fix now prestrategies $\sigma : A \vdash B$, $\tau : B \vdash C$ and $\delta : C \vdash D$. We aim to show that composition is associative up to isomorphism, *i.e.* that there is an isomorphism of prestrategies, called the *associator*:

$$a_{\sigma, \tau, \delta} : (\delta \odot \tau) \odot \sigma \cong \delta \odot (\tau \odot \sigma).$$

Proposition 6.2.22 makes associativity seem almost trivial: as configurations of the composition are identified to certain pairs of configurations, mapping $(\delta \odot \tau) \odot \sigma$ to $\delta \odot (\tau \odot \sigma)$ is only a matter of sending $(x^\delta \odot x^\tau) \odot x^\sigma$ to $x^\delta \odot (x^\tau \odot x^\sigma)$, *i.e.* by associativity of the set-theoretic cartesian product – as for associativity for the composition of spans. But as Proposition 6.2.22 makes formal, composition of strategies is composition of spans, *plus deadlocks*, captured by the *causal compatibility* constraint; and we must also prove that causal compatibility is compatible with associativity.

For this, it is cleaner to introduce *causally compatible triples*:

Definition 6.2.23. Consider $x^\sigma \in \mathcal{C}(\sigma)$, $x^\tau \in \mathcal{C}(\tau)$ and $x^\delta \in \mathcal{C}(\delta)$. We set conditions:

- matching: if $x_B^\sigma = x_B^\tau = x_B$ and $x_C^\tau = x_C^\delta = x_C$,
- causally compatible: if they are matching and the relation $\triangleleft = \triangleleft_\sigma \cup \triangleleft_\tau \cup \triangleleft_\delta$ on $x_A \parallel x_B \parallel x_C \parallel x_D$ defined in Figure 6.5, is acyclic.
- minimal: if events maximal for \triangleleft^* occur in A or D .

where $\partial_\sigma x^\sigma = x_A \parallel x_B^\sigma$, $\partial_\tau x^\tau = x_B^\tau \parallel x_C^\tau$, and $\partial_\delta x^\delta = x_C^\delta \parallel x_D$.

Crucially, ternary causal compatibility is equivalent to the two skewed versions:

Lemma 6.2.24. *For $x^\sigma \in \mathcal{C}(\sigma)$, $x^\tau \in \mathcal{C}(\tau)$, $x^\delta \in \mathcal{C}(\delta)$, the following are equivalent:*

- (1) $x^\sigma, x^\tau, x^\delta$ are minimal causally compatible,
- (2) x^σ, x^τ are min. caus. comp. and $x^\tau \odot x^\sigma, x^\delta$ are min. caus. comp.,
- (3) x^τ, x^δ are min. caus. comp. and $x^\sigma, x^\delta \odot x^\tau$ are min. caus. comp.

Proof. (1) \Rightarrow (2). For x^σ, x^τ causally compatible, a cycle in $\varphi[x^\sigma, x^\tau]$ routinely corresponds to a cycle in $x_A \parallel x_B \parallel x_C$ for the relation $\triangleleft_\sigma \cup \triangleleft_\tau$ defined as:

$$\begin{aligned} (\partial_\sigma \parallel C)(m) \triangleleft_\sigma (\partial_\sigma \parallel C)(m') & \text{ where } m <_{\sigma \parallel C} m' \quad (m, m' \in x^\sigma \parallel x_C), \\ (A \parallel \partial_\tau)(m) \triangleleft_\tau (A \parallel \partial_\tau)(m') & \text{ where } m <_{A \parallel \tau} m' \quad (m, m' \in x_A \parallel x^\tau) \end{aligned}$$

which immediately transports to a cycle in \triangleleft on $x_A \parallel x_B \parallel x_C \parallel x_D$, contradiction. For minimality, consider m maximal in $x_A \parallel x_B \parallel x_C$ for $\triangleleft_\sigma \cup \triangleleft_\tau$; assume, seeking a contradiction, that it is in B . Then this event considered in $x_A \parallel x_B \parallel x_C \parallel x_D$, still denoted by m by abuse of notation, cannot be maximal for \triangleleft by minimality. Thus there is $m \triangleleft n$ for n in $x_A \parallel x_B \parallel x_C \parallel x_D$. But by construction of \triangleleft , n must be in A, B, C and $m \triangleleft_\sigma \cup \triangleleft_\tau n$ in $x_A \parallel x_B \parallel x_C$, contradiction. Hence, x^σ, x^τ minimal.

For $x^\tau \odot x^\sigma, x^\delta$ causally compatible, a cycle in $\varphi[x^\tau \odot x^\sigma, x^\delta]$ corresponds to a cycle in $x_A \parallel x_C \parallel x_D$ for the relation $\triangleleft_{\tau \odot \sigma} \cup \triangleleft_\delta$ obtained as above. But this immediately gives a cycle in $x_A \parallel x_B \parallel x_C \parallel x_D$ for $(\triangleleft_\sigma \cup \triangleleft_\tau)^* \cup \triangleleft_\delta$, which gives a cycle in $x_A \parallel x_B \parallel x_C \parallel x_D$ for \triangleleft , contradiction. For minimality, an event m maximal for $\triangleleft_{\sigma, \tau}$ and \triangleleft_δ cannot be in C : otherwise there is n in A, D such that $m \triangleleft^* n$, but then it is straightforward that $m (\triangleleft_{\sigma, \tau} \cup \triangleleft_\delta)^* n$, contradicting its maximality.

(2) \Rightarrow (1). Consider a cycle $m_1 \triangleleft \dots \triangleleft m_k \triangleleft m_1$. If the cycle remains in A, B, C , we get a contradiction to x^σ, x^τ causally compatible. Otherwise, removing all moves in B , we get a cycle in $(\triangleleft_{\sigma, \tau} \cup \triangleleft_\delta)$ hence a contradiction to $x^\tau \odot x^\sigma, x^\delta$ causally compatible. For minimality, consider m maximal for \triangleleft^* . By minimality of x^σ, x^τ it is not in B , and by minimality of $x^\tau \odot x^\sigma, x^\delta$ it cannot be in C ; hence it is in A or D .

(1) \Rightarrow (3), (3) \Rightarrow (1). Symmetric to the two implications above. \square

Altogether, this gives us all the ingredients required for the associator:

Proposition 6.2.25. *Consider $\sigma : A \vdash B$, $\tau : B \vdash C$ and $\delta : C \vdash D$ prestrategies.*

*Then there is a unique isomorphism, the **associator**:*

$$a_{\sigma, \tau, \delta} : (\delta \odot \tau) \odot \sigma \cong \delta \odot (\tau \odot \sigma),$$

s.t. for all $x^\sigma, x^\tau, x^\delta$ min. caus. comp., $a_{\sigma, \tau, \delta}((x^\delta \odot x^\tau) \odot x^\sigma) = x^\delta \odot (x^\tau \odot x^\sigma)$.

Proof. The two clauses serve as the definition of the action of $a_{\delta, \tau, \sigma}$ on configurations, via Proposition 6.2.22 and Lemma 6.2.24. Compatibility with display maps is from the description of the display map of composition in Proposition 6.2.22. To lift this to a map, we observe that preservation of cardinal is clear by local injectivity and preservation of display maps, and preservation of unions follows from Lemma 6.2.20. Finally, the same arguments holds for the inverse, therefore forming an isomorphism. \square

Composition of prestrategies is associative, but there does not seem to be an identity turning it into a (bi)category – the obvious candidate, the *copycat strategy* that we shall see later on, is only neutral for composition with respect to *strategies*.

6.3 Composition of Strategies

Now we focus on *strategies*, i.e. those prestrategies satisfying the additional conditions of Definition 6.1.6. In this section, we aim in particular to prove Proposition 6.2.1.

6.3.1 Composition of Strategies

Thus, we now turn to the composition of strategies. Our first objective is to establish that they are indeed stable under composition.

Proposition 6.3.1. *Consider $\sigma : A \vdash B$ and $\tau : B \vdash C$ strategies.*

Then, $\tau \odot \sigma : A \vdash C$ is a strategy.

Proof. By Proposition 6.2.22, we have $\tau \odot \sigma : A \vdash C$ a prestrategy.

Courteous. Let $p \rightarrow_{\tau \odot \sigma} p'$ s.t. $\text{pol}(p) = +$ or $\text{pol}(p') = -$. By Lemma 6.2.21,

$$p \rightarrow_{\tau \odot \sigma} q_1 \rightarrow_{\tau \odot \sigma} \dots \rightarrow_{\tau \odot \sigma} q_n \rightarrow_{\tau \odot \sigma} p'$$

for some $q_1, \dots, q_n \in \tau \otimes \sigma$ synchronized.

We show that in fact $n = 0$ and $p \rightarrow_{\tau \odot \sigma} p'$. If $\text{pol}(p) = +$, then $\text{pol}_{\tau \odot \sigma}(q_1) \in \{\ell, r\}$ – assume w.l.o.g. that it is ℓ . Then by Lemma 6.2.15, $p_\sigma, (q_1)_\sigma$ are defined and $p_\sigma \rightarrow_\sigma (q_1)_\sigma$. But by courtesy, this entails $\partial_\sigma(p_\sigma) \rightarrow_{A \vdash B} \partial_\sigma((q_1)_\sigma)$, impossible for q_1 synchronized. Hence $n = 0$, and $p \rightarrow_{\tau \odot \sigma} p'$. If $\text{pol}(p') = -$, then again by Lemma 6.2.15, $\partial_{\tau \odot \sigma}(q_n) \rightarrow_{A \parallel B \parallel C} \partial_{\tau \odot \sigma}(p')$, contradicting q_n synchronized, so $p \rightarrow_{\tau \odot \sigma} p'$.

Now, by Lemma 6.2.14, $p_\sigma \rightarrow_\sigma p'_\sigma$ or $p_\tau \rightarrow_\sigma p'_\tau$, say w.l.o.g. the former. But then $\partial_\sigma(p_\sigma) \rightarrow_{A \vdash B} \partial_\sigma(p'_\sigma)$, from which $\partial_{\tau \odot \sigma}(p) \rightarrow_{A \vdash C} \partial_{\tau \odot \sigma}(p')$ follows.

Receptive. Consider $x^\tau \odot x^\sigma \in \mathcal{C}(\tau \odot \sigma)$, and $\partial_{\tau \odot \sigma}(x^\tau \odot x^\sigma) = x_A^\sigma \parallel x_C^\tau \vdash_{A \vdash C} c^-$, assuming w.l.o.g. it is in C . For *existence*, by receptivity of τ we have $x^\tau \vdash_\tau t$ such that $\partial_\tau(t) = c$. Writing $y^\tau = x^\tau \uplus \{t\}$, it is direct that x^σ, y^τ is minimal causally compatible, and $x^\tau \odot x^\sigma \prec y^\tau \odot x^\sigma$ provides the required extension, via Proposition 6.2.22. For *uniqueness*, assume $x^\tau \odot x^\sigma \prec z^\tau \odot z^\sigma$ via some event p such that $\partial_{\tau \odot \sigma}(p) = c^-$. So,

$$x^\tau \otimes x^\sigma \subseteq u^\tau \otimes u^\sigma \xrightarrow{p} z^\tau \otimes z^\sigma$$

where $x^\tau \otimes x^\sigma \subseteq u^\tau \otimes u^\sigma$ adds a set X of synchronized events, such that for all $q \in X$, $q \leq_{\tau \odot \sigma} p$. But for any $q \rightarrow_{\tau \odot \sigma} p$, we have q in C by Lemma 6.2.15, so X is empty. As p is in C , this entails $x^\sigma = u^\sigma = z^\sigma$ and $x^\tau = u^\tau \prec z^\tau$. So $z^\tau = x^\tau \uplus \{t'\}$ such that $x^\tau \vdash_\tau t'$ and $\partial_\tau(t') = c$, so $t = t'$ by *receptive* for τ ; hence $y^\tau \odot x^\sigma = z^\tau \odot z^\sigma$. \square

We have a well-defined composition operation for *strategies*, but we have not yet arrived at the claimed Proposition 6.2.1. For that, it still remains to investigate the effect of ignoring the trailing Opponent moves, restricting to *+covered configurations*.

6.3.2 Strategies and +-covered configurations

In traditional alternating game semantics, strategies can be defined in two different ways: as done in Definition 3.1.2 by taking plays finishing by moves of arbitrary polarity, but requiring a *receptivity* property for strategies; or – as is more standard – by only considering even-length plays in strategies, that is, plays ending with Player moves. For concurrent strategies, a similar effect can be obtained via *+-covered configurations*:

Definition 6.3.2. Consider $\sigma : A$ a (pre)strategy on game A . A configuration $x \in \mathcal{C}(\sigma)$ is **+-covered** if for any $s \in x$ such that s is \leq_σ -maximal in x , $\text{pol}_\sigma(s) = +$.

We write $\mathcal{C}^+(\sigma)$ for the +-maximal configurations of σ .

We shall see that it suffices to compare strategies on their +-covered configurations. To show this, it is convenient to rely not on the original definition of strategies in Definition 6.1.6, but on the following more “big-step” characterization:

Proposition 6.3.3. Consider $\sigma : A$ a prestrategy on game A . Then it is a strategy iff

- discrete opfibration: for all $x^\sigma \in \mathcal{C}(\sigma)$, for all $\partial_\sigma(x^\sigma) \subseteq^- y_A \in \mathcal{C}(A)$,
there exists a unique $x^\sigma \subseteq y^\sigma \in \mathcal{C}(\sigma)$ s.t. $\partial_\sigma(y^\sigma) = y_A$.
- +discrete fibration: for all $x^\sigma \in \mathcal{C}(\sigma)$, for all $\partial_\sigma(x^\sigma) \supseteq^+ y_A \in \mathcal{C}(A)$,
there exists a unique $x^\sigma \supseteq y^\sigma \in \mathcal{C}(\sigma)$ s.t. $\partial_\sigma(y^\sigma) = y_A$.

where $x \subseteq^p y$ (for $p \in \{-, +\}$) means $x \subseteq y$ and $\text{pol}_A(y \setminus x) \subseteq \{p\}$.

Proof. --discrete opfibration. Existence is clear by repeated applications of *receptive*. For uniqueness, consider $x^\sigma \subseteq y^\sigma, z^\sigma \in \mathcal{C}(\sigma)$ s.t. $\partial_\sigma(y^\sigma) = \partial_\sigma(z^\sigma) = y_A$. Take $m_1 \in y^\sigma \setminus x^\sigma$ and $m_2 \in z^\sigma \setminus x^\sigma$ s.t. $\partial_\sigma(m_1) = \partial_\sigma(m_2)$, w.l.o.g. assume that for any $m'_1 <_\sigma m_1$ and $m'_2 <_\sigma m_2$, $\partial_\sigma(m'_1) = \partial_\sigma(m'_2)$ implies $m'_1 = m'_2$. By *courteous*, m_1 and m_2 have the same immediate causal dependencies, so $m_1 = m_2$ by *receptive*.

+discrete fibration. Existence follows by repeated applications of *courtesy*, which entails that events in $x^\sigma \in \mathcal{C}(\sigma)$ displaying to maximal positive events in $\partial_\sigma(x^\sigma)$ must be maximal in x^σ . Uniqueness is clear by *locally injective*.

In the sequel, we shall only use this direction. So for the sake of economy, for the other direction we simply refer to [Castellan et al., 2017a, Lemma 3.13]. \square

This was first noticed by Winskel in [Winskel, 2013b], in seeking a connection between concurrent strategies as profunctors – indeed, by the lemma above, it also follows that strategies are also characterized by the discrete fibration property for the composite partial order $\sqsubseteq_A = \supseteq^- \subseteq^+$, the so-called “Scott order”, which informs a link with profunctors. In this monograph, we shall merely use the --discrete opfibration property. Relying on that, we refine the mapification lemma in the following way⁴:

Lemma 6.3.4. Consider $\sigma, \tau : A$ two strategies. Assume there is a function

$$f : \mathcal{C}^+(\sigma) \rightarrow \mathcal{C}^+(\tau)$$

compatible with display maps and preserving unions. Then, there is a unique morphism of strategies $\hat{f} : \sigma \Rightarrow \tau$ such that for all $x \in \mathcal{C}^+(\sigma)$, $\hat{f} x = f x$.

⁴Thanks to Victor Blanchi for discovering this nice lemma!

Proof. Uniqueness. Consider g, h two such extensions, and consider $x \in \mathcal{C}(\sigma)$. Removing trailing events from x , there is a unique $y \in \mathcal{C}^+(\sigma)$ such that $y \subseteq^- x$. By hypothesis, $g(y) = h(y) = f(y)$. Since g, h are maps of event structures and preserve display maps, they must preserve inclusion and polarities of these inclusions, so that

$$f(y) \subseteq^- g(x), h(x)$$

but $\partial_\tau g(x) = \partial_\tau h(x)$ by pres. of display maps, so $g(x) = h(x)$ by *--discrete opfibration* of Proposition 6.3.3. So $g(x) = h(x)$ for all $x \in \mathcal{C}(\sigma)$, so $g = h$ by Lemma 6.1.12.

Existence. We first extend f to all configurations: if $x \in \mathcal{C}(\sigma)$, as above we consider the unique $y \in \mathcal{C}^+(\sigma)$ such that $y \subseteq^- x$ obtained by removing trailing negative events in x ; as f preserves display maps we have $\partial_\tau f(y) = \partial_\sigma y \subseteq^- \partial_\sigma x$. By *--discrete opfibration* of Proposition 6.3.3, there is a unique $z \in \mathcal{C}(\tau)$ such that $f(y) \subseteq^- z$ and $\partial_\tau(z) = \partial_\sigma(x)$; we set $f(x) = z$. So defined, $f : \mathcal{C}(\sigma) \rightarrow \mathcal{C}(\tau)$ automatically preserves display maps and hence preserves cardinal (as ∂_σ and ∂_τ , as maps of event structures, are locally injective). By uniqueness in *--discrete opfibration* of Proposition 6.3.3, it is easy to prove that f preserves unions. Hence by Lemma 6.1.11, it extends to a map of event structures in a unique way, concluding the proof.

Rigidity. We prove that any map of event structures $h : \sigma \rightarrow \tau$ such that $\partial_\tau \circ h = \partial_\sigma$ and that sends $+$ -covered configurations to $+$ -covered configurations, is rigid. We actually show that if $s_1 \rightarrow_\sigma s_2$, then $h(s_1) <_\sigma h(s_2)$, which entails rigidity. If $\text{pol}(s_1) = +$ or $\text{pol}(s_2) = -$, then $\partial_\sigma s_1 \rightarrow_A \partial_\sigma s_2$ by *courteous*, and thus $f(s_1) \leq_\tau f(s_2)$ by Lemma 6.1.8 – so that $f(s_1) <_\tau f(s_2)$ by local injectivity. It remains to consider the case where $\text{pol}(s_1) = -$ and $\text{pol}(s_2) = +$. In that case, as $[s_2]_\sigma \in \mathcal{C}^+(\sigma)$, we have $h([s_2]_\sigma) \in \mathcal{C}^+(\tau)$ as well. Now, observe that s_1 is maximal in $[s_2]_\sigma$, otherwise contradicting $s_1 \rightarrow_\sigma s_2$. Hence writing $x = [s_2]_\sigma \setminus \{s_1, s_2\}$, $y = [s_2]_\sigma \setminus \{s_2\}$, we have $x \subset y \subset [s_2]_\sigma$ in $\mathcal{C}(\sigma)$. This is sent to τ via h to the covering sequence in $\mathcal{C}(\tau)$

$$h(x) \xrightarrow{h(s_1)} h(y) \xrightarrow{h(s_2)} h([s_2]_\sigma),$$

which shows in particular that only $h(s_2)$ can possibly depend on $h(s_1)$ within $h([s_2]_\sigma)$. Hence if $\neg(h(s_1) <_\tau h(s_2))$, it immediately follows that $h(s_1)$ is maximal in $h([s_2]_\sigma)$. But $h(s_1)$ is negative, contradicting that $h([s_2]_\sigma)$ is $+$ -covered; thus $h(s_1) <_\sigma h(s_2)$. \square

It is striking that the *rigidity* condition on morphisms is automatic, and follows from the fact that the map of event structures preserves $+$ -covered configurations!

This lemma will be used often, in two particular situations: firstly, to show that two morphisms of strategies coincide, it suffices to show that they coincide on $+$ -covered configurations; secondly, in order to prove σ and τ isomorphic, it suffices to produce an order-iso between $\mathcal{C}^+(\sigma)$ and $\mathcal{C}^+(\tau)$ compatible with display maps (observing that the supremum of two $+$ -covered configurations in $\mathcal{C}^+(\sigma)$, if it exists, is their union).

6.3.3 Composition and $+$ -coveredness

Clearly $+$ -covered configurations are important, in that morphisms between strategies are entirely determined by their action on $+$ -covered configurations. Next we show that

additionally, $+$ -covered configurations interact very well with composition. Let us fix for now $\sigma : A \vdash B$ and $\tau : B \vdash C$ two strategies, and investigate their composition.

The first step is to remark that $+$ -coverdness makes *minimality* redundant:

Lemma 6.3.5. *Consider $x^\sigma \in \mathcal{C}^+(\sigma)$ and $x^\tau \in \mathcal{C}^+(\tau)$ causally compatible. Then, x^σ and x^τ are minimal causally compatible.*

Proof. Seeking a contradiction, assume there is a maximal event of $\varphi[x^\sigma, x^\tau]$ in B , necessarily of the form $((1, s), (2, t))$ for $s \in x^\sigma$ and $t \in x^\tau$ necessarily both maximal. But necessarily one of them is negative, contradicting $+$ -coveredness of x^σ and x^τ . \square

Besides, the resulting configurations of the composition are automatically $+$ -covered:

Lemma 6.3.6. *Consider $x^\sigma \in \mathcal{C}^+(\sigma)$ and $x^\tau \in \mathcal{C}^+(\tau)$ causally compatible. Then, $x^\tau \odot x^\sigma \in \mathcal{C}^+(\tau \odot \sigma)$.*

Proof. First, x^σ, x^τ is minimal causally compatible by Lemma 6.3.5.

Consider $m \in x^\tau \odot x^\sigma$ maximal. This means that m is also maximal in $x^\tau \otimes x^\sigma = [x^\tau \odot x^\sigma]_{\tau \otimes \sigma}$, so corresponds to some pair maximal in $\varphi[x^\sigma, x^\tau]$ via Lemma 6.2.12. Since x^σ, x^τ is minimal causally compatible, this pair is in A or C – say *w.l.o.g.* in C , so it has the form $((2, c), (2, t))$ for $t \in x^\tau$. Necessarily, t is maximal in x^τ , so is positive since $x^\sigma \in \mathcal{C}^+(\tau)$. 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 6.3.7. *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 $p \in x^\tau \otimes x^\sigma$ such that $p_\sigma = s$. If p is maximal in $x^\tau \otimes x^\sigma$, by Lemma 6.2.19 we have $x^\tau \otimes x^\sigma = [x^\tau \odot x^\sigma]_{\tau \otimes \sigma}$, so p must be visible and maximal in $x^\tau \odot x^\sigma$, hence positive by hypothesis, and s is positive.

Otherwise, there is some $p \rightarrow_{\tau \otimes \sigma} q$. By Lemma 6.2.14, $p_\sigma \rightarrow_\sigma q_\sigma$ or $p_\tau \rightarrow_\tau q_\tau$. The former contradicts maximality of $p_\sigma = s$, so $p_\tau \rightarrow_\tau q_\tau$. If s is positive, we are done. Otherwise, p_τ is positive, so q_τ is also in B by *courteous*, and $\partial_\tau(p_\tau) <_{B \vdash C} \partial_\tau(q_\tau)$. Then q_σ is also defined, and by Lemma 6.1.8, $p_\sigma < q_\sigma$, contradicting maximality of s . The symmetric reasoning shows any $t \in x^\tau$ maximal in x^τ is positive. \square

We can now finally deduce our characterization of the composition of strategies:


Proposition 6.2.1. *Consider A, B, C games, and $\sigma : A \vdash B$ and $\tau : B \vdash C$ strategies. Then there is a strategy $\tau \odot \sigma : A \vdash C$, unique up to iso, s.t. there are order-isos:*

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

such that for $x^\sigma \in \mathcal{C}^+(\sigma)$ and $x^\tau \in \mathcal{C}^+(\tau)$ causally compatible,

$$\partial_{\tau \odot \sigma}(x^\tau \odot x^\sigma) = x_A^\sigma \vdash x_C^\tau$$

where we write $\partial_\sigma x^\sigma = x_A^\sigma \vdash x_B^\sigma$ and $\partial_\tau x^\tau = x_B^\tau \vdash x_C^\tau$.

Figure 6.6: The copycat strategy on $(U \multimap U) \multimap U$

Proof. Existence. The composition exists by Proposition 6.2.22, and is a strategy by Proposition 6.3.1. The isos are restrictions of those of Proposition 6.2.22. By Lemma 6.3.5, causally compatible $x^\sigma \in \mathcal{C}^+(\sigma)$ and $x^\tau \in \mathcal{C}^+(\tau)$ are automatically minimal, and by Lemma 6.3.6, $x^\tau \odot x^\sigma \in \mathcal{C}^+(\tau \odot \sigma)$ is $+-$ -covered. Reciprocally, if $x^\tau \odot x^\sigma \in \mathcal{C}^+(\tau \odot \sigma)$ is $+-$ -covered, then by Lemma 6.3.7, so are $x^\sigma \in \mathcal{C}^+(\sigma)$ and $x^\tau \in \mathcal{C}^+(\tau)$.

Uniqueness. By Lemma 6.3.4. \square

This concludes our study of composition of prestrategies and strategies. In the next section, we show that games and strategies may be organized into a bicategory **CG**.

6.4 The Bicategory **CG**

The bicategory will have: as *objects*, games, as *morphisms* from A to B , the strategies on $A \vdash B$, and as *2-cells*, the morphisms of strategies. We have already seen above the composition of strategies and the *unitors*, 2-cells witnessing associativity of composition. It remains to construct *identities*, 2-cells witness identity laws, definition of horizontal composition of 2-cells; after which we must establish the coherence laws.

6.4.1 Copycat

First we define *copycat strategies*, which will provide identities.

Definition. Copycat, illustrated in Figure 6.6, is an *asynchronous forwarder*: it sends each negative event on $A \vdash A$ to the matching positive event on the other side.

We start with the components of copycat, postponing the verification of conditions.

Definition 6.4.1. For any game A , we set $\mathfrak{C}_A : A \vdash A$ with components:

$$\begin{aligned} \text{events: } |\mathfrak{C}_A| &= |A| + |A| \\ \text{causality: } \leq_{\mathfrak{C}_A} &= (\{(1, a), (1, a') \mid a \leq_A a'\} \uplus \\ &\quad \{(2, a), (2, a') \mid a \leq_A a'\} \uplus \\ &\quad \{(1, a), (2, a) \mid \text{pol}_A(a) = +\} \uplus \\ &\quad \{(2, a), (1, a) \mid \text{pol}_A(a) = -\})^* \\ \text{conflict: } \#_{\mathfrak{C}_A} &= \{(e_1, e_2) \mid \exists e'_1 \leq_{\mathfrak{C}_A} e_2, e'_2 \leq_{\mathfrak{C}_A} e_2, e'_1 \#_{A \vdash A} e'_2\} \end{aligned}$$

where $(-)^*$ is the transitive closure; and $\partial_{\mathfrak{C}_A} : |\mathfrak{C}_A| \rightarrow |A \vdash A|$ the identity function.

Copycat is often a challenging landmark in setting up a game semantics, and here it is no exception: it is not obvious at all that it satisfies all the required conditions! We shall slowly examine various aspects of this definition, for a fixed game A .

The event structure \mathfrak{C}_A . We first focus on the basic structure of \mathfrak{C}_A .

Lemma 6.4.2. Consider A a game. Then, $(|\mathfrak{C}_A|, \leq_{\mathfrak{C}_A}, \#_{\mathfrak{C}_A})$ is an event structure.

Moreover, $\mathcal{C}(\mathfrak{C}_A) = \{x \parallel y \mid x, y \in \mathcal{C}(A), x \cap y \subseteq^+ x \ \& \ x \cap y \subseteq^- y\}$.

Proof. Immediate verifications, omitted. \square

It will be useful to understand exactly what are the immediate causal links of \mathfrak{C}_A :

Lemma 6.4.3. Consider A a game, and $(i, a), (j, a') \in \mathfrak{C}_A$. Then $(i, a) \rightarrow_{\mathfrak{C}_A} (j, a')$ iff

- (1) $i = j, a \rightarrow_A a'$, and $(\text{pol}_{\mathfrak{C}_A}(i, a) = + \text{ or } \text{pol}_{\mathfrak{C}_A}(j, a') = -)$, or
- (2) $i \neq j, a = a', \text{pol}_{\mathfrak{C}_A}(i, a) = - \text{ and } \text{pol}_{\mathfrak{C}_A}(j, a') = +$.

This is a direct verification. We also include a direct consequence of this:

Lemma 6.4.4. Consider A a game. Then, $\mathcal{C}^+(\mathfrak{C}_A) = \{x \parallel x \mid x \in \mathcal{C}(A)\}$.

Proof. \supseteq . Clear by Lemma 6.4.2.

\subseteq . Consider $x \parallel y \in \mathcal{C}^+(\mathfrak{C}_A)$; by Lemma 6.4.2, we have $x, y \in \mathcal{C}(A)$ with $x \cap y \subseteq^+ x$ and $x \cap y \subseteq^- y$. Suppose one of these inclusions is strict; for instance with $a \in y$ such that $a \notin x$ (the other case is symmetric). Necessarily, $\text{pol}_A(a) = -$. If a is maximal in y , then $(2, a)$ is maximal in $x \parallel y$ by Lemma 6.4.3 since $a \notin x$. But $\text{pol}_{A \vdash A}(2, a) = -$, contradicting that $x \parallel y$ is $+$ -covered. Otherwise, there is $a' \in y$ such that $a \rightarrow_A a'$. Now if $\text{pol}_A(a') = +$, we must have $a' \in x$, so $a \in x$ as $x \in \mathcal{C}(A)$, contradiction. So, $\text{pol}_A(a') = -$. By iterating this reasoning, we find $a'' \in y \setminus x$ negative and maximal for A ; and by the reasoning above $(2, a'')$ is negative maximal in $x \parallel y$, contradiction. \square

This lemma above will play an important role in this monograph: it shows that when restricting to $+$ -covered configurations (which is always possible when comparing strategies, by Lemma 6.3.4), copycat behaves as the identity relation. For $x \in \mathcal{C}(A)$, we write $\mathfrak{C}_x = x \vdash x \in \mathcal{C}^+(\mathfrak{C}_A)$ for the corresponding $+$ -covered configuration.

Next, we show that copycat indeed is a well-formed strategy.

Proposition 6.4.5. *For any game A , we have $\mathfrak{c}_A : A \vdash A$ a strategy.*

Proof. By Lemma 6.4.2, \mathfrak{c}_A is an event structure. It is direct that $\partial_{\mathfrak{c}_A}$ is a map of event structures. *Courteous* follows from Lemma 6.4.3, and *receptive* from Lemma 6.4.2. \square

In the sequel, we shall mostly reason on copycat via the characterization of its +-covered configurations in Lemma 6.4.4, showing that it acts as an identity relation.

6.4.2 Copycat and Strategies

Now, we show that copycat indeed is neutral for composition with respect to strategies.

The unitors. To define the unitor isomorphisms, from our existing tools, the only thing left to prove is that synchronizations with copycat never induce any deadlock. For this – and for later on in this monograph – we introduce the following lemma:

Lemma 6.4.6. *Take $\sigma : A \vdash B$, $\tau : B \vdash C$, and $x^\sigma \in \mathcal{C}(\sigma)$, $x^\tau \in \mathcal{C}(\tau)$ matching.*

Consider the relation $\triangleleft_\sigma \cup \triangleleft_\tau$ on $x_A \parallel x_B \parallel x_C$ defined with

$$\begin{aligned} (\partial_\sigma \parallel C)(m) &\triangleleft_\sigma (\partial_\sigma \parallel C)(m') \quad \text{where} \quad m <_{\sigma \parallel C} m' \quad (m, m' \in x^\sigma \parallel x_C), \\ (A \parallel \partial_\tau)(m) &\triangleleft_\tau (A \parallel \partial_\tau)(m') \quad \text{where} \quad m <_{A \parallel \tau} m' \quad (m, m' \in x_A \parallel x^\tau) \end{aligned}$$

writing $\partial_\sigma(x^\sigma) = x_A \parallel x_B$ and $\partial_\tau(x^\tau) = x_B \parallel x_C$.

Then, x^σ, x^τ are causally compatible iff $\triangleleft_\sigma \cup \triangleleft_\tau$ has no cycle within x_B .

Proof. If. First, notice that \triangleleft has no direct link between x_A and x_C . The crux of the argument is that if $m \triangleleft m'$ with one of them in A , then $m \triangleleft_\sigma m'$: indeed if they are both in A , then by definition either $m \triangleleft_\sigma m'$ and we are done, or $m <_{A \parallel B \parallel C} m'$ and $m \triangleleft_\sigma m'$ follows from Lemma 6.1.8. Likewise, if $m \triangleleft m'$ with one of them in C , then $m \triangleleft_\tau m'$.

We proceed with the proof. First, it is a simple reformulation of the definition that x^σ, x^τ are causally compatible iff $\triangleleft = \triangleleft_\sigma \cup \triangleleft_\tau$ is acyclic. If there is a cycle in \triangleleft , it must pass through B : otherwise it still entirely in A or C . But by the observation above, this yields a cycle in \triangleleft_σ or \triangleleft_τ , contradiction. So the cycle passes through B . Then, for

$$m^B \triangleleft n_1^A \triangleleft \dots \triangleleft n_k^A \triangleleft (m')^B,$$

again by the observation above all those links are in \triangleleft_σ , hence $m^B \triangleleft_\sigma (m')^B$. Likewise, all sections of the cycle in C may be removed symmetrically, yielding a cycle in x_B . \square

We now focus on the interactions of a strategy with copycat. In particular, the core of the next lemma is that no deadlocks can arise from the interaction of a strategy with copycat, so that all matching pairs of +-covered configurations are causally compatible.

Lemma 6.4.7. *Consider $\sigma : A \vdash B$ a strategy, and $x^\sigma \in \mathcal{C}^+(\sigma)$, $x^{\mathfrak{c}_B} \in \mathcal{C}^+(\mathfrak{c}_B)$.*

Then, $x^\sigma, x^{\mathfrak{c}_B}$ are causally compatible iff $x^{\mathfrak{c}_B} = \mathfrak{c}_{x_B^\sigma}$, for $\partial_\sigma(x^\sigma) = x_A^\sigma \parallel x_B^\sigma$.

Proof. If. Write $\partial_\sigma(x^\sigma) = x_A \parallel x_B$, and seeking a contradiction, consider a cycle in $x_A \parallel x_B^\ell \parallel x_B^r$ – using ℓ, r for disambiguation – in the relation $\triangleleft_\sigma \cup \triangleleft_{\mathfrak{C}_B}$ defined as:

$$\begin{aligned} (\partial_\sigma \parallel B)(m) &\triangleleft_\sigma (\partial_\sigma \parallel B)(m') \quad \text{where} \quad m <_{\sigma \parallel B} m' \quad (m, m' \in x^\sigma \parallel x_B^r) \\ (A \parallel \partial_{\mathfrak{C}_B})(m) &\triangleleft_{\mathfrak{C}_B} (A \parallel \partial_{\mathfrak{C}_B})(m') \quad \text{where} \quad m <_{A \parallel \mathfrak{C}_B} m' \quad (m, m' \in x_A \parallel x_B^\ell \parallel x_B^r), \end{aligned}$$

by Lemma 6.4.6 we can assume that the cycle is entirely in x_B^ℓ . Now, it is straightforward from the definition of \mathfrak{C}_B that if $m \triangleleft_{\mathfrak{C}_B} m'$ with m, m' in x_B^ℓ , then $m <_{A \parallel B \parallel B} m'$ so that $m \triangleleft_\sigma m'$ by Lemma 6.1.8. Thus, this yields a cycle in \triangleleft_σ , contradiction.

Only if. Immediate from the definition and Lemma 6.4.4. \square

Clearly, the symmetric lemma holds, with a symmetric proof:

Lemma 6.4.8. *Consider $\sigma : A \vdash B$ a strategy, and $x^{\mathfrak{C}_A} \in \mathcal{C}^+(\mathfrak{C}_A)$, $x^\sigma \in \mathcal{C}^+(\sigma)$. Then, $x^{\mathfrak{C}_A}, x^\sigma$ are causally compatible iff $x^{\mathfrak{C}_A} = \mathfrak{C}_{x_A^\sigma}$, for $\partial_\sigma(x^\sigma) = x_A^\sigma \parallel x_B^\sigma$.*

This completes the ingredients needed to define the *unitors*:

Proposition 6.4.9. *Consider $\sigma : A \vdash B$ a strategy.*

*Then there are unique isomorphisms of strategies, the **unitors**:*

$$l_\sigma : \mathfrak{C}_B \odot \sigma \cong \sigma \quad r_\sigma : \sigma \odot \mathfrak{C}_A \cong \sigma$$

such that for all $x^\sigma \in \mathcal{C}^+(\sigma)$, we have $l_\sigma(\mathfrak{C}_{x_B^\sigma} \odot x^\sigma) = x^\sigma$ and $r_\sigma(x^\sigma \odot \mathfrak{C}_{x_A^\sigma}) = x^\sigma$.

Proof. Configurations and symmetries of $\mathfrak{C}_B \odot \sigma$ and $\sigma \odot \mathfrak{C}_A$ have this shape by Lemmas 6.4.7 and 6.4.8. These assignments on configurations are order-isos and compatible with display maps, thus extend to an isomorphism of strategies by Lemma 6.3.4. \square

6.4.3 Horizontal Composition and Bicategorical Structure

The missing piece is the horizontal composition of 2-cells, handled by:

Proposition 6.4.10. *Consider $\sigma, \sigma' : A \vdash B$ and $\tau, \tau' : B \vdash C$ strategies, and $f : \sigma \Rightarrow \sigma', g : \tau \Rightarrow \tau'$ morphisms of strategies.*

*There is a unique morphism of strategies, the **horizontal composition** of f and g ,*

$$g \odot f : \tau \odot \sigma \Rightarrow \tau' \odot \sigma',$$

such that $(g \odot f)(x^\tau \odot x^\sigma) = g(x^\tau) \odot f(x^\sigma)$ for all $x^\tau \odot x^\sigma \in \mathcal{C}^+(\tau \odot \sigma)$.

Proof. Uniqueness. Obvious from Lemma 6.3.4.

Existence. Consider $x^\tau \odot x^\sigma \in \mathcal{C}^+(\tau \odot \sigma)$. As f and g preserve display maps, $f(x^\sigma)$ and $g(x^\tau)$ are matching. As f and g are rigid, and by Lemma 6.1.8, they induce order-isos $f : x^\sigma \cong f(x^\sigma)$ and $g : x^\tau \cong g(x^\tau)$, so $f(x^\sigma)$ and $g(x^\tau)$ are still $+$ -covered, and causally compatible. So the action on $+$ -covered configurations $x^\tau \odot x^\sigma \in \mathcal{C}^+(\tau \odot \sigma)$ given above is well-defined. It obviously preserves display maps and preserves unions by Lemma 6.2.20, hence it lifts to a morphism of strategies by Lemma 6.2.20. \square

We are now in position to conclude with the bicategorical structure:

Theorem 6.4.11. *There is **CG**, a bicategory with: objects, all games, morphisms from A to B , strategies on $A \vdash B$, and 2-cells, morphisms of strategies.*

Proof. Identities are copycat strategies, composition is composition of strategies. Associators are defined in Section 6.2.3, unitors in Section 6.4.2. The laws are obtained by checking them pointwise on $+$ -covered configurations and concluding by Lemma 6.3.4; for instance, naturality of the left unitor is established very simply, via:

$$\begin{aligned} l_\sigma \circ (\text{id} \odot f)(\mathfrak{C}_{x_B}^\sigma \odot x^\sigma) &= l_\sigma(\mathfrak{C}_{x_B}^\sigma \odot f(x^\sigma)) \\ &= f(x^\sigma) \\ &= f \circ l_\sigma(\mathfrak{C}_{x_B}^\sigma \odot x^\sigma) \end{aligned}$$

for $f : \sigma \Rightarrow \sigma' : A \vdash B$ a morphism of strategies, and $x^\sigma \in \mathcal{C}^+(\sigma)$. \square

There is, of course, much more to the structure of **CG**. In particular it is a compact closed bicategory, and may be used as a model for a linear programming language. But we shall not investigate this structure further, and will do it instead for **TCG**.

6.4.4 Characterization of Strategies

Before adding symmetry into the picture, we prove an additional property on copycat. Namely, it allows a neat characterization of *strategies* among prestrategies:

Proposition 6.4.12. *Consider $\sigma : A \vdash B$ a prestrategy.*

Then, σ is a strategy iff $\mathfrak{C}_B \odot \sigma \odot \mathfrak{C}_A \cong \sigma$.

Proof. If. It suffices to prove that $\mathfrak{C}_B \odot \sigma \odot \mathfrak{C}_A$ is automatically *receptive* and *courteous*.

For *courteous*, consider $p \rightarrow_{\mathfrak{C}_B \odot \sigma \odot \mathfrak{C}_A} p'$ such that $\text{pol}(p) = +$ or $\text{pol}(p') = -$, say first it is the former. Necessarily p occurs in A or B , say *w.l.o.g.* that it is in B . Up to iso we consider $p \rightarrow_{\mathfrak{C}_B \odot \sigma'} p'$, for $\sigma' = \sigma \odot \mathfrak{C}_A$. By Lemma 6.2.21 we have

$$p \rightarrow_{\mathfrak{C}_B \odot \sigma'} q_1 \rightarrow_{\mathfrak{C}_B \odot \sigma'} \dots \rightarrow_{\mathfrak{C}_B \odot \sigma'} q_n \rightarrow_{\mathfrak{C}_B \odot \sigma'} p'$$

with q_1, \dots, q_n synchronized. By Lemma 6.2.14, we must have $p_{\mathfrak{C}_B} \rightarrow_{\mathfrak{C}_B} (q_1)_{\mathfrak{C}_B}$ – as p_σ does not exist. But as p is positive, this implies $\partial_{\mathfrak{C}_B}(p_{\mathfrak{C}_B}) \rightarrow_{B \vdash B} \partial_{\mathfrak{C}_B}((q_1)_{\mathfrak{C}_B})$, contradicting q_1 synchronized, so $p \rightarrow_{\mathfrak{C}_B \odot \sigma'} p'$ and $p_{\mathfrak{C}_B} \rightarrow_{\mathfrak{C}_B} p'_{\mathfrak{C}_B}$, which entails $\partial_{\mathfrak{C}_B}(p_{\mathfrak{C}_B}) \rightarrow_{B \vdash B} \partial_{\mathfrak{C}_B}(p'_{\mathfrak{C}_B})$ so that $\partial_{\mathfrak{C}_B \odot \sigma'}(p) \rightarrow_{A \vdash B} \partial_{\mathfrak{C}_B \odot \sigma'}(p')$ as required. All the other cases are symmetric, so that $\partial_{\mathfrak{C}_B \odot \sigma \odot \mathfrak{C}_A}(p) \rightarrow_{A \vdash B} \partial_{\mathfrak{C}_B \odot \sigma \odot \mathfrak{C}_A}(p')$ as required.

For *receptive*, consider $z \in \mathcal{C}(\mathfrak{C}_B \odot \sigma \odot \mathfrak{C}_A)$ with $\partial_{\mathfrak{C}_B \odot \sigma \odot \mathfrak{C}_A}(z) \vdash_{A \vdash B} b^-$, *w.l.o.g.* assume b occurs in B , and consider again z , which by Proposition 6.2.22 we may write as $z = (x_B \parallel y_B) \odot x^{\sigma'} \in \mathcal{C}(\mathfrak{C}_B \odot \sigma')$ with $\sigma' = \sigma \odot \mathfrak{C}_A$, with $\partial_{\mathfrak{C}_B \odot \sigma'}(z) = x_A^{\sigma'} \parallel y_B$. So by hypothesis, we have $y_B \vdash_B b^-$. Then it is immediate that $x_B \parallel (y_B \cup \{b\})$ are still minimal causally compatible, so that by Proposition 6.2.22 we have

$$(x_B \parallel y_B) \odot x^{\sigma'} \xrightarrow{p} (x_B \parallel (y_B \cup \{b\})) \odot x^{\sigma'}$$

for some $(x_B \parallel y_B) \odot x^{\sigma'} \vdash_{\mathfrak{A}_B \odot \sigma'} p$ such that $\partial_{\mathfrak{A}_B \odot \sigma'}(p) = b$. Uniqueness is straightforward via Proposition 6.2.22 and definition of $\partial_{\mathfrak{A}_B \odot \sigma'}$.

Receptive and *courteous* are clearly invariant under strong iso, so σ is a strategy.

Only if. By bicategorical laws for **CG**. \square

Thus strategies are exactly those prestrategies invariant under composition with copycat, *i.e.* robust when accessed under asynchronous delay. Conceptually, this vindicates the definition of strategies. Technically, it occasionally provides a convenient way to prove that certain prestrategies are actually strategies. This result appeared as the main theorem in Rideau and Winskel’s original paper on concurrent strategies [Rideau and Winskel, 2011]. A result in the same spirit already appeared in Ghica and Murawski’s interleaving concurrent games model [Ghica and Murawski, 2008].

6.5 Conclusions and Historical Notes

Concurrent games were initiated by Abramsky and Melliès’ seminal fully complete model for MALL [Abramsky and Melliès, 1999]. There, strategies are certain stable closure operators on dI-domains: intuitively, a strategy-as-closure-operator closes the current position under all available Player moves, played in some unspecified order. Concurrent games were originally presented as a way to circumvent the so-called *Blass problem*, *i.e.* the non-associativity of composition in Blass’ model of linear logic, diagnosed in [Abramsky, 2003] as due to an excess of sequentiality.

Asynchronous Games. Following this, Melliès led a research programme on *asynchronous games* [Melliès, 2003, Melliès, 2004a, Melliès, 2004b, Melliès, 2005]. Melliès’ goal was to reproduce and extend the success of concurrent games in a language closer to more standard (sequential) game semantics; as well as put the spotlight on the *causal structures* that lurk behind the traditional notion of *innocence*. Asynchronous games are based on Mazurkiewicz traces: plays are *alternating* but related by *homotopy tiles* expressing causal independence between parts of computation. Plays up to homotopy give rise to a notion of *positions* compatible with the positions of the original concurrent games. Based on alternating asynchronous games, Melliès was able to provide a fully complete model of full propositional linear logic [Melliès, 2005], showing that what made this possible already for closure operators was not in fact concurrency, but *positionality* – see [Clairambault, 2023] for a recent retelling of this story.

Later, Melliès and Mimram developed basic *non-alternating* asynchronous games [Melliès and Mimram, 2007]. While non-alternating asynchronous strategies are non-deterministic as sets of plays, they satisfy conditions ensuring that they are essentially deterministic up to the choice of the scheduler – in particular, it follows a posteriori from the conditions that the set of possible executions is the set of linearizations of a *causality* partial order: provided the external Opponent plays in the same way, all execution paths will ultimately reach the same state.

Strategies as Partial Orders. Independently, a notion of *strategies-as-partial-orders* emerged from Girard’s Ludics [Girard, 2001]. Indeed, *designs* (strategies in Ludics) may be seen as a linear version of Hyland-Ong strategies [Faggian and Hyland, 2002], presented following their *causal* structure rather than the traditional chronologically ordered plays. This inspired *Ludics Nets*, or *L-nets* [Faggian and Maurel, 2005]: a notion of *strategies-as-graph*, aiming at a concurrent generalization of designs (see also [Curien and Faggian, 2005] for links with more standard strategies via proof nets). Inspired by this line of work and ideas from Hyland, Faggian and Piccolo introduced a category of *strategies-as-partial-orders* [Faggian and Piccolo, 2009], which lands close to Mimram and Melliès’ non-alternating asynchronous strategies – but whereas in Mimram and Melliès’ work the causal partial order is derived, here it is primitive.

Strategies as Event Structures. Rideau and Winskel then introduced the bicategory **CG** [Rideau and Winskel, 2011] that we developed in this chapter: as for Faggian and Piccolo, strategies in **CG** are explicitly partially ordered, but they additionally have a fibred structure letting them harmonously handle non-determinism with explicit branching information – this is analogous to other fibred approaches to non-determinism in game semantics [Eberhart et al., 2017, Jacq and Melliès, 2018]. At the same time, Winskel was starting a research project “ECSYM”, seeking a concurrent generalization of domain theory – the main inspirations were the non-deterministic extension of domain theory given by spans of event structures [Saunders-Evans and Winskel, 2006] and *event structures with symmetry* [Winskel, 2007]. But ECSYM soon focused on games and strategies and was the cradle for the first steps of the present developments.

From now on, unless specified otherwise and until the end of this document, we shall use the term *concurrent games* to refer specifically to this latter framework **CG** – keeping in mind all its precursors and the intellectual history behind it.

Chapter 7

Thin Concurrent Games

The main limitation of **CG** is that it is inherently restricted to speaking about systems with linear resource usage. Indeed, games as in Definition 6.1.3 do not support the *bang* construction from linear logic. Of course, one can define in **CG** a construction on games taking countably many copies as illustrated in (6.3), but this operation is incomplete: analogously to categorifications of the relational model (generalized species of structure [Fiore et al., 2008] or template games [Melliès, 2019b]) where *states* form a groupoid, we need to add additional structure to games, expressing which events are “morally the same” by specifying the admissible permutations between events.

Intuitively, the roadmap for building *concurrent games with symmetry* is to follow the same recipe as for **CG**, but replacing *event structures* (and their maps) with Winskel’s *event structures with symmetry* [Winskel, 2007] – which add to event structures precisely the ability to express these missing permutations. But while developing this, it turned out far more subtle than initially anticipated. With respect to this anticipated roadmap, the main surprise of **TCG** is a rediscovery of phenomenon first noticed by Melliès in the setting of asynchronous games: that a *causal, deterministic* account of uniformity required to split symmetries into two, the *positive* symmetries due to Player, and the *negative* symmetries due to Opponent [Melliès, 2003].

In this chapter, we will review a few of the subtleties and arising design choices. Then we shall focus on constructing *Thin Concurrent Games*, under the form of a bicategory **TCG**. We must hence ask the reader to brace for another rather technical journey. Like **CG**, the construction of **TCG** is not easy. But like **CG**, we did our best to engineer its construction in such a way that one can use it later on while leaving most combinatorial details and low-level definitions nicely encapsulated.

7.1 Symmetry in Games and Strategies

In this section, we introduce the basic definitions for games and strategies in the presence of symmetry, leaving for later the compositional structure. Before that, we start

with a few preliminaries on event structures with symmetry.

7.1.1 Event Structures with Symmetry

In order to motivate the definition of event structures with symmetry, we start with the key construction that *introduces* symmetry in games arising from the interpretation of types or linear logic formulas: the *bang* – making formal the illustration of (6.3).

Definition 7.1.1. Consider A a game. Then, we define the **bang** $!A$ with:

$$\begin{array}{llll} \text{events:} & |!A| & = & \mathbb{N} \times |A| \\ \text{causality:} & (i, a_1) \leq_{!A} (j, a_2) & \Leftrightarrow & i = j \wedge a_1 \leq_A a_2 \\ \text{conflict:} & (i, a_1) \#_{!A} (j, a_2) & \Leftrightarrow & i = j \wedge a_1 \#_A a_2 \\ \text{polarities:} & \text{pol}_{!A}(i, a) & = & \text{pol}_A(a). \end{array}$$

Following the methodology of linear logic, the construction $!A$ injects a non-linear behaviour into an otherwise linear setting – one may define the usual, non-linear arrow via $A \rightarrow B = !A \multimap B$. To achieve that, the definition above follows the intuition from AJM games (see Section 3.4) that $!A$ consists in countably many independent copies of the game A , kept apart via integers referred to as *copy indices*.

It will be useful to adopt specific notations for configurations of $!A$. First:

Definition 7.1.2. Consider A a game. We write $\text{Fam}(A)$ for the set of all families $(x_i)_{i \in I}$ where $I \subseteq_f \mathbb{N}$ is a finite set of integers, and $x_i \in \mathcal{C}(A)$ with $x_i \neq \emptyset$ for all $i \in I$.

We call elements of $\text{Fam}(A)$ simply **families** on A .

The set $\text{Fam}(A)$ is partially ordered by $(x_i)_{i \in I} \leq (y_j)_{j \in J}$ iff $I \subseteq J$ and $x_i \subseteq y_i$ for all $i \in I$. Families on A are useful in that they correspond to configurations of $!A$ – the proof of this proposition is straightforward, but observe that it crucially relies on the hypothesis that the configurations appearing in families are *non-empty*.

Proposition 7.1.3. Consider A a game. Then, the function

$$\begin{array}{ll} \text{Fam}(A) & \rightarrow \mathcal{C}(!A) \\ (x_i)_{i \in I} & \mapsto \sum_{i \in I} x_i = \bigsqcup_{i \in I} \{i\} \times x_i \end{array}$$

is an order-isomorphism – we often write $[x_i \mid i \in I] \in \mathcal{C}(!A)$ for $(x_i)_{i \in I} \in \text{Fam}(A)$.

Copy indices are special with respect to other constructions producing new moves. They keep distinct copies apart, but in principle the behaviour of strategies should not depend on the specific integers used: their exact value should not matter, which we express with our convention introduced in Section 3.4 to often write them down in grey. But we need a formal way to express that they are indeed interchangeable:

Definition 7.1.4. Consider A a game, and $x, y \in \mathcal{C}(!A)$. A **plain reindexing** $\theta : x \cong_A y$ is a bijection $\theta : x \simeq y$ such that there is $\pi : I \simeq J$ a permutation satisfying

$$\theta(i, a) = (\pi(i), a)$$

for all $(i, a) \in !A$.

A plain reindexing links two configurations $x, y \in \mathcal{C}(!A)$ which only differ via copy indices for the bang. We need more than an equivalence relation: to express that strategies are “invariant under reindexings”, we need an explicit bijection rather than the mere information that two configurations can be reindexed to each other. Indeed, while in AJM games an equivalence $s \cong_A t$ induces, chronologically, a 1-to-1 matching between the events of s and those of t , an equivalence relation does not suffice anymore for configurations, which are only partially ordered – hence explicit bijections are required.

Plain reindexings only change copy indices coming from the external $!$, but in general reindexings may affect copy indices coming from deep within the type. Thus a type should not be interpreted only with a game, but with a “game with a set of allowed reindexings” – but to what mathematical structure does that correspond?

Isomorphism families. It is an *isomorphism family* on an event structure:

Definition 7.1.5. 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 identities; 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 $\theta \supseteq \theta' \in \mathcal{S}(E)$ s.t. $\theta' : x' \simeq y'$.
- extension: for all $\theta : x \simeq y \in \mathcal{S}(E)$, $x \subseteq x' \in \mathcal{C}(E)$,
there is a (not necessarily unique) $\theta \subseteq \theta' \in \mathcal{S}(E)$ s.t. $\theta' : x' \simeq y'$.

The pair $(E, \mathcal{S}(E))$ is called an *event structure with symmetry (ess)*.

We regard isomorphism families as *proof-relevant* equivalence relations: they convey which configurations are interchangeable, witnessed by an explicit bijection.

If E is an ess, we call the elements of $\mathcal{S}(E)$ **symmetries**. We write $\theta : x \cong_E y$ to mean that $\theta : x \simeq y$ is a bijection such that $\theta \in \mathcal{S}(E)$, and write $x = \text{dom}(\theta)$ and $y = \text{cod}(\theta)$. We also write $x \cong_E y$ to mean that there is a symmetry θ such that $\theta : x \cong_E y$, this yields an equivalence relation on configurations.

It is an easy exercise to prove that symmetries are automatically order-isomorphisms:

Lemma 7.1.6. Consider E an ess, and $\theta : x \cong_E y$ a symmetry.

Then, for all $e_1, e_2 \in x$, we have $e_1 \leq_E e_2$ iff $\theta(e_1) \leq_E \theta(e_2)$.

Proof. Left as an exercise to the reader. □

Maps with symmetry. If event structures are equipped with an isomorphism family, we must adjust the notion of maps of event structures so that symmetry is preserved:

Definition 7.1.7. Consider E and F two ess, and $f : E \rightarrow F$ a map of event structures.

We say that f **preserves symmetry** if for all $\theta \in \mathcal{S}(E)$, we have $f\theta \in \mathcal{S}(F)$, where

$$f\theta = \{(fe, fe') \mid (e, e') \in \theta\}.$$

Then f is a *map of event structures with symmetry* – we write **ESS** for the category.

It is immediate that $f\theta$ may also be expressed as the composition

$$fx \xrightarrow{f^{-1}} x \xrightarrow{\theta} y \xrightarrow{f} fy$$

obtained by exploiting local injectivity of f ; the above then asks that this composition of bijections should be in $\mathcal{S}(F)$ – this observation will be used silently from now on.

In *concurrent games with symmetry*, both games and strategies will be certain event structures with symmetry, and the *display map* of a strategy will have to preserve symmetry. As for plain maps of event structures, maps of event structures with symmetry may be characterized through their action on configurations *and symmetries*:

Lemma 7.1.8. *Consider E and F two ess, and $f : E \rightarrow F$ a map of event structures. Then f preserves symmetry iff there is a (necessarily unique) function*

$$\tilde{f} : \mathcal{S}(E) \rightarrow \mathcal{S}(F)$$

commuting with dom and cod, i.e. $\text{dom} \circ \tilde{f} = f \circ \text{dom}$ and $\text{cod} \circ \tilde{f} = f \circ \text{cod}$.

Proof. Consider $\theta : x \cong_E y$. By Lemma 7.1.6, θ is an order-iso, so we get

$$\emptyset = \theta_0 \xrightarrow{(e_1, e'_1)} \theta_1 \xrightarrow{(e_2, e'_2)} \theta_2 \prec \dots \prec \theta_{n-1} \xrightarrow{(e_n, e'_n)} \theta_n = \theta \quad (7.1)$$

by following any covering chain of x or y . By commutation with dom and cod , it is immediate by induction on i that for all $0 \leq i \leq n$, $\tilde{f}\theta_i = f\theta_i$; in particular $\tilde{f}\theta = f\theta$. This shows uniqueness of \tilde{f} as it must be the action of f on symmetries as specified in Definition 7.1.7. But this means that f sends symmetries to symmetries, thus it preserves symmetry as required. \square

Altogether, we get a symmetry-aware version of the mapification lemma:

Lemma 7.1.9. *Consider E, F two ess, and a function $f : \mathcal{C}(E) \rightarrow \mathcal{C}(F)$.*

If f preserves unions, cardinality, and there is a function $\tilde{f} : \mathcal{S}(E) \rightarrow \mathcal{S}(F)$ s.t.

$$\text{dom} \circ \tilde{f} = f \circ \text{dom}, \quad \text{cod} \circ \tilde{f} = f \circ \text{cod},$$

then there exists a unique map of ess $\hat{f} : E \rightarrow F$ s.t. for all $x \in \mathcal{C}(E)$, $\hat{f}(x) = f(x)$.

Proof. By Lemma 6.1.11, there is $\hat{f} : E \rightarrow F$ a unique map of event structures. By Lemma 7.1.8 it preserves symmetry, which concludes the proof. \square

Observe that by Lemma 7.1.8 the map \tilde{f} , if it exists, is necessarily unique as it must coincide with the action of \hat{f} (the map of event structures) on symmetries as specified by Definition 7.1.7. In the sequel, when applying Lemma 7.1.9, we shall sometimes accordingly omit the definition of \tilde{f} , when we feel its explicit description is not helpful.

Maps up to Symmetry. For now, symmetry on event structures puts a new obligation on maps, which must now preserve it. We shall now see that it also provides us with a new equivalence relation on maps, which may coincide *up to symmetry* only:

Definition 7.1.10. Consider E, F ess, and $f, g : E \rightarrow F$ maps of ess.

We say that f and g are **symmetric**, written $f \sim g$, if the composite bijection

$$f \circ x \stackrel{f^{-1}}{\simeq} x \stackrel{g}{\simeq} g \circ x$$

written $\langle f, g \rangle_x$ and obtained via local injectivity of f and g , is in $\mathcal{S}(F)$ for all $x \in \mathcal{C}(E)$.

Symmetry between maps will be extremely important in the sequel: it allows us to identify maps that are the same only up to reindexing, which will be essential in the construction of models of linear logic. For symmetry to play that role, it must of course be preserved by all constructions on maps. In particular:

Proposition 7.1.11. The category **ESS**, equipped with symmetry on homsets, is enriched over equivalence relations. In particular, composition preserves \sim .

Proof. Straightforward, using that maps preserve symmetry. \square

7.1.2 Saturated Games with Symmetry

There are several ways to extend concurrent games in the presence of symmetry. Though this will not be our final word, we first sketch an approach we call “saturated” or “fat” – mathematically, it is easier, but suffers from significant drawbacks.

Games with symmetry. The motto leading to saturated concurrent games is to simply follow the same recipe as in Section 6.1.2, but building on **ESS** rather than **ES**. Accordingly, we first simply add symmetry to Definition 6.1.3:

Definition 7.1.12. A **game with symmetry (gws)** is an ess $A = (|A|, \leq_A, \#_A, \mathcal{S}(A))$ with

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

a **polarity function** such that for all $\theta : x \cong_A y$ and $a \in x$, $\text{pol}_A(a) = \text{pol}_A(\theta(a))$.

We are not far removed from traditional, sequential notions of games with symmetry from the literature. In particular, the above is a natural concurrent generalization of AJM games – the link can be easily made by first recovering alternating plays:

Definition 7.1.13. An **alternating play** on gws A is a sequence $s = s_1 \dots s_n$ which is:

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

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

It is clear that if A is a gws, then the tuple $(|A|, \text{pol}_A, \Downarrow\text{-Plays}(A))$ forms a simple game in the sense of Definition 3.1.1. To obtain an AJM game, it remains to define:

Definition 7.1.14. *Let A be a gws and $s, t \in \Downarrow\text{-Plays}(A)$.*

*We say s and t are **symmetric**, written $s \cong_A t$, if s and t have the same length, and*

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

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

As expected, we then have the AJM “unfolding” of a game with symmetry:

Proposition 7.1.15. *For any gws A , $(|A|, \text{pol}_A, \Downarrow\text{-Plays}(A), \cong_A)$ is an AJM game.*

Proof. That \cong_A is an equivalence relation follows directly from the *groupoid* condition on event structures with symmetry. There are three more conditions to check.

For (1), consider $s_1 \dots s_n \cong_A t_1 \dots t_p$. By definition, $n = p$ and

$$\theta_{s,t}^n = \{(s_i, t_i) \mid 1 \leq i \leq n\} : \{s_1, \dots, s_n\} \cong_A \{t_1, \dots, t_n\}$$

is a symmetry in $\mathcal{S}(A)$. So for all $1 \leq i \leq n$, $\text{pol}_A(s_i) = \text{pol}_A(\theta_{s,t}^n(s_i)) = \text{pol}_A(t_i)$.

Likewise, (2) is immediate from the definition.

For (3), consider $s = s_1 \dots s_n \cong t_1 \dots t_n = t$ with $sa \in \Downarrow\text{-Plays}(A)$. By definition,

$$\theta_{s,t}^n = \{(s_i, t_i) \mid 1 \leq i \leq n\} : \{s_1, \dots, s_n\} \cong_A \{t_1, \dots, t_n\}$$

is a symmetry in $\mathcal{S}(A)$. By hypothesis, $x = \{s_1, \dots, s_n\} \in \mathcal{C}(A)$ extends to $x' = x \uplus \{a\} \in \mathcal{C}(A)$. Therefore, by condition *extension*, writing $y = \{t_1, \dots, t_n\}$, there is some $y \subseteq y' \in \mathcal{C}(A)$ and an extension $\theta_{s,t}^n \subseteq \theta' \in \mathcal{S}(A)$ such that $\theta' : x' \cong_A y'$. But θ' is a bijection, so we must have $y' = y \uplus \{\theta'(a)\}$, which we write $b = \theta'(a)$. It is then immediate that $tb \in \Downarrow\text{-Plays}(A)$ with $sa \cong_A tb$ as required. \square


For the expert reader, we mention that this unfolding does not respect the usual AJM game constructions as an alternating play on a parallel composition $A \parallel B$ (though we have not yet defined this in the presence of symmetry) might not be locally alternating on A and B – this is analogous to the situation in Hyland-Ong games.

Saturated strategies. Following the same motto as before, a *concurrent prestrategy with symmetry* on gws A should simply add symmetry to Definition 6.1.5, *i.e.* should consist in a tuple $\sigma = (|\sigma|, \leq_\sigma, \#_\sigma, \mathcal{S}(\sigma), \partial_\sigma)$ where $(|\sigma|, \leq_\sigma, \#_\sigma, \mathcal{S}(\sigma))$ is an ess, and

$$\partial_\sigma : \sigma \rightarrow A$$

preserves symmetry. But while a good first step, this is not yet a working notion of (pre)strategy with symmetry: for instance, it accepts a prestrategy as in Figure 7.1, playing on $!N$, and which simply outputs the copy index used by Opponent to interrogate the natural number¹. This breaks *uniformity*: as in AJM games (see Section 3.4), in

¹Though the \mathbf{q}^- are to be symmetric in the game, nothing forces them to be symmetric in σ .

Figure 7.1: An invalid prestrategy with symmetry on $!N$


order to satisfy the expected laws of a model of linear logic we must ensure that the behaviour of strategies should not depend on Opponent's choice of copy indices.

In *saturated concurrent games*, uniformity is achieved by enforcing *isofibration*:

Definition 7.1.16. Consider A a gws. A **saturated prestrategy** on A is a tuple $\sigma = (|\sigma|, \leq_\sigma, \#_\sigma, \partial_\sigma)$ where $(|\sigma|, \leq_\sigma, \#_\sigma)$ is an ess, $\partial_\sigma : \sigma \rightarrow A$ is a map of ess, and s.t.:


isofibration: for all $x^\sigma \in \mathcal{C}(\sigma)$, for all $\theta : y_A \cong_A \partial_\sigma x^\sigma$,
there is a unique $\varphi : y^\sigma \cong_\sigma x^\sigma$ such that $\partial_\sigma \varphi = \theta$.

This removes the behaviour of Figure 7.1 – or rather, it forces there to also be



a countable non-deterministic choice letting the strategy pick a random natural number in reaction to any copy index (although the diagram only displays conflicts between contiguous numbers, it is meant that all natural numbers are pairwise conflicting). In effect, *isofibration* enforces uniformity by ensuring that the behaviour of a strategy is invariant under the addition of *noise*, non-deterministically scrambling all copy indices.

The impact of saturation. *Isifibration* is a rather neat condition mathematically speaking, but it is computationally more debatable. To see the issue, consider



a typical configuration of the *copycat strategy* on $!U$, anticipating on some definitions to come. As in AJM games, in symmetry-free concurrent games, copycat is a simple deterministic strategy merely copying Opponent moves from one side to the other.

But, it does not satisfy *isofibration*: indeed, that would force copycat to also include configurations featuring reindexing, as in Figure 7.2 but with all choices of copy indices. This “saturated copycat” would be a non-deterministic strategy, as illustrated by Figure







Figure 7.2: Reindexing Figure 7.3: Local conflict Figure 7.4: Non-local conflict

7.3, as a copy index can be propagated to any copy index on the other side but must be propagated only once. The emerging patterns of conflict can be quite complex: for instance, Figure 7.4 illustrates that by *local injectivity*, two threads independently coming up with the same index must conflict, even though they should be computationally independent if the strategies are to satisfy something like single-threadedness.

Saturated concurrent games ensure uniformity by adding non-deterministic noise. All the non-deterministic choices arising in saturated copycat yield symmetric moves, so it is in principle possible to express that the saturated copycat is deterministic “up to symmetry”, but it is quite technical and the model loses part of the concrete appeal of CG. We now explore an alternative: *thin concurrent games*.

7.1.3 Thin Concurrent Games

We wish to construct an alternative setting for strategies with symmetry, that avoids *isofibration* and instead recognizes the old – deterministic – copycat as uniform. In what sense is the deterministic copycat uniform? Well, as in AJM games: in the diagram



if Opponent reindexes their move to q_1^- , there is a *unique* matching Player reindexing of q_0^+ to q_1^+ so that we get a configuration symmetric to the original. It is this unique Player reindexing that must be captured for *thin* games. The solution is dual to saturation: in saturated games, *all* reindexings are allowed, giving rise to a non-deterministic choice. In contrast in thin games, only *one canonical* reindexing is allowed; conflicts are actual, computationally meaningful non-deterministic choices and not artifacts from symmetry.

But what is an “Opponent reindexing”, and a “Player reindexing”? Our first step will be to refine games with symmetry to include chosen symmetries that capture this.

Thin Concurrent Games. We start with the definition:

Definition 7.1.17. A *thin concurrent game (tcg)* is a gws $A = (|A|, \leq_A, \#_A, \text{pol}_A, \mathcal{S}(A))$




Figure 7.5: Positive and negative symmetries

with isomorphism families $\mathcal{S}_+(A), \mathcal{S}_-(A)$ s.t. $\mathcal{S}_+(A) \subseteq \mathcal{S}(A), \mathcal{S}_-(A) \subseteq \mathcal{S}(A)$, s.t.

- orthogonality: for all $\theta \in \mathcal{S}(A)$, if $\theta \in \mathcal{S}_+(A) \cap \mathcal{S}_-(A)$, then $\theta = \text{id}_x$ for $x \in \mathcal{C}(A)$,
- receptivity: if $\theta \in \mathcal{S}_-(A)$ and $\theta \subseteq^- \theta' \in \mathcal{S}(A)$, then $\theta' \in \mathcal{S}_-(A)$,
- +receptivity: if $\theta \in \mathcal{S}_+(A)$ and $\theta \subseteq^+ \theta' \in \mathcal{S}(A)$, then $\theta' \in \mathcal{S}_+(A)$,

where $\theta \subseteq^p \theta'$ means that $\theta \subseteq \theta'$ adding only (pairs of) events of polarity p .

Symmetries in $\mathcal{S}_+(A)$ are called **positive symmetries**, and intuitively correspond to reindexing only positive events. If $\theta : x \cong_A y$ is positive, we also write $\theta : x \cong_A^+ y$. Likewise, symmetries in $\mathcal{S}_-(A)$ are the **negative symmetries**, corresponding to reindexing only negative events, and we write $\theta : x \cong_A^- y$ if θ is negative.

Orthogonality means that if a symmetry reindexes *only* Player moves and *only* Opponent moves, it must be an identity symmetry. A negative symmetry that extends in $\mathcal{S}(A)$ with negative events must remain negative, and symmetrically.

We show in Figure 7.5 examples of positive and negative symmetries. It looks like the first should not be negative, as we rename \checkmark_0^+ to \checkmark_1^+ . However, this index is only a repetition of 0 in q_0^- : it originates in the ! on the rhs of \vdash , which is attached to q^- .

This split of symmetries between *positive* and *negative* symmetries is the core of thin concurrent games. This simple definition has the following pleasant consequence:

Lemma 7.1.18. Consider A a tcg, and $\theta : x \cong_A y$ a symmetry.

Then there are unique $z \in \mathcal{C}(A)$, $\theta^- : x \cong_A^- z$ and $\theta^+ : z \cong_A^+ y$ s.t. $\theta = \theta^+ \circ \theta^-$.

Proof. Existence. By induction on the size (i.e. cardinality) of θ . If $\theta = \emptyset$, it is clear. Take $\theta : x \cong_A y$ factoring as $\theta = \theta^+ \circ \theta^-$ for $\theta^- : x \cong_A^- z$, $\theta^+ : z \cong_A^+ y$, and assume

$$\theta \uplus \{(a, b)\} : x \uplus \{a\} \cong_A y \uplus \{b\}.$$

Since symmetries preserve polarity, $\text{pol}_A(a) = \text{pol}_A(b)$ – say *w.l.o.g.* it is positive, the other case is symmetric. Then, by *extension* on $\mathcal{S}_-(A)$, there is some extension

$$\theta^- \uplus \{(a, c)\} : x \uplus \{a\} \cong_A^- z \uplus \{c\}$$

for some c a positive event. But then, by *groupoid* laws for $\mathcal{S}(A)$ we have

$$(\theta \uplus \{(a, b)\}) \circ (\theta^- \uplus \{(a, c)\})^{-1} : z \uplus \{c\} \cong_A y \uplus \{b\}$$

which extends θ^+ with $\{(c, b)\}$, so by $+$ -receptivity, $\theta^+ \uplus \{(c, b)\} \in \mathcal{S}_+(A)$ as required.

Uniqueness. Now, assume there is an alternative factorization $\theta = \varphi^+ \circ \varphi^-$ for

$$\varphi^- : x \cong_A^- z' \quad \varphi^+ : z' \cong_A^+ y,$$

then we have $\varphi^- \circ (\theta^-)^{-1} = (\varphi^+)^{-1} \circ \theta^+ : z \cong_A z'$; but the former is negative and the latter positive, so those compositions must be the identity by *orthogonality*. \square

This is illustrated in Figure 7.5. Intuitively, any reindexing may be obtained by first reindexing Opponent moves, then reindexing Player moves, in a unique way. Obviously, by groupoid laws we can deduce the dual positive-negative factorization.

Thin concurrent strategies. Now, we define *thin concurrent strategies* on a tcg A .

As before, we give the definition in two stages. First, we define *prestrategies*:

Definition 7.1.19. A *prestrategy* on tcg A comprises an ess $(|\sigma|, \leq_\sigma, \#_\sigma, \mathcal{S}(\sigma))$ with

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

a map of event structures with symmetry called the *display map*, subject to:

- \sim -receptive: for $\theta : x \cong_\sigma y$, 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 $\theta \in \mathcal{S}(\sigma)$, if $\partial_\sigma \theta \in \mathcal{S}_+(A)$, then $\theta = \text{id}_x$ for some $x \in \mathcal{C}(\sigma)$.

where for $\psi \in \mathcal{S}(A)$, $\psi \vdash_{\mathcal{S}(A)} (a_1, a_2)$ iff $(a_1, a_2) \notin \psi$ and $\psi \uplus \{(a_1, a_2)\} \in \mathcal{S}(A)$.

We call this a *prestrategy*, keeping the same terminology as without symmetry – this should not cause confusion: if the tcg is merely a game (meaning, its symmetries are all identities), these additional conditions vacuously hold. We may still sometimes say *thin (pre)strategy*, to insist on the presence of these additional conditions.

As without symmetry, *prestrategies* are not yet adequate as strategies, but they do capture *uniformity* with respect to symmetry. In particular, \sim -receptivity forces σ to internally consider as symmetric any negative moves symmetric in the game, e.g.

$$\{\mathbf{q}_0^-\} \cong_{!B}^- \{\mathbf{q}_1^-\} \quad \Rightarrow \quad \{\mathbf{q}_0^-\} \cong_\sigma \{\mathbf{q}_1^-\}$$

for $\sigma : !B$ (provided σ has unique events matching $\mathbf{q}_0^-, \mathbf{q}_1^-$ referred to here by the same name). Then, the *extension* axiom of isomorphism families on $\mathcal{S}(\sigma)$ ensure there is

$$\begin{array}{ccc} \{\mathbf{q}_0^-\} & \cong_\sigma & \{\mathbf{q}_1^-\} \\ \downarrow \cap_+ & \Rightarrow & \downarrow \cap_+ \quad \downarrow \cap_+ \\ x & & x \cong_\sigma y \end{array}$$

then *thin* ensures that this extension is unique; another such extension y' yields

$$\theta : y \cong_\sigma y'$$




Figure 7.6: Two thin concurrent strategies

obtained by going back and forth via x . But then, displaying to the game yields

$$\text{id}_{\{q_1^-\}} \subseteq^+ \partial_\sigma \theta$$

but $\text{id}_{\{q_1^-\}} \in \mathcal{S}_+(\text{!B})$ by groupoid laws, so that $\partial_\sigma \theta \in \mathcal{S}_+(\text{!B})$ by $+$ -receptivity of positive symmetries. But then by *thin*, θ must be an identity bijection. Together, those conditions \sim -receptivity, extension, thin are a powerful way to ensure uniformity of σ .

Strategies are obtained by adding the same conditions as in Definition 6.1.6:

Definition 7.1.20. A prestrategy σ on tcg A is a **strategy** if it satisfies:

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

We write $\sigma : A$ to mean that σ is a strategy on game A .

As an illustration, we show in Figure 7.6 two examples of concurrent strategies. On the left hand side, the strategy is displayed exhaustively – its symmetries are all reduced to identities. On the right hand side, the representation is symbolic: Opponent may ask arbitrarily many times the moves $q_{0,i}^-$, $q_{1,j}^-$ and $q_{1,k}^-$. Symmetries are not represented in the diagram, but would be all bijections between configurations preserving and reflecting \rightarrow , and only changing copy indices i, j and k . This strategy is that for the term

$$\lambda f^{\text{B} \rightarrow \text{B}}. \text{if } (f \text{ tt}) \\ (\text{if } (f \text{ tt}) \perp \text{tt}) \\ (\text{if } (f \text{ ff}) \text{tt} \perp)$$

with respect to the interpretation given in Section 9.1. The left hand side strategy is the affine version of one definable with state, as shall be seen in Section 9.1.

7.1.4 Constructions on Thin Concurrent Games

We start with the basic tcgs from which all others will be defined. The **empty tcg**, written **1**, has no events. The tcgs matching *ground types* are U , B and N respectively




Figure 7.7: U




Figure 7.8: B




Figure 7.9: N

displayed in Figures 7.7, 7.8 and 7.9 – in all three cases, all symmetries are identities.

Basic constructions. First, we introduce the *simple parallel composition*. In this constructions, tcgs A_1 and A_2 are simply put side by side with all components inherited:

Definition 7.1.21. Consider two tcgs A_1 and A_2 .

Then, we define their *simple parallel composition* $A_1 \parallel A_2$ as:

events:	$ A_1 \parallel A_2 $	$=$	$ A_1 + A_2 $
causality:	$(i, a) \leq_{A_1 \parallel A_2} (j, a')$	\Leftrightarrow	$i = j \ \& \ a \leq_{A_i} a'$
conflict:	$(i, a) \#_{A_1 \parallel A_2} (j, a')$	\Leftrightarrow	$i = j \ \& \ a \#_{A_i} a'$,
symmetry:	$\theta \in \mathcal{S}(A_1 \parallel A_2)$	\Leftrightarrow	$\exists \theta_i \in \mathcal{S}(A_i), \theta = \theta_1 \parallel \theta_2$,
polarities:	$\text{pol}_{A_1 \parallel A_2}(i, a)$	$=$	$\text{pol}_{A_i}(a)$
pos. symmetries:	$\theta_1 \parallel \theta_2 \in \mathcal{S}_+(A_1 \parallel A_2)$	\Leftrightarrow	$\theta_1 \in \mathcal{S}_+(A_1) \ \& \ \theta_2 \in \mathcal{S}_+(A_2)$
neg. symmetries:	$\theta_1 \parallel \theta_2 \in \mathcal{S}_-(A_1 \parallel A_2)$	\Leftrightarrow	$\theta_1 \in \mathcal{S}_-(A_1) \ \& \ \theta_2 \in \mathcal{S}_-(A_2)$.

where, if $\theta_i : x_i \cong_{A_i} y_i$ for $i \in \{1, 2\}$, we set $(\theta_1 \parallel \theta_2)(i, e) = (i, \theta_i(e))$.

It is direct by construction that all conditions of tcgs are satisfied. Note that only taking the first four components into account, this definition also applies to plain ess.

Recall that Section 6.1.3 that configurations $x \in \mathcal{C}(A \parallel B)$ are exactly those of the form $x = x_A \parallel x_B = x_A + x_B$ where $x_A \in \mathcal{C}(A)$ and $x_B \in \mathcal{C}(B)$. We now have additional analogous decompositions for *symmetries*:

Lemma 7.1.22. Consider A and B ess. Then, there are order-isomorphisms

$$\begin{aligned}
 (- \parallel -) &: \mathcal{C}(A) \times \mathcal{C}(B) \cong \mathcal{C}(A \parallel B) \\
 (- \parallel -) &: \mathcal{S}(A) \times \mathcal{S}(B) \cong \mathcal{S}(A \parallel B) \\
 (- \parallel -) &: \mathcal{S}_-(A) \times \mathcal{S}_-(B) \cong \mathcal{S}_-(A \parallel B) \\
 (- \parallel -) &: \mathcal{S}_+(A) \times \mathcal{S}_+(B) \cong \mathcal{S}_+(A \parallel B)
 \end{aligned}$$

which commute with dom and cod, in the sense that for all $\theta_A \in \mathcal{S}(A)$ and $\theta_B \in \mathcal{S}(B)$,

$$\begin{aligned}
 \text{dom}(\theta_A \parallel \theta_B) &= \text{dom}(\theta_A) \parallel \text{dom}(\theta_B), \\
 \text{cod}(\theta_A \parallel \theta_B) &= \text{cod}(\theta_A) \parallel \text{cod}(\theta_B).
 \end{aligned}$$

Proof. Straightforward with $\theta_A \parallel \theta_B$ comprising all $((1, a), (1, a'))$ for $(a, a') \in \theta_A$ and $((2, b), (2, b'))$ for $(b, b') \in \theta_B$; and all sets ordered by componentwise inclusion. \square

From now on, we shall routinely decompose symmetries as above. Next, we introduce the other elementary operation on tcgs, the *dual*:

Definition 7.1.23. Consider A a tcg. Its *dual* A^\perp has all components as for A except:

$$\begin{aligned} \text{pol}_{A^\perp}(a) &= -\text{pol}_A(a), & (a \in A) \\ \mathcal{S}_+(A^\perp) &= \mathcal{S}_-(A), \\ \mathcal{S}_-(A^\perp) &= \mathcal{S}_+(A). \end{aligned}$$

It is clear by construction that if A is a tcg, then so is A^\perp . As for plain games, if A and B are tcgs then we set $A \vdash B$ to be $A^\perp \parallel B$. Clearly, the order-isomorphisms of Lemma 7.1.22 adapt smoothly to this construction, and let us *e.g.* write any symmetry on $A \vdash B$ as $\theta_A \vdash \theta_B : x_A \vdash x_B \cong_{A \vdash B} y_A \vdash y_B$ for $\theta_A : x_A \cong_A y_A$ and $\theta_B : x_B \cong_B y_B$.

Constructions for PCF types. In this chapter, we shall not yet detail the interpretation of programming languages. Nevertheless, we introduce here the two additional constructions required to interpret PCF types, so that as to make precise the games in the examples appearing so far. For the remaining constructions on tcgs, see Chapter 8.

The first construction is the *linear arrow* applied to tcgs with a structural constraint:

Definition 7.1.24. We define the following conditions on a thin concurrent game A :

- negative: for all $a \in \min(A)$, we have $\text{pol}_A(a) = -$,
- well-opened: A is negative, and there is exactly one $a \in \min(A)$.

where $\min(A)$ is the set of minimal events in A .

As it turns out, the interpretation of PCF types only yields well-opened tcgs, whose unique minimal events is the *initial* move by Opponent, prompting computation.

To interpret the arrow type, we first define the *linear arrow*. Intuitively, in game semantics, strategies from A to B are defined to play on A and B simultaneously, with however the polarity on A reversed to reflect its contravariance. This suggests using $A \vdash B$ as arrow type – and we shall see indeed that strategies from A to B do play on this compound game. But as in the traditional games introduced in Part I, types should yield *negative* games (at least as long as we are considering call-by-name languages); while $A \vdash B$ is not negative unless A is empty. So we set:

Definition 7.1.25. Consider A a negative tcg, and B a well-opened tcg.

Their *linear arrow* $A \multimap B$ has all components set as $A \vdash B$ except for:

$$\text{causality: } \leq_{A \multimap B} = \leq_{A \vdash B} \uplus \{((2, b_0), (1, a)) \mid a \in A\}$$

writing $\min(B) = \{b_0\}$, yielding a well-opened tcg.

This corrects the non-negativity of $A \vdash B$, by forcing the missing dependency.

Finally, as for AJM games, the last missing component of the arrow tcg is the *exponential* operation from linear logic; which finally has a non-trivial action on symmetry:

Definition 7.1.26. Consider A a negative tcg. Then, we define the **bang** $!A$ with:

$$\begin{array}{lll}
\text{events:} & |!A| & = \mathbb{N} \times |A| \\
\text{causality:} & (i, a_1) \leq_{!A} (j, a_2) & \Leftrightarrow i = j \wedge a_1 \leq_A a_2 \\
\text{conflict:} & (i, a_1) \#_{!A} (j, a_2) & \Leftrightarrow i = j \wedge a_1 \#_A a_2 \\
\text{symmetries:} & \theta \in \mathcal{S}(!A) & \Leftrightarrow \exists \pi : \mathbb{N} \simeq \mathbb{N}, \quad \exists (\theta_n)_{n \in \mathbb{N}} \in \mathcal{S}(A)^\mathbb{N} \\
& & \forall (i, a) \in \text{dom}(\theta), \theta(i, a) = (\pi(i), \theta_i(a)) \\
\text{polarities:} & \text{pol}_{!A}(i, a) & = \text{pol}_A(a) \\
\text{pos. symmetries:} & \theta \in \mathcal{S}_+(!A) & \Leftrightarrow \exists (\theta_n)_{n \in \mathbb{N}} \in \mathcal{S}_+(A)^\mathbb{N}, \\
& & \forall (i, a) \in \text{dom}(\theta), \theta(i, a) = (i, \theta_i(a)) \\
\text{neg. symmetries:} & \theta \in \mathcal{S}_-(!A) & \Leftrightarrow \exists \pi : \mathbb{N} \simeq \mathbb{N}, \exists (\theta_n)_{n \in \mathbb{N}} \in \mathcal{S}_-(A)^\mathbb{N}, \\
& & \forall (i, a) \in \text{dom}(\theta), \theta(i, a) = (\pi(i), \theta_i(a))
\end{array}$$

This yields a negative tcg $!A$.

Positive and negative symmetries for the bang were illustrated in Figure 7.5. We skip the routine verification that this defines a tcg, but it is worth observing that this depends on *negative*. If $\Theta \oplus$ is a game with two independent events with the indicated polarities, then $!(\Theta \oplus)$ is *not* a tcg: the positive symmetry $\oplus_0 \cong^+ \oplus_0$ extends with (\oplus_1, \oplus_2) to a valid symmetry, so by $+$ -receptivity it should be a positive symmetry – but it is not².

Here, we keep our convention to write copy indices in grey. With only the first four clauses, this also applies to define the bang $!E$ of a plain ess E – notice then the similarity between the definition of symmetries and Definition 3.4.3. But for a tcg, the definitions of positive and negative symmetries are asymmetric: as the bang applies to negative tcgs, the permutation π must be understood as reindexing Opponent moves – accordingly, it is arbitrary for $\mathcal{S}_-(!A)$ but restricted to the identity for $\mathcal{S}_+(!A)$.

Altogether, this lets us interpret any PCF type as a well-opened tcg, with $\llbracket \mathbb{U} \rrbracket = \mathbb{U}$, $\llbracket \mathbb{B} \rrbracket = \mathbb{B}$, $\llbracket \mathbb{N} \rrbracket = \mathbb{N}$, and $\llbracket A \rightarrow B \rrbracket = \llbracket A \rrbracket \multimap \llbracket B \rrbracket$ – we stop there for now.

7.2 Mediating Between Strategies

In Chapter 6, we have seen that (plain) concurrent games and strategies naturally organized themselves into a *bicategory*, where strategies are related by certain *morphisms* (Definition 6.1.7) – the corresponding *isomorphisms* provide an adequate equivalence on strategies. As in **CG**, strategies on thin concurrent games are to be compared via adequate notions of morphisms. But unlike in **CG**, there are quite a few candidates for those morphisms and generated equivalences; hence we start by exploring the options.

7.2.1 A Zoology of Morphisms and Equivalences

Morphisms between strategies. It is natural to first relate strategies with the same morphisms, importing Definition 6.1.7, simply replacing morphisms of event structures

²The bang in thin concurrent games behaves well only in situations where one deals with *polarized* games, *i.e.* all games are either positive or negative: no games have minimal events of mixed polarity. In the rare situations where one needs a bang on non-polarized games [Baillot et al., 1997a, Melliès, 2019b], saturated concurrent games should be used instead – in that case the reader is referred to [Castellan et al., 2014].

with morphisms of event structures with symmetry – we call those *strong morphisms*.


Definition 7.2.1. Consider σ and τ two (pre)strategies on a tcg A .

A **strong morphism** from σ to τ , is a map of ess $f : \sigma \rightarrow \tau$ satisfying:

$$\begin{aligned} \text{compatibility with display maps: } & \partial_\tau \circ f = \partial_\sigma, \\ \text{rigidity: } & \text{for all } s \leq_\sigma t, \text{ we have } fs \leq_\tau ft. \end{aligned}$$

The associated **strong isomorphism** between (pre)strategies, is written $f : \sigma \cong \tau$. The induced equivalence relation is a useful equivalence between strategies, understood as σ and τ behaving in exactly the same way – including the choice of copy indices – up to renaming of internal events. With this definition, one can replay Chapter 6 based on tcgs and obtain a bicategory of tcgs, strategies, and strong morphisms.

However, strong morphisms miss the point of symmetry, which is to identify strategies that only differ in their choice of copy indices. For instance, we wish to identify



which are two candidates for the *dereliction* strategy on U . The dereliction must choose a copy index, but it should not matter at all which one is selected – being able to perform such identifications is vital in order to get the expected laws for the exponential (!).

To address that, we weaken the commutation property, using Definition 7.1.10:

Definition 7.2.2. Consider $\sigma, \tau : A$ two (pre)strategies.

A **weak morphism** from σ to τ is a map of ess $f : \sigma \rightarrow \tau$ such that

$$\begin{aligned} \text{weak commutation: } & \partial_\tau \circ f \sim \partial_\sigma, \\ \text{rigidity: } & \text{for all } s \leq_\sigma t, \text{ we have } fs \leq_\tau ft. \end{aligned}$$

As announced, this authorizes reindexings, of both Player and Opponent moves. But intuitively, we wish to relate strategies by *only letting Player* change their indices. In tcgs, this might be achieved by appealing to the two *positive* and *negative* sub-symmetries. This is done by first considering a refinement of Definition 7.1.10:

Definition 7.2.3. Consider E an ess, A a tcg, and $f, g : E \rightarrow A$ maps of ess. Then f, g are **positively symmetric**, written $f \sim^+ g$, iff $\forall x \in \mathcal{C}(E), \langle f, g \rangle_x \in \mathcal{S}_+(A)$.

Intuitively, two parallel maps that target a tcg are *positively symmetric* if they only differ by a reindexing of Player moves. This induces a final notion of morphism:

Definition 7.2.4. Consider $\sigma, \tau : A$ two (pre)strategies.

A **positive morphism** from σ to τ , written $f : \sigma \Rightarrow \tau$, is a map $f : \sigma \rightarrow \tau$ s.t.:

$$\begin{aligned} \text{positive commutation: } & \partial_\tau \circ f \sim^+ \partial_\sigma, \\ \text{rigidity: } & \text{for all } s \leq_\sigma t, \text{ we have } fs \leq_\tau ft. \end{aligned}$$

Altogether we have *three* notions of mediating maps between (pre)strategies: *strong morphisms*, *weak morphisms* and *positive morphisms*. In the end, positive morphisms will turn out to be the most important; but they will all play a role in the development.

Equivalences. This zoology gets even more abundant when looking at the induced equivalences (ignoring strong isomorphism, as it does not support reindexing). Indeed, considering $f : \sigma \rightarrow \tau$ and $g : \tau \rightarrow \sigma$, we can require them to be inverses on the nose; or *up to symmetry*, i.e. $g \circ f \sim \text{id}_\sigma$ and $f \circ g \sim \text{id}_\tau$. This means we have four options:

$$\begin{array}{llll} \text{weak equivalence:} & \partial_\tau \circ f \sim \partial_\sigma & \partial_\sigma \circ g \sim \partial_\tau & g \circ f \sim \text{id}_\sigma \quad f \circ g \sim \text{id}_\tau \\ \text{positive equivalence:} & \partial_\tau \circ f \sim^+ \partial_\sigma & \partial_\sigma \circ g \sim^+ \partial_\tau & g \circ f \sim \text{id}_\sigma \quad f \circ g \sim \text{id}_\tau \\ \text{weak isomorphism:} & \partial_\tau \circ f \sim \partial_\sigma & \partial_\sigma \circ g \sim \partial_\tau & g \circ f = \text{id}_\sigma \quad f \circ g = \text{id}_\tau \\ \text{positive isomorphism:} & \partial_\tau \circ f \sim^+ \partial_\sigma & \partial_\sigma \circ g \sim^+ \partial_\tau & g \circ f = \text{id}_\sigma \quad f \circ g = \text{id}_\tau. \end{array}$$

Fortunately, we shall see that they all end up inducing the same equivalence relation on (pre)strategies, which we call **thin equivalence** and write $\sigma \approx \tau$. This is thanks to some important properties of positive morphisms, that we explore now.

7.2.2 A Study of Positive Morphisms

We aim to prove that these equivalences actually coincide. On our way there, we shall develop a few properties of thin concurrent strategies that have proved crucial.

Strictification of positive equivalence. Our first step here is an easy consequence of *thin* that illustrates very well what that condition achieves: the observation that a positive equivalence actually *already is* an isomorphism.

Lemma 7.2.5. *Consider $\sigma, \tau : A$ (pre)strategies on tcg A .*

Any positive equivalence $(f : \sigma \Rightarrow \tau, g : \tau \Rightarrow \sigma)$ is a positive isomorphism.

Proof. From the hypotheses, we get $g \circ f \sim \text{id}_\sigma$, meaning

$$\langle g \circ f, \text{id}_\sigma \rangle_x = \{(g(fs), s) \mid s \in x\} \in \mathcal{S}(\sigma)$$

for all $x \in \mathcal{C}(\sigma)$. But we also get $\partial_\sigma \circ g \circ f \sim^+ \partial_\sigma$, meaning

$$\langle \partial_\sigma \circ g \circ f, \partial_\sigma \rangle_x = \partial_\sigma \langle g \circ f, \text{id}_\sigma \rangle_x \in \mathcal{S}_+(A)$$

for all $x \in \mathcal{C}(\sigma)$ – but so by *thin*, $\theta_x = \text{id}_x$ for all $x \in \mathcal{C}(\sigma)$. So for all $x \in \mathcal{C}(\sigma)$ and $s \in x$, $g(f(s)) = s$. But any event $s \in \sigma$ appears in at least one configuration $[s] \in \mathcal{C}(\sigma)$, so $g \circ f = \text{id}_\sigma$. Symmetrically, $f \circ g = \text{id}_\tau$. \square

Going back and forth around a positive equivalence yields a map $g \circ f \sim \text{id}_\sigma$. But as the symmetries witnessing this display to positive symmetries in A , they must be identities as σ is thin. This property will be very important in the forthcoming development: the horizontal composition of positive isomorphisms will only yield a positive equivalence, and it is only via this lemma that it will turn out to be an isomorphism.

This strictification only holds for *positive* equivalences. To generalize it further, we shall prove that any mediating map between (pre)strategies can be made positive.

A fibration-like property. Let us investigate this *positivisation* process, and see how its crux is a fundamental *fibration-like* property of thin strategies.

Intuitively, the idea is simple: for σ and τ thin (pre)strategies, a weak morphism

$$f : \sigma \rightarrow \tau$$

changes copy indices for both Player and Opponent. The idea is, for each $s^- \in \sigma$, to set $f^+(s^-)$ as the unique event symmetric to $f(s)$ with the same copy index as s . By uniformity, we expect that the assignment f^+ should be completable to Player moves. But this would be intractable to formalize directly on events, so instead we reason on configurations. That f is only a *weak* morphism means that for $x \in \mathcal{C}(\sigma)$,

$$\langle \partial_\sigma, \partial_\tau \circ f \rangle_x : \partial_\sigma x \cong_A \partial_\tau f(x)$$


may not be in $\mathcal{S}_+(A)$. However, by Lemma 7.1.18 it factors uniquely as

$$\partial_\sigma x \xrightarrow[\cong_A^-]{\theta^+} y \xrightarrow[\cong_A^+]{\theta^-} \partial_\tau f(x);$$


where y is understood as $\partial_\tau f(x)$ with negative copy indices “set back” to those of $\partial_\sigma(x)$. Now, we would like to set $f^+(x)$ by *transporting* $f(x)$ along this symmetry θ^- . In the presence of the condition *isofibration* of saturated strategies (Definition 7.1.16), this would be straightforward, but our thin strategies do not satisfy *isofibration*.

Instead, we shall prove and exploit an alternative fibration-like property for thin strategies: we shall find $z \in \mathcal{C}(\tau)$ such that $\partial_\tau(z)$ is not *equal* to y , but *positively symmetric* to y . This is best first illustrated on an example. Consider

$$!(U \rightarrow U) \multimap U$$



a symbolic representation of the strategy $\sigma : A$ (with $A = !(B \multimap B) \multimap B$) obtained for instance by interpreting $\lambda f^{U \rightarrow U}. f(f \perp); \perp$. So σ admits as configuration:



Now, assume we have the negative symmetry reindexing the bottom move:

$$x = \begin{array}{c} \mathbf{q}_- \\ \swarrow \quad \searrow \\ \mathbf{q}_0^+ \quad \mathbf{q}_5^+ \\ \vdots \quad \vdots \\ \mathbf{q}_{0,4}^- \end{array} \cong_A^- \begin{array}{c} \mathbf{q}_- \\ \swarrow \quad \searrow \\ \mathbf{q}_0^+ \quad \mathbf{q}_5^+ \\ \vdots \quad \vdots \\ \mathbf{q}_{0,8}^- \end{array} = y$$

Clearly, σ cannot match y , but it can match it up to *positive symmetry*, with

$$\begin{array}{c} \text{!(U} \rightarrow \text{U)} \multimap \text{U} \\ \begin{array}{c} \mathbf{q}_- \\ \swarrow \quad \searrow \\ \mathbf{q}_0^+ \quad \mathbf{q}_5^+ \\ \vdots \quad \vdots \\ \mathbf{q}_{0,4}^- \end{array} \end{array} \cong_\sigma \begin{array}{c} \text{!(U} \rightarrow \text{U)} \multimap \text{U} \\ \begin{array}{c} \mathbf{q}_- \\ \swarrow \quad \searrow \\ \mathbf{q}_0^+ \quad \mathbf{q}_5^+ \\ \vdots \quad \vdots \\ \mathbf{q}_{0,8}^- \end{array} \end{array} \xrightarrow{\partial_\sigma} \begin{array}{c} \mathbf{q}_- \\ \swarrow \quad \searrow \\ \mathbf{q}_0^+ \quad \mathbf{q}_9^+ \\ \vdots \quad \vdots \\ \mathbf{q}_{0,8}^- \end{array}$$

positively isomorphic to y : Player can adapt to the change of Opponent copy indices by performing his own reindexing. This can be done systematically by exploiting uniformity of σ ; or more precisely by a straightforward successive application of \sim -receptivity and *extension*. Finally, the process results in a unique configuration via *thin*. Altogether:

Lemma 7.2.6. *Consider A a tcg, $\sigma : A$ a (pre)strategy, $x \in \mathcal{C}(\sigma)$, $\theta^- : \partial_\sigma x \cong_A^- y$.*

Then, there are unique $\varphi : x \cong_\sigma z$ and $\theta^+ : y \cong_A^+ \partial_\sigma z$ such that

$$\partial_\sigma \varphi = \theta^+ \circ \theta^- : \partial_\sigma x \cong_A \partial_\sigma z.$$

Proof. Existence. By induction on (the size of) x . If x is empty, this is clear. Otherwise, $x = x' \uplus \{s_1\}$ with s_1 maximal in x . By *restriction*, θ^- restricts to $\vartheta^- : \partial_\sigma x' \cong_A^- y'$. By induction hypothesis, there are $\varphi' : x' \cong_\sigma z'$ and $\vartheta^+ : y' \cong_A^+ \partial_\sigma z'$ such that

$$\partial_\sigma \varphi' = \vartheta^+ \circ \vartheta^- : \partial_\sigma x' \cong_A \partial_\sigma z'.$$

Furthermore, let us write $\theta^- = \vartheta^- \uplus \{(\partial_\sigma s_1, a_1)\} : \partial_\sigma x \cong_A^- y$ with $y = y' \uplus \{a_1\}$, by construction of ϑ^- . Now there are two cases, depending on the polarity of s_1 . If s_1, a_1 are negative, then by *extension* for $\mathcal{S}_+(A)$, there is an extension $\vartheta^+ \uplus \{(a_1^-, a_2^-)\} \in \mathcal{S}_+(A)$, set $\theta^+ = \vartheta^+ \uplus \{(a_1^-, a_2^-)\} : y \cong_A^+ \partial_\sigma z' \uplus \{a_2^-\}$. To sum up,

$$\varphi' : x' \cong_\sigma z', \quad x' \vdash_\sigma s_1^-, \quad \partial_\sigma \varphi' \vdash (\partial_\sigma(s_1^-), a_2^-),$$

so by \sim -receptivity there is a unique $z' \vdash_\sigma s_2^-$ such that $\varphi' \vdash_{\mathcal{S}(\sigma)} (s_1^-, s_2^-)$ and $\partial_\sigma(s_2^-) = a_2^-$. Setting $\varphi = \varphi' \uplus \{(s_1^-, s_2^-)\}$ provides the extension of φ' as required.

Now, if s_1 is positive, then by *extension* on $\mathcal{S}(\sigma)$ there is $\varphi = \varphi' \uplus \{(s_1^+, s_2^+)\} \in \mathcal{S}(\sigma)$ with $z = z' \uplus \{s_2^+\} \in \mathcal{C}(\sigma)$. But writing $a_2^+ = \partial_\sigma s_2^+$, this means that we may set

$$\theta^+ = y' \uplus \{a_1^+\} \xrightarrow{(\theta^-)^{-1}} \cong_A \partial_\sigma x \xrightarrow{\partial_\sigma \varphi} \cong_A \partial_\sigma z' \uplus \{a_2^+\}$$

in $\mathcal{S}(A)$ by *groupoid* and which extends $\vartheta^+ : y' \cong_A^+ \partial_\sigma z'$ with (a_1^+, a_2^+) , so also $\theta^+ \in \mathcal{S}_+(A)$ by *+receptivity*. By construction, we have $\partial_\sigma \varphi = \theta^+ \circ \theta^-$ as required.

Uniqueness. Consider alternative $\phi : x \cong_\sigma w$ and $\vartheta^+ : y \cong_A^+ \partial_\sigma w$ such that $\partial_\sigma \phi = \vartheta^+ \circ \theta^- : \partial_\sigma x \cong_A \partial_\sigma w$. Then we have the following commuting diagram

$$\begin{array}{ccccc} \partial_\sigma z & \xleftarrow{\partial_\sigma \varphi} & \partial_\sigma x & \xrightarrow{\partial_\sigma \phi} & \partial_\sigma w \\ & \searrow \theta^+ & \downarrow \theta^- & \nearrow \vartheta^+ & \\ & & y & & \end{array}$$

in $\mathcal{S}(A)$, showing that $\partial_\sigma(\phi \circ \varphi^{-1}) \in \mathcal{S}_+(A)$, so $z = w$ and $\phi \circ \varphi^{-1} = \text{id}_z$ by *thin*. It follows that $\varphi = \phi$, and then that $\theta^+ = \vartheta^+$, as required. \square

This is reminiscent of Mellies' notion of *uniformity by bi-invariance* [Mellies, 2003].

Via this fundamental lemma, we can mimic the consequences of the *isofibration* of saturated strategies, without saturation. We also mention the following generalization, which gives a similar fibration-like property without assuming that the symmetry along which the reindexing happens is negative – it is an easy consequence of the above.

Lemma 7.2.7. *Consider A a tcg, $\sigma : A$ a (pre)strategy, $x \in \mathcal{C}(\sigma)$, $\theta : \partial_\sigma x \cong_A y$.*

Then, there are unique $\varphi : x \cong_\sigma z$ and $\vartheta^+ : y \cong_A^+ \partial_\sigma z$ such that

$$\partial_\sigma \varphi = \vartheta^+ \circ \theta : \partial_\sigma x \cong_A \partial_\sigma z.$$

Proof. Existence. First, we factor $\theta = \theta^+ \circ \theta^-$ for some $\theta^- : \partial_\sigma x \cong_A^- w$ and $\theta^+ : w \cong_A^+ y$ as prescribed by Lemma 7.1.18. By Lemma 7.2.6, there are $\varphi : x \cong_\sigma z$ and $\psi^+ : w \cong_A^+ \partial_\sigma z$ such that $\psi^+ \circ \theta^- = \partial_\sigma \varphi$. But then, setting $\vartheta^+ = \psi^+ \circ (\theta^+)^{-1} : y \cong_A^+ \partial_\sigma z$, we have $\vartheta^+ \circ \theta = \psi^+ \circ \theta^- = \partial_\sigma \varphi$ as required.

Uniqueness. As in Lemma 7.2.6, for alternative $\phi : x \cong_\sigma w$ and $\psi^+ : y \cong_A^+ \partial_\sigma w$ such that $\partial_\sigma \phi = \psi^+ \circ \theta : \partial_\sigma x \cong_A \partial_\sigma w$, we have the following commuting diagram

$$\begin{array}{ccccc} \partial_\sigma z & \xleftarrow{\partial_\sigma \varphi} & \partial_\sigma x & \xrightarrow{\partial_\sigma \phi} & \partial_\sigma w \\ & \searrow \vartheta^+ & \downarrow \theta & \nearrow \psi^+ & \\ & & y & & \end{array}$$

in $\mathcal{S}(A)$, showing $\partial_\sigma(\phi \circ \varphi^{-1}) \in \mathcal{S}_+(A)$, which entails uniqueness by *thin*. \square

7.2.3 All Equivalences Coincide

Positivization of mediating maps. We deduce the desired positivization property:

Proposition 7.2.8. *Consider $\sigma, \tau : A$ (pre)strategies, $f : \sigma \rightarrow \tau$ a weak morphism.*

Then, there is a unique positive morphism $f^+ : \sigma \Rightarrow \tau$ such that $f \sim f^+$.

Proof. Existence. Consider $x \in \mathcal{C}(\sigma)$. By definition of weak morphisms we have

$$\langle \partial_\sigma, \partial_\tau \circ f \rangle_x = \{(\partial_\sigma(s), \partial_\tau(f(s))) \mid s \in x\} : \partial_\sigma x \cong_A \partial_\tau(f(x)),$$

and by Lemma 7.2.7, there are unique $\varphi_x : f(x) \cong_\tau z_x$ and $\theta_x^+ : \partial_\sigma x \cong_A^+ \partial_\tau z_x$ s.t.

$$\partial_\tau \varphi_x = \theta_x^+ \circ \langle \partial_\sigma, \partial_\tau \circ f \rangle_x^{-1},$$

let us call all that a *solution* for x . We set $f^+(x) = z_x$. This defines an action on configurations; we aim to induce a map by Lemma 7.1.9. Clearly, f^+ preserves cardinality as $x \simeq f(x)$ by local injectivity of f and $f(x) \cong_\tau z$; we show that f^+ preserves unions.

Consider $x, y \in \mathcal{C}(\sigma)$ such that $x \cup y \in \mathcal{C}(\sigma)$. As above, we have

$$\varphi_{x \cup y} : f(x \cup y) \cong_\tau z_{x \cup y}, \quad \theta_{x \cup y}^+ : \partial_\sigma(x \cup y) \cong_A^+ \partial_\tau z_{x \cup y}$$

such that $\partial_\tau \varphi_{x \cup y} = \theta_{x \cup y}^+ \circ \langle \partial_\sigma, \partial_\tau \circ f \rangle_{x \cup y}^{-1}$, i.e. a solution for $x \cup y$. But as $f(x \cup y) = f(x) \cup f(y)$, $\partial_\sigma(x \cup y) = \partial_\sigma(x) \cup \partial_\sigma(y)$ and $\langle \partial_\sigma, \partial_\tau \circ f \rangle_{x \cup y} = \langle \partial_\sigma, \partial_\tau \circ f \rangle_x \cup \langle \partial_\sigma, \partial_\tau \circ f \rangle_y$, by the *restriction* axiom it restricts to solutions for x and y , and must hence be their componentwise union. In particular, $f^+(x \cup y) = f^+(x) \cup f^+(y)$ as needed.

For the action on symmetries, for $\omega : x \cong_\sigma x'$ we set as $\widehat{f^+}(\omega)$ the composite

$$f^+(x) \xrightarrow{\varphi_x^{-1}} f(x) \xrightarrow{f(\omega)} f(x') \xrightarrow{\varphi_{x'}} f^+(x'),$$

clearly commuting with domain and codomain, so that by Lemma 7.1.9 there is a unique map of ess $f^+ : \sigma \rightarrow \tau$ extending the action of f^+ on configurations. For all $x \in \mathcal{C}(\sigma)$,

$$\langle \partial_\sigma, \partial_\tau \circ f^+ \rangle_x = \{(\partial_\sigma s, \partial_\tau(f^+(s))) \mid s \in x\} = \partial_\tau(\varphi_x) \circ \langle \partial_\sigma, \partial_\tau \circ f \rangle_x = \theta_x^+ \in \mathcal{S}_+(A)$$

so that $f^+ : \sigma \Rightarrow \tau$ is a positive morphism as required.

Uniqueness. Consider $f_1^+, f_2^+ : \sigma \Rightarrow \tau$ two symmetric positive morphisms. This means for all $x \in \mathcal{C}(\sigma)$ we have $\langle f_1^+, f_2^+ \rangle_x \in \mathcal{S}(\tau)$, together with

$$\begin{aligned} \langle \partial_\sigma, \partial_\tau \circ f_1^+ \rangle_x & : \partial_\sigma(x) \cong_A^+ \partial_\tau(f_1^+(x)), \\ \langle \partial_\sigma, \partial_\tau \circ f_2^+ \rangle_x & : \partial_\sigma(x) \cong_A^+ \partial_\tau(f_2^+(x)), \end{aligned}$$

but $\langle \partial_\sigma, \partial_\tau \circ f_2^+ \rangle_x \circ \langle \partial_\sigma, \partial_\tau \circ f_1^+ \rangle_x^{-1} = \partial_\tau \langle f_1^+, f_2^+ \rangle_x$, so $\partial_\tau \langle f_1^+, f_2^+ \rangle_x \in \mathcal{S}_+(A)$, so that $\langle f_1^+, f_2^+ \rangle_x$ identity by *thin*. Thus $f_1^+(x) = f_2^+(x)$ and $f_1^+ = f_2^+$ by Lemma 6.1.12. \square

From this, we may finally deduce the following:

Corollary 7.2.9. *For $\sigma, \tau : A$ (pre)strategies on tsg A , the following are equivalent:*

- (1) σ and τ are weakly equivalent,
- (2) σ and τ are positively equivalent,
- (3) σ and τ are weakly isomorphic,
- (4) σ and τ are positively isomorphic,

and we write $\sigma \approx \tau$ if that is the case.

Proof. (1) \Rightarrow (2) by Proposition 7.2.8; (2) \Rightarrow (4) by Lemma 7.2.5; (4) \Rightarrow (3) by definition; and (3) \Rightarrow (1) by definition. \square

These four variants induce the same equivalence relation on (pre)strategies, but the equivalences themselves (with the mediating maps) are not the same! In particular, going from weak equivalence to positive equivalence requires changing the maps.

From all this, it should be clear why we adopt positive morphisms and positive isomorphisms as the canonical device to compare (pre)strategies, used throughout this monograph. The developments in this section will be used often, and notably to ensure that positive morphisms are preserved by forthcoming operations on strategies.

7.2.4 Constructing Positive Morphisms

In plain concurrent games, we have argued that working with morphisms of strategies following their concrete definition – that is, as functions acting on the sets of events – is often impractical. For this purpose we introduced Lemma 6.3.4, which shows that in order to build a morphism between strategies, one can ignore trailing Opponent events: it suffices to define the action of a morphism on $+$ -covered configurations only. In the presence of symmetry, this need for tools to help construct morphisms is strengthened, so we need a similar lemma. As with Lemma 6.3.4, positive morphisms will act on $+$ -covered configurations; but also on $+$ -covered *symmetries*:

Definition 7.2.10. Consider A a tcg, and σ a (pre)strategy on A .

A symmetry $\theta \in \mathcal{S}(\sigma)$ is **$+$ -covered** if its domain (equivalently, its codomain) is. We write $\mathcal{S}^+(\sigma)$ for the set of $+$ -covered symmetries of σ .

The reader should take particular care not to confuse the set of $+$ -covered symmetries of a thin prestrategy $\sigma : A$, written $\mathcal{S}^+(\sigma)$, with the set of positive symmetries of a tcg A , written $\mathcal{S}_+(A)$. The former asks that maximal events are positive, relying on the fact that symmetries are order-isomorphisms and so preserve maximal events; while the latter is part of the data of a tcg, intuitively capturing those symmetries only reindexing Player moves. Note that only a tcg has positive and negative symmetries, not a (pre)strategy.

With this notion, we prove the following generalization of Lemma 6.3.4:

Lemma 7.2.11. Consider $\sigma, \tau : A$ two strategies. Assume there is a function $f : \mathcal{C}^+(\sigma) \rightarrow \mathcal{C}^+(\tau)$ preserving unions, and monotone functions

$$\tilde{f} : \mathcal{S}^+(\sigma) \rightarrow \mathcal{S}^+(\tau), \quad \check{f} : \mathcal{C}^+(\sigma) \rightarrow \mathcal{S}^+(A),$$

compatible with dom and cod as in Lemma 7.1.9, such that for all $x \in \mathcal{C}^+(\sigma)$,

$$\check{f}(x) : \partial_\sigma x \cong_A^+ \partial_\tau f(x),$$

and such that for all $\theta : x \cong_\sigma y \in \mathcal{S}^+(\sigma)$, we have $\partial_\tau(\check{f}(\theta)) = \check{f}(y) \circ (\partial_\sigma \theta) \circ \check{f}(x)^{-1}$.

Then, there is a unique positive $\hat{f} : \sigma \Rightarrow \tau$ such that for all $x \in \mathcal{C}^+(\sigma)$, $\hat{f}x = f x$.

Proof. Existence. We first extend f to all configurations, via the following elaboration of the reasoning in Lemma 6.3.4. Consider $x \in \mathcal{C}(\sigma)$. There is a unique $y \in \mathcal{C}^+(\sigma)$ such that $y \subseteq^- x$, obtained by removing trailing events. Now, by hypothesis we have

$$\check{f}(y) : \partial_\sigma y \cong_A^+ \partial_\tau f(y),$$

and $\partial_\sigma y \subseteq^- \partial_\sigma x$, thus by *extension*, there is $\check{f}(y) \subseteq^- \theta : \partial_\sigma x \cong_A^+ z$; it is actually unique because if we also have $\check{f}(y) \subseteq^- \theta' : \partial_\sigma x \cong_A^+ z'$, then $\theta' \circ \theta^{-1} : z \cong_A^+ z'$ is a negative extension of $\text{id}_{\partial_\sigma x}$ hence is also a negative symmetry, and thus must be an identity by *orthogonality*. Now that $\theta : \partial_\sigma x \cong_A^+ z$ is uniquely determined, we note that $\partial_\tau f(y) \subseteq^- z$ inducing a unique $f(y) \subseteq^- u \in \mathcal{C}(\tau)$ so that $\partial_\tau u = z$ by *--discrete opfibration*; we set $f(x) = u$ and $\check{f}(x) = \theta$. From the uniqueness properties of z, θ, u it is straightforward that f , extended in this way, preserves unions.

We must similarly extend \tilde{f} to all symmetries. If $\theta : x \cong_\sigma x'$, then as above there is a unique $\varphi \subseteq^- \theta$ such that $\varphi \in \mathcal{S}^+(\sigma)$, write $\tilde{\varphi} : y \cong_\sigma y'$, obtained by removing trailing (pairs of) negative moves. We have then $f(\varphi) : f(y) \cong_\tau f(y')$. Then

$$\psi = \partial_\tau f(x) \stackrel{\check{f}(x)^{-1}}{\cong_A^+} \partial_\sigma x \stackrel{\partial_\sigma \theta}{\cong_A} \partial_\sigma x' \stackrel{\check{f}(x')}{\cong_A^+} \partial_\tau f(x')$$

is a negative extension of $\check{f}(y') \circ (\partial_\sigma \varphi) \circ \check{f}(y)^{-1}$, which by hypothesis is $\partial_\tau(\tilde{f}(\theta))$. Thus by *~receptivity* of τ , there is a unique $\tilde{f}(\theta) \subseteq^- \theta'$ s.t. $\partial_\tau \theta' = \psi$, and we set $\tilde{f}(\theta) = \theta'$.

Clearly, this assignment is compatible with the extension of f with respect to domain and codomain, and thus by Lemma 7.1.9, there is a unique map of ess $\hat{f} : \sigma \rightarrow \tau$ extending f . We must show that $\partial_\sigma \sim^+ \partial_\tau \circ \hat{f}$, which amounts to showing

$$\check{f}(x) = \langle \partial_\sigma, \partial_\tau \circ \hat{f} \rangle_x \quad (7.2)$$

for all $x \in \mathcal{C}(\sigma)$, which we prove by induction on x . If x is empty, there is nothing to prove. If x is not $+$ -covered, then there is unique $y \in \mathcal{C}^+(\sigma)$ such that $y \subset^- x$. By induction hypothesis, $\check{f}(y) = \langle \partial_\sigma, \partial_\tau \circ \hat{f} \rangle_y$, and there is a covering chain

$$y \xrightarrow{s_1^-} \dots \xrightarrow{s_n^-} x$$

by induction on which it is immediate that (7.2) holds for x as well by definition of the extension of \check{f} . Now if x is $+$ -covered, it has a maximal positive event s . Setting $y = x \setminus \{s\}$, by induction hypothesis (7.2) holds for $y \in \mathcal{C}(\sigma)$, but then it holds for x by monotonicity of \check{f} . Thus we have $\partial_\sigma \sim^+ \partial_\tau \circ \hat{f}$ as required.

Rigidity. Similar to the reasoning in Lemma 6.3.4.

Uniqueness. Consider g another such extension. We first show by induction on x that for all $x \in \mathcal{C}(\sigma)$, we have $f(x) = g(x)$ and $\check{f}(x) = \langle \partial_\sigma, \partial_\tau \circ g \rangle_x$. If x is empty there is nothing to prove. If x is not $+$ -covered, there is $y \in \mathcal{C}^+(\sigma)$ such that $y \subset^- x$. By induction hypothesis, we have $f(y) = g(y)$ and $\check{f}(y) = \langle \partial_\sigma, \partial_\tau \circ g \rangle_y$. Now we have

$$\langle \partial_\sigma, \partial_\tau \circ g \rangle_y \subset^- \langle \partial_\sigma, \partial_\tau \circ g \rangle_x \in \mathcal{S}^+(A),$$

forcing $g(x)$ to coincide with the construction of $f(x)$ above; that $\check{f}(x) = \langle \partial_\sigma, \partial_\tau \circ g \rangle_x$ is immediate from the construction of $\check{f}(x)$ following a covering chain of $y \subset^- x$. Finally, if x is $+$ -covered, the result follows by $f(x) = g(x)$ and monotonicity of \check{f} . \square

As before, when using this lemma we will not always explicitly spell out the action on $+$ -covered symmetries when it is transparent from that on $+$ -covered configurations, similarly we will detail \tilde{f} only when it is non-trivial. Though unlike for Lemma 7.1.9, it remains unclear if \tilde{f} and \hat{f} are necessarily unique here.

We mention in passing two direct consequences of the above. First, we have:

Corollary 7.2.12. *Consider $f, g : \sigma \Rightarrow \tau$ two positive morphisms between strategies. If for all $x^\sigma \in \mathcal{C}^+(\sigma)$, we have $f(x^\sigma) = g(x^\sigma)$, then $f = g$.*

Proof. By uniqueness in Lemma 7.2.11. □

And then, the following is sometimes useful to construct strong isomorphisms:

Corollary 7.2.13. *Consider $\sigma, \tau : A$ strategies on tcg A . Assume there are*

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

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

Then, there is a strong isomorphism $\hat{f} : \sigma \cong \tau$ s.t. for all $x \in \mathcal{C}^+(\sigma)$, $\hat{f}(x) = f(x)$.

7.3 Composition and Copycat

Now that we have introduced thin strategies and the required tools to mediate between them, we are in position to compose them and introduce copycat. For both, the idea is to keep the structure developed in Chapter 6, only enriching it with symmetry.

For composition, the symmetries will exactly be pairs of causally compatible symmetries, as summarized by the following extension of Proposition 6.2.1:

Proposition 7.3.1. *Consider A, B, C tcgs, and $\sigma : A \vdash B$ and $\tau : B \vdash C$ strategies. Then there is a strategy $\tau \odot \sigma : A \vdash C$, unique up to strong iso, such that there are*

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

order-isomorphisms commuting with dom , cod , and such that we have

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

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

As without symmetry, this proposition packages almost all the “working concurrent game semanticist” needs to know about the composition of thin strategies. We shall gradually work our way towards its proof, and the concrete construction of composition.

In the statement, we have used for the display of symmetries the same notation as for the display of configurations in Chapter 6, that is to say, we typically write $\theta^\sigma \in \mathcal{S}(\sigma)$ for symmetries in σ for σ a (pre)strategy, and we write its display as $\partial_\sigma \theta^\sigma = \theta_A^\sigma \vdash \theta_B^\sigma$. We have also used the notion of causal compatibility for symmetries:

Definition 7.3.2. Take $\sigma : A \vdash C$, $\tau : B \vdash C$ prestrategies, $\theta^\sigma \in \mathcal{S}(\sigma)$, $\theta^\tau \in \mathcal{S}(\tau)$.

We say that θ^σ and θ^τ are **matching** if $\theta_B^\sigma = \theta_B^\tau$; and **causally compatible** if the domains $\text{dom}(\theta^\sigma)$, $\text{dom}(\theta^\tau)$ are (equivalently, if the codomains $\text{cod}(\theta^\sigma)$, $\text{cod}(\theta^\tau)$ are).

7.3.1 Interactions with Symmetry

We now set to construct concretely the composition of thin (pre)strategies. The first step is simply to enrich Definition 6.2.3 with symmetry, for A , B and C tegs.

Definition 7.3.3. An **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 $s_1 = s_2$,
- \sim -preserving: for all $\theta \in \mathcal{S}(\mu)$, $\partial(\theta) \in \mathcal{S}(A \parallel B \parallel C)$,
- thin: for all $\theta \in \mathcal{S}(\mu)$, writing $\partial(\theta) \in \theta_A \parallel \theta_B \parallel \theta_C$,
if $\theta_A \in \mathcal{S}_-(A)$ and $\theta_C \in \mathcal{S}_+(C)$, then $\theta = \text{id}_x$ for some $x \in \mathcal{C}(\mu)$.

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

Fix for now prestrategies $\sigma : A \vdash B$ and $\tau : B \vdash C$. Following the developments in Chapter 6, we have already constructed an interaction $\tau \otimes \sigma$ in the sense of Definition 6.2.3, that we need to enrich with symmetry. This relies on the following observation, letting us “zip” causally compatible symmetries into a bijection in $\tau \otimes \sigma$:

Lemma 7.3.4. Consider $\theta^\sigma : x^\sigma \cong_\sigma y^\sigma$ and $\theta^\tau : x^\tau \cong_\tau y^\tau$ causally compatible.

Then, there exists a unique bijection $\theta^\tau \otimes \theta^\sigma : x^\tau \otimes x^\sigma \simeq y^\tau \otimes y^\sigma$ such that

$$\begin{array}{ccccc}
 x^\sigma \parallel x_C^\tau & \xleftarrow{\Pi_\sigma} & x^\tau \otimes x^\sigma & \xrightarrow{\Pi_\tau} & x_A^\sigma \parallel x^\tau \\
 \theta^\sigma \parallel \theta_C^\tau \downarrow & & \downarrow \theta^\tau \otimes \theta^\sigma & & \downarrow \theta_A^\sigma \parallel \theta^\tau \\
 y^\sigma \parallel y_C^\tau & \xleftarrow{\Pi_\sigma} & y^\tau \otimes y^\sigma & \xrightarrow{\Pi_\tau} & y_A^\sigma \parallel y^\tau
 \end{array} \tag{7.3}$$

commutes in the category of finite sets and bijections.

Proof. Existence. As θ^σ and θ^τ are matching, the outer diagram commutes by Lemma 6.2.12. This lets us define the bijection $\theta^\tau \otimes \theta^\sigma : x^\tau \otimes x^\sigma \simeq y^\tau \otimes y^\sigma$ as either path along the diagram above. The condition holds by construction.

Uniqueness. Clearly, $\theta^\tau \otimes \theta^\sigma$ is uniquely determined by the diagram. \square

This invites us to set $\mathcal{S}(\tau \otimes \sigma)$ as including all $\theta^\tau \otimes \theta^\sigma$ for $\theta^\sigma \in \mathcal{S}(\sigma)$ and $\theta^\tau \in \mathcal{S}(\tau)$ causally compatible. It remains to prove that this satisfies the required conditions.

Lemma 7.3.5. *Setting $\mathcal{S}(\tau \otimes \sigma)$ to comprise all $\theta^\tau \otimes \theta^\sigma$ for $\theta^\sigma \in \mathcal{S}(\sigma)$ and $\theta^\tau \in \mathcal{S}(\tau)$ causally compatible defines an *ess* $(|\tau \otimes \sigma|, \leq_{\tau \otimes \sigma}, \#_{\tau \otimes \sigma}, \mathcal{S}(\tau \otimes \sigma))$, yielding*

$$\begin{array}{ccc}
 & \tau \otimes \sigma & \\
 \Pi_\sigma \swarrow & \vee & \searrow \Pi_\tau \\
 \sigma \parallel C & & A \parallel \tau \\
 \partial_\sigma \parallel C \searrow & & \swarrow A \parallel \partial_\tau \\
 & A \parallel B \parallel C &
 \end{array} \quad (7.4)$$

a pullback in the category of *ess* and their maps.

Proof. We verify the axioms of isomorphism families. First, *groupoid* axioms follow from the definition. For *restriction*, by definition of θ^σ and θ^τ matching we have

$$\begin{array}{ccc}
 x^\sigma \parallel x_C^\tau & \xleftarrow{\text{dom}} \varphi[x^\sigma, x^\tau] \xrightarrow{\text{cod}} & x_A^\sigma \parallel x^\tau \\
 \theta^\sigma \parallel \theta_C^\tau \downarrow & & \downarrow \theta_A^\sigma \parallel \theta^\tau \\
 y^\sigma \parallel y_C^\tau & \xleftarrow{\text{dom}} \varphi[y^\sigma, y^\tau] \xrightarrow{\text{cod}} & y_A^\sigma \parallel y^\tau
 \end{array} \quad (7.5)$$

commuting; assume now that $u^\sigma \otimes u^\tau \subseteq x^\sigma \otimes x^\tau$; i.e. we have $u^\sigma \subseteq x^\sigma$ and $u^\tau \subseteq x^\tau$ causally compatible. But then by *restriction* for σ and τ , we have

$$\begin{array}{ccc}
 u^\sigma \parallel u_C^\tau & \xleftarrow{\text{dom}} \varphi[u^\sigma, u^\tau] \xrightarrow{\text{cod}} & u_A^\sigma \parallel u^\tau \\
 \vartheta^\sigma \parallel \vartheta_C^\tau \downarrow & & \downarrow \vartheta_A^\sigma \parallel \vartheta^\tau \\
 v^\sigma \parallel v_C^\tau & \xleftarrow{\text{dom}} \varphi[v^\sigma, v^\tau] \xrightarrow{\text{cod}} & v_A^\sigma \parallel v^\tau
 \end{array}$$

with $\vartheta^\sigma \subseteq \theta^\sigma$ and $\vartheta^\tau \subseteq \theta^\tau$; in particular the diagram still commutes. *Restriction* follows by definition and Lemma 6.2.12.

For *extension*, consider again the diagram (7.5). It suffices to deal with one-step extensions, the general case follows by induction. Furthermore, we focus on an extension

$$\varphi[x^\sigma, x^\tau] \xrightarrow{((1,s),(2,t))} \varphi[u^\sigma, u^\tau],$$

with $x^\sigma \xrightarrow{s} u^\sigma$ and $x^\tau \xrightarrow{t} u^\tau$, i.e. a synchronized event between σ and τ – where $\partial_\sigma s$ and $\partial_\tau t$ occur in B – the other cases are simpler. Moreover, we assume $\text{pol}_\sigma(s) = +$, the other case is symmetric. Now, by *extension* for σ , there is $\theta^\sigma \dashv \vartheta^\sigma : u^\sigma \cong_\sigma v^\sigma$. By \sim -preserving for σ , $\partial_\sigma \vartheta^\sigma = \theta_A^\sigma \parallel \vartheta_B^\sigma$ with $\theta_B^\sigma \dashv \vartheta_B^\sigma : u_B^\sigma \cong_B v_B^\sigma$. But then, \sim -receptivity of τ exactly gives us a matching extension $\theta^\tau \subseteq \vartheta^\tau$ such that

$$\begin{array}{ccc}
 u^\sigma \parallel x_C^\tau & \xleftarrow{\text{dom}} \varphi[u^\sigma, u^\tau] \xrightarrow{\text{cod}} & x_A^\sigma \parallel u^\tau \\
 \vartheta^\sigma \parallel \vartheta_C^\tau \downarrow & & \downarrow \theta_A^\sigma \parallel \vartheta^\tau \\
 v^\sigma \parallel y_C^\tau & \xleftarrow{\text{dom}} \varphi[v^\sigma, v^\tau] \xrightarrow{\text{cod}} & y_A^\sigma \parallel v^\tau
 \end{array}$$

commutes, which provides us with the required extension by Lemma 6.2.12.

It is clear by definition (see the diagram of Lemma 7.3.4) that Π_σ and Π_τ preserve symmetry. We skip the universal property, straightforward via Lemma 7.1.9. \square

At last, this provides us with the ingredients for the interaction.

Proposition 7.3.6. *There is an interaction $\tau \otimes \sigma$, unique up to iso, s.t. 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. The ess $\tau \otimes \sigma$ is defined in Lemma 7.3.5, with display map $\partial_{\tau \otimes \sigma} : \tau \otimes \sigma \rightarrow A \parallel B \parallel C$ defined as either way around (7.4). The order-iso for configurations comes from Proposition 6.2.8, while for symmetries it follows from Lemma 7.3.4 (it is immediate from (7.3) that θ^σ and θ^τ can be recovered from $\theta^\tau \otimes \theta^\sigma$).

For *thin*, take $\theta^\tau \otimes \theta^\sigma \in \mathcal{S}(\tau \otimes \sigma)$ s.t. $\theta_A^\sigma \in \mathcal{S}_-(A)$ and $\theta_C^\tau \in \mathcal{S}_+(A)$. Consider:

$$\begin{array}{ccc} x^\sigma \parallel x_C^\tau & \xleftarrow{\text{dom}} \varphi[x^\sigma, x^\tau] \xrightarrow{\text{cod}} & x_A^\sigma \parallel x^\tau \\ \theta^\sigma \parallel \theta_C^\tau \downarrow & & \downarrow \theta_A^\sigma \parallel \theta^\tau \\ y^\sigma \parallel y_C^\tau & \xleftarrow{\text{dom}} \varphi[y^\sigma, y^\tau] \xrightarrow{\text{cod}} & y_A^\sigma \parallel y^\tau \end{array}$$

and seeking a contradiction, assume the induced $\psi : \varphi[x^\sigma, x^\tau] \simeq \varphi[y^\sigma, y^\tau]$ is not an identity. By *restriction*, we may assume w.l.o.g. that there is exactly one non-identity $((m, m'), (n, n')) \in \psi$. Several cases arise. If $m = (2, c)$, $n = (2, t)$ with $m' = (2, c')$ and $n' = (2, t')$, then $\partial_\tau \theta^\tau = \theta_B \parallel \theta_C^\tau$ with $\theta_B \in \mathcal{S}_-(B)$ as an identity symmetry and $\theta_C^\tau \in \mathcal{S}_+(C)$ by hypothesis, so θ^τ is an identity as τ is *thin*, contradiction. The case $m = (1, s)$, $n = (1, a)$, $m' = (1, s')$ and $n' = (1, a')$ is symmetric. The last case has $m = (1, s)$, $n = (2, t)$, $m' = (1, s')$ and $n' = (2, t')$. If $\text{pol}_\sigma(s) = \text{pol}_\sigma(s') = +$, then

$$\partial_\sigma \theta^\sigma = \theta_A^\sigma \parallel \theta_B^\sigma = \text{id}_{x_A^\sigma} \parallel (\text{id}_{x_B^\sigma} \uplus \{(b, b')\})$$

with $\partial_\sigma(s) = (2, b)$ and $\partial_\sigma(s') = (2, b')$. Then, $\text{pol}_B(b) = \text{pol}_B(b') = +$, so $\theta_B^\sigma \in \mathcal{S}_+(B)$ by *+receptivity*. Of course, $\theta_A^\sigma \in \mathcal{S}_-(A)$ by hypothesis, so θ^σ is an identity by *thin*, contradiction. The final case, with $\text{pol}_\tau(t) = \text{pol}_\tau(t') = +$, is symmetric.

Uniqueness. Immediate from the order-isos, by Lemma 7.1.9. \square

Here, an isomorphism of interactions simply means an iso of ess commuting with the display maps on the nose. We may now give the concrete definition of composition.

7.3.2 Composition with Symmetry

To defined composition, we simply add symmetry to Definition 6.2.16:

Definition 7.3.7. The *composition* of $\sigma : A \vdash B$ and $\tau : B \vdash C$ comprises the components of Definition 6.2.16, enriched with a set of symmetries, defined as:

$$\theta : x \cong_{\tau \odot \sigma} y \Leftrightarrow \exists \theta' \subseteq \theta : x' \cong_{\tau \otimes \sigma} y'.$$

We have yet to show that composition yields a (pre)strategy; and first of all, while the axioms of event structures are immediate, it is not obvious at all that $\mathcal{S}(\tau \odot \sigma)$ is an isomorphism family: we need a tighter characterization of symmetries:

Lemma 7.3.8. Consider $x, y \in \mathcal{C}(\tau \odot \sigma)$, and a symmetry $\theta : x \cong_{\tau \odot \sigma} y$.
There exists a unique $[\theta]_{\tau \otimes \sigma}$, the *interaction witness* of θ , s.t. $\theta \subseteq [\theta]_{\tau \otimes \sigma}$, and:

$$[\theta]_{\tau \otimes \sigma} : [x]_{\tau \otimes \sigma} \cong_{\tau \otimes \sigma} [y]_{\tau \otimes \sigma}.$$

Proof. Existence. By definition, there is $\theta \subseteq \theta' : x' \cong_{\tau \otimes \sigma} y'$. But $[x]_{\tau \otimes \sigma} \subseteq x'$, so by restriction we have $\theta'' : [x]_{\tau \otimes \sigma} \cong_{\tau \otimes \sigma} y''$ – note θ'' still contains θ as $\text{dom}(\theta) = x \subseteq [x]_{\tau \otimes \sigma}$. Now, by Lemma 7.1.6 θ'' is an order-isomorphism. From this, it is easy to show that maximal events of y'' are in y , i.e. $y'' = [y]_{\tau \otimes \sigma}$.

Uniqueness. Assume $\vartheta, \vartheta' : [x]_{\tau \otimes \sigma} \cong_{\tau \otimes \sigma} [y]_{\tau \otimes \sigma}$ such that $\theta \subseteq \vartheta, \vartheta'$. Then

$$\partial_{\tau \otimes \sigma} \vartheta = \theta_A \parallel \vartheta_B \parallel \theta_C \quad \partial_{\tau \otimes \sigma} \vartheta' = \theta_A \parallel \vartheta'_B \parallel \theta_C$$

only differ on B as $\theta \subseteq \vartheta, \vartheta'$ only adding (pairs of) synchronized events. But then,

$$\partial_{\tau \otimes \sigma}(\vartheta' \circ \vartheta^{-1}) = \text{id}_{[y]_{\tau \otimes \sigma}} \parallel (\vartheta'_B \circ \vartheta_B^{-1}) \parallel \text{id}_{[x]_{\tau \otimes \sigma}}$$

so $\vartheta' \circ \vartheta^{-1}$ is an identity since $\tau \otimes \sigma$ is *thin*; therefore $\vartheta = \vartheta'$. \square

We can now conclude that the composition yields an event structure with symmetry:

Proposition 7.3.9. We have $\tau \odot \sigma = (|\tau \odot \sigma|, \leq_{\tau \odot \sigma}, \#_{\tau \odot \sigma}, \mathcal{S}(\tau \odot \sigma))$, an *ess*.

Moreover, we have order-isomorphisms commuting with *dom* and *cod*:

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

where $\theta^\sigma, \theta^\tau$ *minimal* means $\text{dom}(\theta^\sigma), \text{dom}(\theta^\tau)$ (or $\text{cod}(\theta^\sigma), \text{cod}(\theta^\tau)$) *minimal*.

Proof. The definition of the event structure and order-isomorphism on configurations are developed in Proposition 6.2.18, which we only need to extend with symmetries. The conditions for isomorphism families follow directly from those for $\mathcal{S}(\tau \otimes \sigma)$ in combination with Lemma 7.3.8. For the order-isomorphism on symmetries, first we note that the isomorphism of Proposition 7.3.6 refine to an order-isomorphism:

$$(- \otimes -) : \{(\theta^\tau, \theta^\sigma) \in \mathcal{S}(\tau) \times \mathcal{S}(\sigma) \mid \theta^\sigma, \theta^\tau \text{ min. caus. comp.}\} \simeq \mathcal{S}^v(\tau \otimes \sigma)$$

where $\mathcal{S}^v(\tau \otimes \sigma)$ refers to symmetries whose maximal (pairs of) events are *visible*. This is clear by Lemma 6.2.12. Next, we have an order-isomorphism formed with:

$$\begin{array}{ccc} \mathcal{S}^v(\tau \otimes \sigma) & \rightarrow & \mathcal{S}(\tau \odot \sigma) \\ \theta & \mapsto & \theta \cap V^2 \end{array} \quad \begin{array}{ccc} \mathcal{S}(\tau \odot \sigma) & \rightarrow & \mathcal{S}^v(\tau \otimes \sigma) \\ \theta & \mapsto & [\theta]_{\tau \otimes \sigma} \end{array}$$

where V is the set of visible events – it is easy that this defines an order-iso. \square

This defines the composition of σ and τ as an ess, to which we may adjoin a display map specified via Lemma 7.1.9 by completing $\partial_{\tau \odot \sigma}$ with $\partial_{\tau \odot \sigma}(\theta^\tau \odot \theta^\sigma) = \theta_A^\sigma \vdash \theta_C^\tau$.

Preservation of prestrategies. We still have not accounted for a few of the additional conditions on (pre)strategies. Before wrapping up the composition of strategies, it is sometimes technically convenient to be able to compose *prestrategies*. This holds only modulo an additional condition on σ and τ , called *span-courtesy*:

Definition 7.3.10. Consider $\sigma : A \vdash B$. We define the following condition:

span-courtesy: for all $s_1 \rightarrow_\sigma s_2$, if $\text{pol}_\sigma(s_1) = +$ or $\text{pol}_\sigma(s_2) = -$, then writing $\partial_\sigma(s_i) = (k_i, c_i)$, we have $k_1 = k_2$.

This tames failures of courtesy, expressing that they cannot span across the components A, B – without it, composition of prestrategies may fail \sim -receptivity.

Now, we have the ingredients to prove the following defining property:

Proposition 7.3.11. Take $\sigma : A \vdash B$ and $\tau : B \vdash C$ span-courteous prestrategies.

Then, there is a prestrategy $\tau \odot \sigma$, unique up to strong iso, s.t. there are order-isos:

$$\begin{aligned} (- \odot -) & : \{(x^\tau, x^\sigma) \in \mathcal{C}(\tau) \times \mathcal{C}(\sigma) \mid x^\sigma, x^\tau \text{ min. caus. comp.}\} \simeq \mathcal{C}(\tau \odot \sigma) \\ (- \odot -) & : \{(\theta^\tau, \theta^\sigma) \in \mathcal{S}(\tau) \times \mathcal{S}(\sigma) \mid \theta^\sigma, \theta^\tau \text{ min. caus. comp.}\} \simeq \mathcal{S}(\tau \odot \sigma) \end{aligned}$$

commuting with dom and cod; such that for $\theta^\sigma \in \mathcal{S}(\sigma), \theta^\tau \in \mathcal{S}(\tau)$ minimal causally compatible, $\partial_{\tau \odot \sigma}(\theta^\tau \odot \theta^\sigma) = \theta_A^\sigma \parallel \theta_C^\tau$. Moreover, $\tau \odot \sigma$ is span-courteous.

Proof. Existence. The ess $\tau \odot \sigma$ is defined in Proposition 7.3.9, and the display map as above. *Thin* is immediate from *thin* for interactions, and Proposition 7.3.6.

For \sim -receptive, consider $\theta^\tau \odot \theta^\sigma : x^\tau \odot x^\sigma \cong_{\tau \odot \sigma} y^\tau \odot y^\sigma$ with extensions

$$x^\tau \odot x^\sigma \vdash_{\tau \odot \sigma} p_1^-, \quad \partial_{\tau \odot \sigma}(\theta^\tau \odot \theta^\sigma) \vdash_{\tau \odot \sigma} (\partial_{\tau \odot \sigma}(p_1^-), (2, c_2^-)),$$

assuming w.l.o.g. that $\partial_{\tau \odot \sigma}(p_1) = (2, c_1)$ occurs in C , witnessed in the interaction by

$$x^\tau \otimes x^\sigma \subseteq y^\tau \otimes y^\sigma \vdash_{\tau \otimes \sigma} p_1^-,$$

where the inclusion only adds synchronized events, which w.l.o.g. we may assume to be causal dependencies of the extension p_1 . In fact, then, $x^\tau \otimes x^\sigma = y^\tau \otimes y^\sigma$. Indeed, taking $p_0 \in y^\tau \otimes y^\sigma$ such that $p_0 \rightarrow_{\tau \otimes \sigma} p_1$, by Lemma 6.2.15 this entails that $(p_0)_\tau$

is defined and $(p_0)_\tau \rightarrow_\tau (p_1)_\tau$. But as τ is *span-courteous*, $\partial_\tau(p_0)_\tau = (2, c_0)$, so p_0 is visible. Thus $p_0 \in x^\tau \otimes x^\sigma$ necessarily. So in fact, we have:

$$x^\tau \otimes x^\sigma \vdash_{\tau \otimes \sigma} p_1^-, \quad \partial_{\tau \otimes \sigma}(\theta^\tau \otimes \theta^\sigma) \vdash_{\tau \otimes \sigma} ((3, c_1), (3, c_2))$$

which projects to $x^\tau \vdash_\tau (p_1)_\tau$ and $\partial_\tau(\theta_\tau) \vdash_\tau (\partial_\tau((p_1)_\tau), (2, c_2))$, and by \sim -receptive for τ , there is a unique $\theta^\tau \vdash_\tau ((p_1)_\tau, t_2)$ such that $\partial_\tau(t_2) = c_2$. Now, by Proposition 7.3.9 we have $y^\tau \odot y^\sigma \dashv\vdash (y^\tau \uplus \{t_2\}) \odot y^\sigma$ – those being easily seen minimal causally compatible – write $p_2 \in \tau \odot \sigma$ the added move. Likewise, we have

$$\theta^\tau \odot \theta^\sigma \dashv\vdash^{(p_1, p_2)} (\theta^\tau \uplus \{((p_1)_\tau, t_2)\}) \odot \theta^\sigma$$

which fits the requirements; uniqueness is direct from uniqueness of t_2 .

For *span-courteous*, take $p_1 \rightarrow_{\tau \odot \sigma} p_2$ s.t. $\text{pol}(p_1) = +$ or $\text{pol}(p_2) = -$, say first $\text{pol}(p_2) = -$, and *w.l.o.g.* p_2 occurs in C , *i.e.* $\partial_{\tau \odot \sigma}(p_2) = (2, c_2^-)$. By Lemma 6.2.21,

$$p_1 \rightarrow_{\tau \otimes \sigma} q_1 \rightarrow_{\tau \otimes \sigma} \dots \rightarrow_{\tau \otimes \sigma} q_n \rightarrow_{\tau \otimes \sigma} p_2$$

with all q_i synchronized, if any. But by Lemma 6.2.15, $(q_n)_\tau$ exists and $(q_n)_\tau \rightarrow_\tau (p_2)_\tau$. But as τ is *span-courteous*, this implies $\partial_\tau((q_n)_\tau) = (2, c)$, so is q_n visible – this entails $p_1 = q_n$, so $\partial_{\tau \odot \sigma} = (2, c)$ as required. The case where $\text{pol}(p_1) = +$ is analogous.

Uniqueness up to strong iso is straightforward by Lemma 7.1.9. \square

Without *span-courtesy*, \sim -receptivity may fail to be preserved under composition – a counter-example is not terribly hard to find; we omit it for the economy of presentation.

Composition of strategies. We may finally prove our main proposition:

Proposition 7.3.1. *Consider A, B, C tcgs, and $\sigma : A \vdash B$ and $\tau : B \vdash C$ strategies. Then there is a strategy $\tau \odot \sigma : A \vdash C$, unique up to strong iso, such that there are*

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

order-isomorphisms commuting with dom, cod, and such that we have

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

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

Proof. Existence. Being *courteous*, strategies are in particular *span-courteous*, which applies to show that $\tau \odot \sigma$ is a prestrategy. It is a strategy by Proposition 6.3.1, as the additional conditions do not involve symmetry. The order-isomorphisms follow from Proposition 7.3.11 specialized to $+$ -covered configurations – Lemmas 6.3.5, 6.3.6 and 6.3.7 directly extend to symmetries as they are order-isomorphisms.

Uniqueness. By Corollary 7.2.13. \square

7.3.3 Copycat

Now, we introduce the *copycat* thin strategy. As for composition, copycat will simply be the copycat strategy of Chapter 6, adequately enriched with symmetry.

Recall from Lemma 6.4.2 that for a plain game A , the configurations of \mathfrak{C}_A are those $x \vdash y \in \mathcal{C}(A \vdash A)$ such that $x \supseteq^+ x \cap y \subseteq^- y$; here the intuition is that y comprises positive events imported from x , negative events already propagated to x , but also negative events that have yet to be forwarded to x – and the situation is symmetric for x .

So what should be a symmetry

$$\theta : x_1 \vdash x_2 \cong_{\mathfrak{C}_A} y_1 \vdash y_2 ?$$

As it should display to a symmetry on $A \vdash A$, we know it should have the form $\theta_1 \vdash \theta_2$ for $\theta_1 : x_1 \cong_A y_1$ and $\theta_2 : x_2 \cong_A y_2$. But that is not all: additionally, θ_1 and θ_2 must *agree* on events that have already been propagated. We formalize this as:

Definition 7.3.12. For any tcg A , $\mathfrak{C}_A : A \vdash A$ is defined with Definition 6.4.1, with:

$$\begin{aligned} \text{symmetry: } \mathcal{S}(\mathfrak{C}_A) = \{ \theta_1 \vdash \theta_2 \in \mathcal{S}(A \vdash A) \mid \\ \text{dom}(\theta_1 \vdash \theta_2) \in \mathcal{C}(\mathfrak{C}_A), \text{cod}(\theta_1 \vdash \theta_2) \in \mathcal{C}(\mathfrak{C}_A), \\ \forall a \in \text{dom}(\theta_1) \cap \text{dom}(\theta_2), \theta_1(a) = \theta_2(a) \}. \end{aligned}$$

Note that though we have required $\theta_1(a) = \theta_2(a)$ on any element of the common domain. This looks asymmetric, but it will follow from the next lemma that we could have equivalently asked $\theta_1^{-1}(a) = \theta_2^{-1}(a)$ for any a in the common codomain.

Now we show that this gives a well-formed ess, and characterize its symmetries:

Lemma 7.3.13. Consider A a tcg. Then, $(|\mathfrak{C}_A|, \leq_A, \#_A, \mathcal{S}(\mathfrak{C}_A))$ is an ess.

Moreover, $\mathcal{S}(\mathfrak{C}_A) = \{ \theta_1 \vdash \theta_2 \mid \theta_1, \theta_2 \in \mathcal{S}(A), \theta_1 \cap \theta_2 \subseteq^+ \theta_1 \text{ \& } \theta_1 \cap \theta_2 \subseteq^- \theta_2 \}$.

Proof. We first prove the characterization of symmetries. For \subseteq , clearly $\theta_1 \cap \theta_2 \subseteq \theta_1$ and take $(a, a') \in \theta_1$ with $\text{pol}_A(a) = -$. Since $\theta_1 \vdash \theta_2 \in \mathcal{S}(\mathfrak{C}_A)$, in particular $\text{dom}(\theta_1 \vdash \theta_2) \in \mathcal{C}(\mathfrak{C}_A)$. So by Lemma 6.4.2, $a \in \text{dom}(\theta_2)$. Now, by definition of $\mathcal{S}(\mathfrak{C}_A)$ again, $\theta_1(a) = \theta_2(a)$, so $(a, a') \in \theta_1 \cap \theta_2$. The proof that $\theta_1 \cap \theta_2 \subseteq^- \theta_2$ is symmetric.

For \supseteq , consider $a \in \text{dom}(\theta_1) \cap \text{dom}(\theta_2)$ and *w.l.o.g.* assume $\text{pol}_A(a) = +$; the other case is symmetric. Since $a \in \text{dom}(\theta_2)$, there is $(a, a') \in \theta_2$. But since $\theta_1 \cap \theta_2 \subseteq^- \theta_2$, it follows that $(a, a') \in \theta_1$ as well, so $\theta_1(a) = \theta_2(a)$ as required. Next, we show $\text{dom}(\theta_1 \vdash \theta_2) \in \mathcal{C}(\mathfrak{C}_A)$ – let us write $\theta_1 : x_1 \cong_A y_1$ and $\theta_2 : x_2 \cong_A y_2$ and reason via Lemma 6.4.2. Consider $a^+ \in x_2$; we must show $a^+ \in x_1$ as well. But there is some $(a^+, b^+) \in \theta_2$, so by hypothesis $(a^+, b^+) \in \theta_1$ as well, so $a^+ \in x_1$. Symmetrically, any $a^- \in x_1$ must be in x_2 ; so $x_1 \vdash x_2 \in \mathcal{C}(\mathfrak{C}_A)$. Symmetrically, $y_1 \vdash y_2 \in \mathcal{C}(\mathfrak{C}_A)$.

Next come the axioms for isomorphism families. *Groupoid* and *restriction* follow directly from the characterization of symmetries and Lemma 6.4.2. However, *extension* is subtle – to keep the development focused we postpone it to Lemma 7.3.17. \square

Extension is really subtle. Indeed, the first significant surprise when starting the design of concurrent games with symmetry, was that copycat failed *extension* with games

plain ess with polarities (see Section A.4 in [Castellan et al., 2019]); this is a large part of what first steered us towards saturation in [Castellan et al., 2014], and of what later prompted us to look for what would eventually become the definition of tcgs. Here we opt to keep going with the development of copycat which is largely independent of the proof of *extension*, and come back to extension later on.

We mention the property on symmetries reflecting Lemma 6.4.4:

Lemma 7.3.14. *Consider A a tcg. Then, $\mathcal{S}^+(\mathfrak{C}_A) = \{\theta \vdash \theta \mid \theta \in \mathcal{S}(A)\}$.*

If $\theta \in \mathcal{S}(A)$, we write $\mathfrak{C}_\theta = \theta \vdash \theta \in \mathcal{S}^+(\mathfrak{C}_A)$ for this $+$ -covered symmetry.

Proof. Immediate from Lemmas 7.3.13 and 6.4.4. \square

Finally, we show that copycat indeed is a well-formed strategy.

Proposition 7.3.15. *For any tcg A , we have $\mathfrak{C}_A : A \vdash A$ a strategy.*

Proof. By Lemma 7.3.13, \mathfrak{C}_A is an ess. We first check the conditions for a prestrategy (i.e. Definition 7.1.19). *Rule-abiding*, *locally injective*, *\sim -preserving*, and *\sim -receptive* are direct. For *thin*, consider $\theta_1 \vdash \theta_2 : x_1 \vdash x_2 \cong_{\mathfrak{C}_A} y_1 \vdash y_2$ such that $\theta_1 \in \mathcal{S}_-(A)$ and $\theta_2 \in \mathcal{S}_+(A)$. By *restriction*, this means that $\theta_1 \cap \theta_1 \in \mathcal{S}_-(A) \cap \mathcal{S}_+(A)$. By *orthogonality*, this means that $\theta_1 \cap \theta_2 = \text{id}_x$ for some $x \in \mathcal{C}(A)$. But then, by Lemma 7.3.13, $\text{id}_x \subseteq^- \theta_2$, so by *--receptive*, $\theta_2 \in \mathcal{S}_-(A)$. So again, $\theta_2 \in \mathcal{S}_-(A) \cap \mathcal{S}_+(A)$ must be an identity by *orthogonality* – the symmetric argument shows that θ_1 is an identity.

Finally, we show the conditions of Definition 7.1.20: *courteous* follows from Lemma 6.4.3, while *receptive* is direct from Lemma 6.4.2. \square

In the sequel, we shall mostly reason on copycat via the characterizations of its $+$ -covered configurations and symmetries in Lemmas 6.4.2 and 7.3.13.

The extension property. Now, we set towards completing the proof of Lemma 7.3.13 and proving the *extension* property for $\mathcal{S}(\mathfrak{C}_A)$. This is a bit technical, but this development will not be referred to later on in this monograph, so skipping to Section 7.4 should not impede understanding of the rest of the text.

First, we find it interesting to start with a counter-example, to see where is the tension and how *close* extension is to fail. Consider $!(\Theta \oplus)$ obtained via the clauses of Definition 7.1.26, ignoring the negativity assumption and the positive and negative symmetries (recall that we noted below Definition 7.1.26 that those do *not* form a tcg). We may define $\mathfrak{C}_{!(\Theta \oplus)}$ exactly as in Definition 7.3.12, but it does *not* form an event structure with symmetry, as it fails *extension*! To see the problem, consider the symmetry

$$\begin{array}{ccc} \oplus_1 & \vdash & \ominus_1 \\ \theta_1 \cong & & \cong \theta_2 \\ \oplus_1 & \vdash & \ominus_2 \end{array}$$

linking configurations of $\mathfrak{C}_{!(\Theta \oplus)}$, where no event has been forwarded. The issue is that θ_1 and θ_2 have made irreconcilable choices regarding copy indices. It causes no issue

yet as they live in separated components, but the upper configuration may move to

$$\Theta_1 \oplus_1 \quad \vdash \quad \Theta_1$$

which, by *extension*, should trigger a corresponding extension of the bottom configuration, and of the symmetry. But the bottom configuration may only extend to $\Theta_2 \oplus_1 \vdash \Theta_2$ (by propagating to the left hand side the event Θ_2 from the right hand side), and $\Theta_1 \oplus_1 \not\equiv_{!(\Theta \oplus)} \Theta_1 \oplus_2$, because a symmetry in $!(\Theta \oplus)$ must follow a permutation of copy indices! Thus, extension fails. Fortunately for tgs this cannot happen, due to:

Lemma 7.3.16. *Consider A a tcg, $\theta, \theta_1, \theta_2 \in \mathcal{S}(A)$ such that $\theta \subseteq^+ \theta_1$ and $\theta \subseteq^- \theta_2$.*

If $\text{dom}(\theta_1)$ and $\text{dom}(\theta_2)$ are compatible, then so are θ_1 and θ_2 .

Proof. First, we prove this property for the *positive* symmetries, thus assume we have $\theta, \theta_1, \theta_2 \in \mathcal{S}_+(A)$ such that $\theta \subseteq^+ \theta_1$ and $\theta \subseteq^- \theta_2$. By *extension* for $\mathcal{S}_+(A)$, we have

$$\begin{array}{ccc} \theta'_1 : x_1 \cup x_2 \cong_A^+ y'_1 & & \theta'_2 : x_1 \cup x_2 \cong_A^+ y'_2 \\ \downarrow \cup & & \downarrow \cup \\ \theta_1 : x_1 \cong_A^+ y_1 & & \theta_2 : x_2 \cong_A^+ y_2 \\ \swarrow \times & & \searrow \subset \\ & \theta : x \cong_A^+ y & \end{array}$$

so by *groupoid*, we can form $\varphi = \theta'_2 \circ \theta'_1{}^{-1} : y'_1 \cong_A^+ y'_2$. As both θ'_1 and θ'_2 contain θ , it follows that $\text{id}_y \subseteq \varphi$. Now, by *restriction*, consider $\psi : x'_1 \cong_A^+ y_2$ obtained by restricting φ to y_2 on the right. Since $y \subseteq y_2$, we still have $\text{id}_y \subseteq \psi$. But actually that inclusion is negative, which – as $\text{id}_y \in \mathcal{S}_-(A)$ – entails that $\psi \in \mathcal{S}_-(A)$ as well. Thus by *orthogonality*, $x'_1 = y_2$ and $\psi = \text{id}_{y_2}$. Summing up, if $(a, a') \in \theta_2$, then $(a, a') \in \theta'_2$, and there is $(a'', a') \in \varphi$ for some $(a'', a) \in \theta'_1{}^{-1}$. But then $(a'', a') \in \psi$ also, hence $a'' = a'$ and $(a, a') \in \theta'_1$; so we have proved $\theta_2 \subseteq \theta'_1$, establishing compatibility of θ_1 and θ_2 . Symmetrically, the same property holds for $\mathcal{S}_-(A)$.

Now, from this we must conclude that the property holds for $\mathcal{S}(A)$, so take $\theta, \theta_1, \theta_2 \in \mathcal{S}(A)$ such that $\theta \subseteq^+ \theta_1$ and $\theta \subseteq^- \theta_2$. By Lemma 7.1.18, we have unique factorizations

$$\theta = \theta^+ \circ \theta^-, \quad \theta_1 = \theta_1^+ \circ \theta_1^-, \quad \theta_2 = \theta_2^+ \circ \theta_2^-.$$

By *uniqueness* in Lemma 7.1.18, it is immediate that this factorization is monotone, i.e. $\theta^+ \subseteq \theta_1^+$ and $\theta^+ \subseteq \theta_2^+$, so by the reasoning above we have θ_1^+ and θ_2^+ compatible in $\mathcal{S}_+(A)$, i.e. $\theta_1^+ \cup \theta_2^+ \in \mathcal{S}_+(A)$. Likewise, $\theta_1^- \cup \theta_2^- \in \mathcal{S}_-(A)$. Now it is immediate that $\theta_1^- \cup \theta_2^-$ and θ_1^+ and θ_2^+ are composable, and that we have

$$(\theta_1^+ \cup \theta_2^+) \circ (\theta_1^- \cup \theta_2^-) = (\theta_1^+ \circ \theta_1^-) \cup (\theta_2^+ \circ \theta_2^-) = \theta_1 \cup \theta_2$$

which must hence be a valid symmetry, concluding that θ_1 and θ_2 are compatible. \square

With the key lemma above, it is straightforward to finally deduce:

Lemma 7.3.17. *Consider A a tcg. Then, $\mathcal{S}(\mathfrak{A}_A)$ satisfies extension.*

Proof. Consider a symmetry on copycat, which may be written as

$$\begin{array}{ccc} x_1 & \vdash & x_2 \\ \theta_1 \downarrow & & \downarrow \theta_2 \\ y_1 & \vdash & y_2 \end{array}$$

and consider an extension of $x_1 \vdash x_2$. Without loss of generality, we can consider that only x_2 is extended, and with moves of the same polarity – this leaves two cases, for $x_2 \subseteq^- x'_2$ and $x_2 \subseteq^+ x'_2$. In the former case, then by *extension* we have $\theta_2 \subseteq \theta'_2 : x'_2 \cong_A y'_2$, and $\theta_1 \subseteq^+ \theta_1 \cap \theta_2 \subseteq^- \theta_2 \subseteq^+ \theta'_2$, so $\theta_1 \vdash \theta'_2 \in \mathcal{S}(\mathfrak{C}_A)$ by Lemma 7.3.13.

In the latter case, writing $x'_1 = x_1 \cap x'_2$, and considering the restriction θ'_1 of θ_1 to x'_1 ,

$$\theta'_1 \supseteq^+ \theta_1 \cap \theta_2 \subseteq^- \theta_2,$$

so that $\theta'_2 = \theta'_1 \cup \theta_2 \in \mathcal{S}(A)$ by Lemma 7.3.16; and $\theta_1 \vdash \theta'_2$ is the desired extension. \square

At last, this concludes the construction of the thin copycat strategy.

7.4 The Bicategory TCG

We have most of the ingredients for our bicategory **TCG**: a notion of objects (thin concurrent games), morphisms (thin strategies), 2-cells (positive morphisms). We have the data for identities (copycat strategies) and composition.

Structural 2-cells. We must still extend from **CG** the 2-cells for associativity and unity laws, which is direct from the developments above. For associators, we have:

Proposition 7.4.1. *Consider thin prestrategies $\sigma : A \vdash B, \tau : B \vdash C$ and $\delta : C \vdash D$. Then, there is a strong isomorphism of strategies, the **associator***

$$a_{\sigma, \tau, \delta} : (\delta \odot \tau) \odot \sigma \cong \delta \odot (\tau \odot \sigma),$$

such that for all $(x^\delta \odot x^\tau) \odot x^\sigma \in \mathcal{C}((\delta \odot \tau) \odot \sigma)$ and $(\theta^\delta \odot \theta^\tau) \odot \theta^\sigma \in \mathcal{S}((\delta \odot \tau) \odot \sigma)$,

$$\begin{aligned} a_{\sigma, \tau, \delta}((x^\delta \odot x^\tau) \odot x^\sigma) &= x^\delta \odot (x^\tau \odot x^\sigma) \\ a_{\sigma, \tau, \delta}((\theta^\delta \odot \theta^\tau) \odot \theta^\sigma) &= \theta^\delta \odot (\theta^\tau \odot \theta^\sigma). \end{aligned}$$

Proof. To apply Lemma 7.1.9 we show the associator preserves symmetry, which is clear by the description of the symmetries in the composition in Proposition 7.3.9. \square

Likewise, the developments above make it easy to define the *unitors*:

Proposition 7.4.2. *Consider $\sigma : A \vdash B$ a thin strategy.*

*Then there are strong isomorphisms, the **unitors**:*


$$l_\sigma : \mathfrak{C}_B \odot \sigma \cong \sigma \quad r_\sigma : \sigma \odot \mathfrak{C}_A \cong \sigma$$

such that for all $x^\sigma \in \mathcal{C}^+(\sigma)$ and $\theta^\sigma \in \mathcal{S}^+(\sigma)$,

$$\begin{aligned} l_\sigma(\mathfrak{C}_{x_B^\sigma} \odot x^\sigma) &= x^\sigma & r_\sigma(x^\sigma \odot \mathfrak{C}_{x_A^\sigma}) &= x^\sigma \\ l_\sigma(\mathfrak{C}_{\theta_B^\sigma} \odot \theta^\sigma) &= \theta^\sigma, & r_\sigma(\theta^\sigma \odot \mathfrak{C}_{\theta_A^\sigma}) &= \theta^\sigma. \end{aligned}$$

Proof. Obvious by Corollary 7.2.13. \square

Horizontal composition. There is still one piece of data missing: the horizontal composition of positive morphisms, letting us assemble two positive morphisms as in



into $g \odot f : \tau \odot \sigma \Rightarrow \tau' \odot \sigma'$. In **CG** this was simple, as we could just set

$$(g \odot f)(x^\tau \odot x^\sigma) = g(x^\tau) \odot f(x^\sigma)$$

defining a 2-cell in **CG** via Lemma 6.3.4. In **TCG** this would also work, for *strong* morphisms. But in general, for *positive* morphisms, $f(x^\sigma) \in \mathcal{C}(\sigma')$ and $g(x^\tau) \in \mathcal{C}(\tau')$ have no reason to be matching as we only have $\partial_{\sigma'} \circ f \sim^+ \partial_\sigma$ and $\partial_{\tau'} \circ g \sim^+ \partial_\tau$.

By the definition of positively symmetric maps, for all $x^\tau \odot x^\sigma \in \mathcal{C}^+(\tau \odot \sigma)$, there are (necessarily unique) $\theta_A^- \in \mathcal{S}_-(A)$, $\theta_B^+ \in \mathcal{S}_+(B)$, $\theta_B^- \in \mathcal{S}_-(B)$ and $\theta_C^+ \in \mathcal{S}_+(C)$ s.t.

$$\begin{array}{ccc} x^\sigma & \xrightarrow{f} & f(x^\sigma) \\ \partial_\sigma \downarrow & & \downarrow \partial_{\sigma'} \\ x_A^\sigma \vdash x_B^\sigma & \xrightarrow{\theta_A^- \vdash \theta_B^+} & f(x^\sigma)_A \vdash f(x^\sigma)_B \end{array} \quad \begin{array}{ccc} x^\tau & \xrightarrow{g} & g(x^\tau) \\ \partial_\tau \downarrow & & \downarrow \partial_{\tau'} \\ x_B^\tau \vdash x_C^\tau & \xrightarrow{\theta_B^- \vdash \theta_C^+} & g(x^\tau)_B \vdash g(x^\tau)_C \end{array} \quad (7.6)$$

commute in the category of finite sets of bijections (f, g and the display maps being locally injective). Let us write $f[x^\sigma]_A = \theta_A^-$, $f[x^\sigma]_B = \theta_B^+$, $g[x^\tau]_B = \theta_B^-$ and $g[x^\tau]_C = \theta_C^+$ to emphasize that those symmetries are part of f and g , and the dependency on x^σ and x^τ . So, though we lack $f(x^\sigma)_B = g(x^\tau)_B$ as *matching* would require, we do have

$$\theta : f(x^\sigma)_B \cong_B g(x^\tau)_B$$

where $\theta = g[x^\tau]_B \circ f[x^\sigma]_B^{-1}$, and using that $x_B^\sigma = x_B^\tau$: $f(x^\sigma)$ and $g(x^\tau)$ are *matching up to symmetry*. The fundamental property we must show, is that such a situation can be massaged to extract configurations of σ' and τ' matching *on the nose*.

This is our final required conceptual tool, *synchronization up to symmetry*:

7.4.1 Synchronization up to Symmetry

Fix $\sigma : A \vdash B$ and $\tau : B \vdash C$ two (pre)strategies.

First, we must specify under which conditions we expect to be able to synchronize two configurations up to symmetry. As without symmetry, there needs to be some causal compatibility condition – accordingly, we first generalize the definition of *causally compatible pairs* (Definition 6.2.2) to account for a mediating symmetry:

Definition 7.4.3. Consider $x^\sigma \in \mathcal{C}(\sigma)$, $x^\tau \in \mathcal{C}(\tau)$ and $\theta : x_B^\sigma \cong_B x_B^\tau$.

We say that the triple $(x^\sigma, \theta, x^\tau)$ is **causally compatible** if the composite bijection

$$\varphi[x^\sigma, \theta, x^\tau] : x^\sigma \parallel x_C^\tau \xrightarrow{\partial_\sigma \parallel x_C^\tau} x_A^\sigma \parallel x_B^\sigma \parallel x_C^\tau \xrightarrow{x_A^\sigma \parallel \theta \parallel x_C^\tau} x_A^\sigma \parallel x_B^\tau \parallel x_C^\tau \xrightarrow{\partial_\tau^{-1} \parallel x_C^\tau} x_A^\sigma \parallel x^\tau$$

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) $\varphi[x^\sigma, \theta, x^\tau]$ from the causal constraints of σ and τ , is acyclic.

In that case, the reflexive transitive closure of \triangleleft is a partial order $\leq_{\varphi[x^\sigma, \theta, x^\tau]}$, and $(x^\sigma, \theta, x^\tau)$ is **minimal** if the maximal elements of $\varphi[x^\sigma, \theta, x^\tau]$ occur in A or C .

At the core of **TCG** is the following property, stating that any causally compatible triple can be massaged to produce a unique equivalent causally compatible pair:

Proposition 7.4.4. Consider $(x^\sigma, \theta, x^\tau)$ a causally compatible triple.

Then, there are unique $y^\sigma \in \mathcal{C}(\sigma)$ and $y^\tau \in \mathcal{C}(\tau)$ causally compatible, and

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

such that we have $\varphi_A^\sigma \in \mathcal{S}_-(A)$, $\varphi_C^\tau \in \mathcal{S}_+(C)$, and the following diagram commutes:

$$\begin{array}{ccccccc} x^\sigma \parallel x_C^\tau & \xrightarrow{\partial_\sigma \parallel x_C^\tau} & x_A^\sigma \parallel x_B^\sigma \parallel x_C^\tau & \xrightarrow{x_A^\sigma \parallel \theta \parallel x_C^\tau} & x_A^\sigma \parallel x_B^\tau \parallel x_C^\tau & \xrightarrow{x_A^\sigma \parallel \partial_\tau^{-1}} & x_A^\sigma \parallel x^\tau \\ & \searrow \varphi^\sigma \parallel \varphi_C^\tau & \searrow \varphi_A^\sigma \parallel \varphi_B^\sigma \parallel \varphi_C^\tau & \searrow \varphi_A^\sigma \parallel \varphi_B^\tau \parallel \varphi_C^\tau & \searrow \varphi_A^\sigma \parallel \varphi_B^\tau \parallel \varphi_C^\tau & \searrow \varphi_A^\sigma \parallel \varphi^\tau & \\ y^\sigma \parallel y_C & \xrightarrow{\partial_\sigma \parallel y_C} & y_A \parallel y_B \parallel y_C & \xrightarrow{y_A \parallel \partial_\tau^{-1}} & y_A \parallel y^\tau & & \end{array}$$

(in particular, the middle triangle commutes – the squares come for free).

Proof. Securedness follows the obvious generalization of Definition 6.2.2.

Existence. For a secured bijection ϕ (top of the diagram) as above, assume adequate $\varphi^\sigma : x^\sigma \cong_\sigma y^\sigma$ and $\varphi^\tau : x^\tau \cong_\tau y^\tau$ are provided by induction hypothesis. Consider:

$$\phi \xrightarrow{(m,n)} \phi'$$

with ϕ' a composite bijection. We first consider the key case, with $m = (1, s)$ for $x^\sigma \vdash_\sigma s^+$, and $n = (2, t)$ for $x^\tau \vdash_\tau t^-$ synchronized. Let us write their displays as

$$\partial_\sigma(x^\sigma \uplus \{s\}) = x_A^\sigma \vdash (x_B^\sigma \uplus \{b_1\}), \quad \partial_\tau(x^\tau \uplus \{t\}) = (x_B^\tau \uplus \{b_2\}) \vdash x_C^\tau$$

synchronizing through $\theta \uplus \{(b_1, b_2)\} \in \mathcal{S}(B)$.

By *extension* for σ , we may propagate this extension through φ^σ to obtain:

$$\psi^\sigma = \varphi^\sigma \uplus \{(s, s')\} : x^\sigma \uplus \{s\} \cong_\sigma y^\sigma \uplus \{s'\}$$

displaying to $\partial_\sigma(\psi^\sigma) = \varphi_A^\sigma \vdash (\varphi_B^\sigma \uplus \{(b_1, b')\}) \in \mathcal{S}(A \vdash B)$. But then:

$$\varphi_B^\tau \uplus \{(b_2, b')\} \in \mathcal{S}(B)$$

as it may be obtained as the composition $(\varphi_B^\sigma \uplus \{(b_1, b')\}) \circ (\theta \uplus \{(b_1, b_2)\})^{-1}$ of known symmetries in B . Then \sim -receptivity directly applies to provide $\psi^\tau = \varphi^\tau \uplus \{(t, t')\}$ matching this extension, concluding this case – missing verifications are immediate.

The case with s negative and t positive is symmetric. Next, consider the case $m = (2, c)$ and $n = (2, t)$ with $x^\tau \vdash_\tau t^-$, *i.e.* a visible, negative move in C . Its display is

$$\partial_\tau(x^\tau \uplus \{t\}) = x_B^\tau \vdash (x_C^\tau \uplus \{c\}),$$

but then *extension* of $\mathcal{S}_+(C)$ applies to provide an extension $\varphi_C^\tau \uplus \{(c, c')\} \in \mathcal{S}_+(C)$. By \sim -receptivity of τ , there is an (unique) matching extension $\psi^\tau = \varphi^\tau \uplus \{(t, t')\} \in \mathcal{S}(\tau)$, concluding this case – other components are unchanged.

Consider then the case $m = (2, c)$ and $n = (2, t)$ with $x^\tau \vdash_\tau t^+$. In that case,

$$\psi^\tau = \varphi^\tau \uplus \{(t, t')\} \in \mathcal{S}(\tau)$$

is provided by *extension* for τ . Considering its display $\partial_\tau(\psi^\tau) = \varphi_B^\tau \vdash (\varphi_C^\tau \uplus \{(c, c')\})$, we do have $\psi_C^\tau \in \mathcal{S}_+(C)$ as $\varphi_C^\tau \in \mathcal{S}_+(C)$ and by $+$ -receptivity of tcgs. This concludes this case – other components are unchanged. All missing cases are symmetric. Finally, *securedness* follows from the diagram and securedness of $(x^\sigma, \theta, x^\tau)$, as all vertical morphisms in the diagram are order-isomorphisms.

Uniqueness. Consider now $\psi^\sigma : x^\sigma \cong_\sigma z^\sigma$ and $\psi^\tau : x^\tau \cong_\tau z^\tau$ also satisfying the hypotheses. But then, by composition of symmetries we also have

$$\omega^\sigma = \psi^\sigma \circ (\varphi^\sigma)^{-1} : y^\sigma \cong_\sigma z^\sigma, \quad \omega^\tau = \psi^\tau \circ (\varphi^\tau)^{-1} : y^\tau \cong_\tau z^\tau,$$

but additionally these symmetries are causally compatible, since

$$\begin{aligned} \omega_B^\sigma &= \psi_B^\sigma \circ (\varphi_B^\sigma)^{-1} \\ &= \psi_B^\tau \circ \theta \circ (\varphi^\tau \circ \theta)^{-1} \\ &= \omega_B^\tau \end{aligned}$$

hence we have $\omega^\tau \otimes \omega^\sigma : y^\tau \otimes y^\sigma \cong_{\tau \otimes \sigma} z^\tau \otimes z^\sigma$ by Proposition 7.3.6. But additionally,

$$(\omega^\tau \otimes \omega^\sigma)_A = \psi_A^\sigma \circ (\varphi_A^\sigma)^{-1} \in \mathcal{S}_-(A),$$

and likewise $(\omega^\tau \otimes \omega^\sigma)_C \in \mathcal{S}_+(C)$, so by *thin* for $\tau \otimes \sigma$, $\omega^\tau \otimes \omega^\sigma = \text{id}_{y^\tau \otimes y^\sigma}$. It follows from Proposition 7.3.6 that ω^σ and ω^τ are identities, so $\varphi^\sigma = \psi^\sigma$ and $\varphi^\tau = \psi^\tau$. \square

The above may be applied to compose configurations up to symmetry:

Corollary 7.4.5. Consider $(x^\sigma, \theta, x^\tau)$ a minimal causally compatible triple.

Then, the pair (y^σ, y^τ) given by Proposition 7.4.4 is minimal causally compatible.

Proof. By Proposition 7.4.4, we have $\varphi^\sigma : x^\sigma \cong_\sigma y^\sigma$ and $\varphi^\tau : x^\tau \cong_\tau y^\tau$ making

$$\begin{array}{ccc} x^\sigma \parallel x_C^\tau & \xrightarrow{\varphi[x^\sigma, \theta, x^\tau]} & x_A^\sigma \parallel x^\tau \\ \varphi^\sigma \parallel \varphi_C^\tau \downarrow & & \downarrow \varphi_A^\sigma \parallel \varphi^\tau \\ y^\sigma \parallel y_C^\tau & \xrightarrow{\varphi[y^\sigma, y^\tau]} & y_A^\sigma \parallel y^\tau \end{array}$$

commute. But the vertical morphisms are order-isomorphisms, inducing an order-iso between the (graphs of) $\varphi[x^\sigma, \theta, x^\tau]$ and $\varphi[y^\sigma, y^\tau]$. In particular it preserves maximal elements, so that the minimality of (y^σ, y^τ) follows from that of $(x^\sigma, \theta, x^\tau)$. \square

This finally lets us define the composition of configurations up to symmetry:

Definition 7.4.6. Consider $(x^\sigma, \theta, x^\tau)$ a minimal causally compatible triple. We set

$$x^\tau \odot_\theta x^\sigma = y^\tau \odot y^\sigma \in \mathcal{C}(\tau \odot \sigma)$$

for y^σ, y^τ the minimal causally compatible pair given by Corollary 7.4.5.

It will be useful later on that this construction preserves unions:

Lemma 7.4.7. Consider $(x_1^\sigma, \theta_1, x_1^\tau), (x_2^\sigma, \theta_2, x_2^\tau)$ minimal causally compatible triples, also compatible with each other in the sense that $(x_1^\sigma \cup x_2^\sigma, \theta_1 \cup \theta_2, x_1^\tau \cup x_2^\tau)$ is a minimal causally compatible triple. Then, $x_1^\tau \odot_{\theta_1} x_1^\sigma$ and $x_2^\tau \odot_{\theta_2} x_2^\sigma$ are compatible, and

$$(x_1^\tau \cup x_2^\tau) \odot_{\theta_1 \cup \theta_2} (x_1^\sigma \cup x_2^\sigma) = (x_1^\tau \odot_{\theta_1} x_1^\sigma) \cup (x_2^\tau \odot_{\theta_2} x_2^\sigma)$$

Proof. Applying Proposition 7.4.4 to $(x^\sigma, \theta, x^\tau) = (x_1^\sigma \cup x_2^\sigma, \theta_1 \cup \theta_2, x_1^\tau \cup x_2^\tau)$, we get

$$\begin{array}{ccccccc} x^\sigma \parallel x_C^\tau & \xrightarrow{\partial_\sigma \parallel x_C^\tau} & x_A^\sigma \parallel x_B^\sigma \parallel x_C^\tau & \xrightarrow{x_A^\sigma \parallel \theta \parallel x_C^\tau} & x_A^\sigma \parallel x_B^\tau \parallel x_C^\tau & \xrightarrow{x_A^\sigma \parallel \partial_\tau^{-1}} & x_A^\sigma \parallel x^\tau \\ \searrow \varphi^\sigma \parallel \varphi_C^\tau & & \searrow \varphi_A^\sigma \parallel \varphi_B^\sigma \parallel \varphi_C^\tau & & \searrow \varphi_A^\sigma \parallel \varphi_B^\tau \parallel \varphi_C^\tau & & \searrow \varphi_A^\sigma \parallel \varphi^\tau \\ y^\sigma \parallel y_C^\tau & \xrightarrow{\partial_\sigma \parallel y_C^\tau} & y_A^\sigma \parallel y_B^\sigma \parallel y_C^\tau & \xrightarrow{y_A^\sigma \parallel \partial_\tau^{-1}} & y_A^\sigma \parallel y_B^\tau \parallel y_C^\tau & \xrightarrow{y_A^\sigma \parallel \partial_\tau^{-1}} & y_A^\sigma \parallel y^\tau \end{array}$$

with similar diagrams for $(x_1^\sigma, \theta_1, x_1^\tau)$ and $(x_2^\sigma, \theta_2, x_2^\tau)$ yielding causally compatible pairs (y_1^σ, y_1^τ) and (y_2^σ, y_2^τ) . But the diagram above restricts to diagrams for $(x_1^\sigma, \theta_1, x_1^\tau)$ and $(x_2^\sigma, \theta_2, x_2^\tau)$, so by the uniqueness clause of Proposition 7.4.4 it must be the componentwise union of the diagrams for $(x_1^\sigma, \theta_1, x_1^\tau)$ and $(x_2^\sigma, \theta_2, x_2^\tau)$. Thus, we have

$$\varphi[y^\sigma, y^\tau] = \varphi[y_1^\sigma, y_1^\tau] \cup \varphi[y_2^\sigma, y_2^\tau]$$

which, restricting to prime sub-secured bijections with visible top, yields

$$x^\tau \odot_\theta x^\sigma = (x_1^\tau \odot_{\theta_1} x_1^\sigma) \cup (x_2^\tau \odot_{\theta_2} x_2^\sigma)$$

as required, implying also their compatibility. \square

7.4.2 Horizontal Composition

We can now use this to define horizontal composition of positive morphisms:

Proposition 7.4.8. *Consider $\sigma, \sigma' : A \vdash B$ and $\tau, \tau' : B \vdash C$ (pre)strategies, along with $f : \sigma \Rightarrow \sigma'$ and $g : \tau \Rightarrow \tau'$ positive morphisms.*

Then, there is a unique positive morphism $g \odot f : \tau \odot \sigma \rightarrow \tau' \odot \sigma'$ such that

$$(g \odot f)(x^\tau \odot x^\sigma) = g(x^\tau) \odot_{g[x^\tau]_B \circ f[x^\sigma]_B^{-1}} f(x^\sigma)$$

for all $x^\tau \odot x^\sigma \in \mathcal{C}(\tau \odot \sigma)$ where $f[x^\sigma]_B$ and $g[x^\tau]_B$ arise as below (7.6).

Proof. Existence. Unfortunately, we cannot directly apply Lemma 7.2.11, because we want this proposition to handle positive morphisms between *prestrategies*. Thus, we construct $g \odot f$ via Lemma 7.1.9. Its action on configurations is given by definition – in particular, it follows immediately from *rigidity* that if (x^σ, x^τ) is minimal causally compatible, then $(f(x^\sigma), g[x^\tau]_B \circ f[x^\sigma]_B^{-1}, g(x^\tau))$ is minimal causally compatible. Preservation of unions follows from Lemma 7.4.7 and the fact that

$$g[x^\tau]_B \circ f[x^\sigma]_B^{-1} = (g[x_1^\tau]_B \circ f[x_1^\sigma]_B^{-1}) \cup (g[x_2^\tau]_B \circ f[x_2^\sigma]_B^{-1})$$

whenever $x^\sigma = x_1^\sigma \cup x_2^\sigma$ and $x^\tau = x_1^\tau \cup x_2^\tau$. Preservation of cardinality follows from the preservation of the display to the game (see *positivity* below).

We must show that this action extends to symmetries. Consider

$$\psi^\tau \odot \psi^\sigma : x^\tau \odot x^\sigma \cong_{\tau \odot \sigma} y^\tau \odot y^\sigma$$

a symmetry in $\tau \odot \sigma$; so $\psi^\sigma : x^\sigma \cong_\sigma y^\sigma$, $\psi^\tau : x^\tau \cong_\tau y^\tau$ with $\psi_B^\sigma = \psi_B^\tau$. We name:

$$\begin{array}{ccc} f(x^\sigma)_B & \xleftarrow{(\varphi_B^\sigma)^{-1}} x_B^\sigma = x_B^\tau & \xrightarrow{g[x^\tau]_B} g(x^\tau)_B \\ & \searrow (\nu_B^\tau)^{-1} & \swarrow \varphi_B^\tau \\ & z_B^{\sigma'} = z_B^{\tau'} & \end{array} \quad \begin{array}{ccc} f(y^\sigma)_B & \xleftarrow{f[y^\sigma]_B} y_B^\sigma = y_B^\tau & \xrightarrow{g[y^\tau]_B} g(y^\tau)_B \\ & \searrow \nu_B^\sigma & \swarrow \nu_B^\tau \\ & u_B^{\sigma'} = u_B^{\tau'} & \end{array}$$

writing $z^{\tau'} \odot z^{\sigma'} = g(x^\tau) \odot_\theta f(x^\sigma)$ for $\theta = g[x^\tau]_B \circ f[x^\sigma]_B^{-1}$, and $u^\tau \odot u^\sigma = g(y^\tau) \odot_\vartheta f(y^\sigma)$ for $\vartheta = g[y^\tau]_B \circ f[y^\sigma]_B^{-1}$. Next we note that the diagram below commutes

$$\begin{array}{ccccc} f(x^\sigma)_B & \xleftarrow{f[x^\sigma]_B} x_B^\sigma = x_B^\tau & \xrightarrow{g[x^\tau]_B} & g(x^\tau)_B & \\ f(\psi^\sigma)_B \downarrow & & \psi_B^\sigma = \psi_B^\tau \downarrow & & \downarrow g(\psi^\tau)_B \\ f(y^\sigma)_B & \xleftarrow{f[y^\sigma]_B} y_B^\sigma = y_B^\tau & \xrightarrow{g[y^\tau]_B} & g(y^\tau)_B & \end{array}$$

by definition of the components involved. But this means that the following diagram

$$\begin{array}{ccc}
 & z_B^{\sigma'} = z_B^{\tau'} & \\
 (\varphi_B^\sigma)^{-1} \swarrow & & \searrow (\varphi_B^\tau)^{-1} \\
 f(x^\sigma)_B & \xrightarrow{g[x^\tau]_B \circ f[x^\sigma]_B^{-1}} & g(x^\tau)_B \\
 f(\psi^\sigma)_B \downarrow & & \downarrow g(\psi^\tau)_B \\
 f(y^\sigma)_B & \xrightarrow{g[y^\tau]_B \circ f[y^\sigma]_B^{-1}} & g(y^\tau)_B \\
 & \swarrow \nu_B^\sigma \quad \nwarrow \nu_B^\tau & \\
 & u_B^{\sigma'} = u_B^{\tau'} &
 \end{array}$$

commutes as well; so we may set the image $(g \odot f)(\psi^\tau \odot \psi^\sigma)$ of $\psi^\tau \odot \psi^\sigma$ as

$$(\nu^\sigma \circ f(\psi^\sigma) \circ (\varphi^\sigma)^{-1}) \odot (\nu^\tau \circ g(\psi^\tau) \circ (\varphi^\tau)^{-1}) : (g \odot f)(x^\tau \odot x^\sigma) \cong_{\tau' \odot \sigma'} (g \odot f)(y^\tau \odot y^\sigma)$$

and as Proposition 7.4.4 entails that φ^σ is uniquely determined by $f(\psi^\sigma)$ and similarly for φ^τ , ν^σ and ν^τ , it follows that this construction is monotone.

Positivity. Write $y^{\tau'} \odot y^{\sigma'} = (g \odot f)(x^\tau \odot x^\sigma)$. We have

$$\begin{array}{ccc}
 x^\sigma & \xrightarrow{f} & f(x^\sigma) \xrightarrow{\varphi^{\sigma'}} y^{\sigma'} \\
 \partial_\sigma \downarrow & & \downarrow \partial_{\sigma'} \\
 \cdot & \xrightarrow{f[x^\sigma]_A \vdash f[x^\sigma]_B} & \cdot \xrightarrow{\varphi_A^{\sigma'} \vdash \varphi_B^{\sigma'}} \cdot
 \end{array}
 \quad
 \begin{array}{ccc}
 x^\tau & \xrightarrow{f} & g(x^\tau) \xrightarrow{\varphi^{\tau'}} y^{\tau'} \\
 \partial_\tau \downarrow & & \downarrow \partial_{\tau'} \\
 \cdot & \xrightarrow{f[x^\tau]_A \vdash f[x^\tau]_B} & \cdot \xrightarrow{\varphi_A^{\tau'} \vdash \varphi_B^{\tau'}} \cdot
 \end{array}$$

where in each case the left hand square commutes by definition of positive morphisms, and $\varphi^{\sigma'}$, $\varphi^{\tau'}$ provided by Proposition 7.4.4 – which also ensure that $\varphi_A^{\sigma'} \in \mathcal{S}_-(A)$ and $\varphi_C^{\tau'} \in \mathcal{S}_+(C)$. Now from the above, by definition and monotonicity of the constructions,

$$\begin{array}{ccc}
 x^\tau \odot x^\sigma & \xrightarrow{g \odot f} & (g \odot f)(x^\tau \odot x^\sigma) \\
 \partial_\sigma \downarrow & & \downarrow \partial_{\tau' \odot \sigma'} \\
 \cdot & \xrightarrow{\varphi_A^{\sigma'} \circ f[x^\sigma]_A \vdash \varphi_C^{\tau'} \circ g[x^\tau]_C} & \cdot
 \end{array} \tag{7.7}$$

commutes as well, and $\varphi_A^{\sigma'} \circ f[x^\sigma]_A \in \mathcal{S}_-(A)$ and $\varphi_C^{\tau'} \circ g[x^\tau]_C \in \mathcal{S}_+(C)$ as required.

Uniqueness. Direct consequence of Lemma 6.1.12. \square

In particular, this entails that thin equivalence is a congruence, *i.e.* if $\sigma \approx \sigma' : A \vdash B$ and $\tau \approx \tau' : B \vdash C$, then $\tau \odot \sigma \approx \tau' \odot \sigma' : A \vdash C$. This will be useful when, later on, we truncate our bicategory **TCG** into a \sim -category – see Section 7.4.4.

We prove one final lemma on horizontal composition:

Lemma 7.4.9. *Consider $f : \sigma \rightarrow \sigma'$ and $g : \tau \rightarrow \tau'$ positive morphisms.*

Then, for all minimal causally compatible triple $(x^\sigma, \theta, x^\tau)$, we have:

$$(g \odot f)(x^\tau \odot_\theta x^\sigma) = g(x^\tau) \odot_{g[x^\tau]_B \circ \theta \circ f[x^\sigma]_B^{-1}} f(x^\sigma).$$

Proof. Let us write $y^\tau \odot y^\sigma = x^\tau \odot_\theta x^\sigma$ and $z^\tau \odot z^\sigma = (g \odot f)(y^\tau \odot y^\sigma)$; by definition of $x^\tau \odot_\theta x^\sigma$ and of $(g \odot f)(y^\tau \odot y^\sigma)$, we have symmetries $\varphi^\sigma : x^\sigma \cong_\sigma y^\sigma$, $\varphi^\tau : x^\tau \cong_\tau y^\tau$, $\nu^{\sigma'} : f(y^\sigma) \cong_{\sigma'} z^{\sigma'}$ and $\nu^{\tau'} : g(y^\tau) \cong_{\tau'} z^{\tau'}$, satisfying the triangles:

$$\begin{array}{ccc} x_B^\sigma & \xrightarrow{\theta} & x_B^\tau \\ & \searrow \varphi_B^\sigma & \swarrow \varphi_B^\tau \\ & y_B & \end{array} \quad \begin{array}{ccc} f(y^\sigma)_B & \xrightarrow{g[y^\tau]_B \circ f[y^\sigma]_B^{-1}} & g(y^\tau)_B \\ & \searrow \nu_B^{\sigma'} & \swarrow \nu_B^{\tau'} \\ & z_B & \end{array}$$

It follows that the following diagram commutes as well:

$$\begin{array}{ccccc} f(x^\sigma)_B & \xrightarrow{f[x^\sigma]_B^{-1}} & x_B^\sigma & \xrightarrow{\theta} & x_B^\tau & \xrightarrow{g[x^\tau]_B} & g(x^\tau)_B \\ f(\varphi^\sigma)_B \downarrow & & \searrow \varphi_B^\sigma & & \swarrow \varphi_B^\tau & & \downarrow g(\varphi^\tau)_B \\ f(y^\sigma)_B & \xrightarrow{f[y^\sigma]_B^{-1}} & y_B & \xrightarrow{g[y^\sigma]_B} & g(y^\tau)_B & & \\ & & \searrow \nu_B^{\sigma'} & & \swarrow \nu_B^{\tau'} & & \\ & & z_B & & & & \end{array}$$

which proves our equality by uniqueness of Proposition 7.4.4. \square

7.4.3 Bicategorical Laws

Proving the bicategorical laws requires one final technical development, establishing an associativity property of synchronization up to symmetry.

Ternary synchronization up to symmetry. Fix three strategies

$$\sigma : A \vdash B, \quad \tau : B \vdash C, \quad \delta : C \vdash D$$

with configurations $x^\sigma \in \mathcal{C}(\sigma)$, $x^\tau \in \mathcal{C}(\tau)$, $x^\delta \in \mathcal{C}(\delta)$, and mediating symmetries

$$\theta : x_B^\sigma \cong_B x_B^\tau, \quad \omega : x_C^\tau \cong_C x_C^\delta.$$

What can we say about their synchronizations up to symmetry? Assuming $(x^\sigma, \theta, x^\tau)$ is minimal causally compatible, by Proposition 7.4.4 there are unique

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

such that $\varphi_A^\sigma \in \mathcal{S}_-(A)$, $\varphi_C^\tau \in \mathcal{S}_+(C)$ and the following diagram commutes

$$\begin{array}{ccc} x_B^\sigma & \xrightarrow{\theta} & x_B^\tau \\ & \searrow \varphi_B^\sigma & \swarrow \varphi_B^\tau \\ & y_B & \end{array}$$

and by Corollary 7.4.5, y^σ, y^τ are minimal causally compatible. With these notations:

Lemma 7.4.10. Assume $(x^\sigma, \theta, x^\tau)$ and $(x^\tau \odot_\theta x^\sigma, \omega \circ (\varphi_C^\tau)^{-1}, x^\delta)$ are min. caus. comp. Then, so is $(x^\tau, \omega, x^\delta)$.

Proof. In that case, then by Proposition 7.4.4 there are unique

$$\Psi^{\sigma, \tau} : y^\tau \odot y^\sigma \cong_{\tau \odot \sigma} u^\tau \odot u^\sigma, \quad \Psi^\delta : x^\delta \cong_\delta u^\delta$$

such that $\Psi_A^{\sigma, \tau} \in \mathcal{S}_-(A)$, $\Psi_D^\delta \in \mathcal{S}_+(D)$ and the following diagram commutes:

$$\begin{array}{ccc} (y^\tau \odot y^\sigma)_C & \xrightarrow{\omega \circ (\varphi_C^\tau)^{-1}} & x_C^\delta \\ \Psi_C^{\sigma, \tau} \searrow & & \swarrow \Psi_C^\delta \\ & u_C & \end{array}$$

recalling that $\Psi^{\sigma, \tau} : y^\tau \odot x^\sigma \cong_{\tau \odot \sigma} u^\tau \odot u^\sigma$ consists in $\Psi^\sigma : x^\sigma \cong_\sigma u^\sigma$ and $\Psi^\tau : x^\tau \cong_\tau u^\tau$ matching on B . By Corollary 7.4.5, $u^\tau \odot u^\sigma$ and u^δ are minimal causally compatible. By Lemma 6.2.24, it follows that u^τ, u^δ are minimal causally compatible. But then

$$\Psi^\tau \circ \varphi^\tau : x^\tau \cong_\tau u^\tau, \quad \Psi^\delta : x^\delta \cong_\delta u^\delta$$

are order-isomorphisms such that the following triangle commutes:

$$\begin{array}{ccc} x_C^\tau & \xrightarrow{\omega} & x_C^\delta \\ (\Psi^\tau \circ \varphi^\tau)_C \searrow & & \swarrow \Psi_C^\delta \\ & u_C & \end{array}$$

Following these isomorphisms, it is a direct verification to establish causal compatibility of $(x^\tau, \omega, x^\delta)$ from that of (u^τ, u^δ) . \square

Now, assuming that the conditions of this lemma are satisfied, we know that $(x^\tau, \omega, x^\delta)$ are causally compatible. This means, again by Proposition 7.4.4, that there are

$$\psi^\tau : x^\tau \cong_\tau z^\tau, \quad \psi^\delta : x^\delta \cong_\delta z^\delta$$

such that the following triangle commutes:

$$\begin{array}{ccc} x_B^\tau & \xrightarrow{\omega} & x_B^\delta \\ \psi_B^\tau \searrow & & \swarrow \psi_B^\delta \\ & z_B & \end{array}$$

Using these notations, we can then prove the following lemma:

Lemma 7.4.11. *Assume $(x^\sigma, \theta, x^\tau)$ and $(x^\tau \odot_\theta x^\sigma, \omega \circ (\varphi_C^\tau)^{-1}, x^\delta)$ are min. caus. comp.*

Then, so is $(x^\sigma, \psi_B^\tau \circ \theta, x^\delta \odot_\omega x^\tau)$.

Proof. By Corollary 7.4.5, u^σ and $u^\delta \odot u^\tau$ are minimal causally compatible. As for the previous lemma, the idea is to provide symmetries making minimal causal compatibility of $(x^\sigma, \psi_B^\tau \circ \theta, x^\delta \odot_\omega x^\tau)$ boil down to that of $(u^\sigma, u^\delta \odot u^\tau)$. First there is:

$$\Psi^\sigma \circ \varphi^\sigma : x^\sigma \cong_\sigma u^\sigma,$$

and likewise we have two symmetries

$$\Psi^\tau \circ \varphi^\tau \circ (\psi^\tau)^{-1} : z^\tau \cong_\tau u^\tau, \quad \Psi^\delta \circ (\psi^\delta)^{-1} : z^\delta \cong_\delta u^\delta$$

which can be checked to have the same projection on C . So their synchronization yields

$$\Phi = (\Psi^\delta \circ (\psi^\delta)^{-1}) \odot (\Psi^\tau \circ \varphi^\tau \circ (\psi^\tau)^{-1}) : x^\delta \odot_\omega x^\tau \cong_{\delta \odot \tau} u^\delta \odot u^\tau.$$

Finally, a direct verification shows that the triangle below commutes:

$$\begin{array}{ccc} x_B^\sigma & \xrightarrow{\psi_B^\tau \circ \theta} & (x^\delta \odot_\omega x^\tau)_B \\ & \searrow \Psi^\sigma \circ \varphi^\sigma & \swarrow \Phi_B \\ & u_B & \end{array} \quad (7.8)$$

via which the announced min. caus. comp. boils down to that of $(u^\sigma, u^\delta \odot u^\tau)$. \square

At this point we are equipped to state the sought associativity property:

Lemma 7.4.12. *If $(x^\sigma, \theta, x^\tau)$ and $(x^\tau \odot_\theta x^\sigma, \omega \circ (\varphi_C^\tau)^{-1}, x^\delta)$ are min. caus. comp.,*

$$x^\delta \odot_{\omega \circ (\varphi_C^\tau)^{-1}} (x^\tau \odot_\theta x^\sigma) = a_{\sigma, \tau, \delta}((x^\delta \odot_\omega x^\tau) \odot_{\psi_B^\tau \circ \theta} x^\sigma).$$

Proof. Keeping the notations above, $x^\delta \odot_{\omega \circ (\varphi_C^\tau)^{-1}} (x^\tau \odot_\theta x^\sigma) = u^\delta \odot (u^\tau \odot u^\sigma)$. Now,

$$\Psi : x^\delta \odot_\omega x^\tau \cong_{\delta \odot \tau} u^\delta \odot u^\tau, \quad \Psi^\sigma \circ \varphi^\sigma : x^\sigma \cong_\sigma u^\sigma$$

are symmetries satisfying (7.8); and it is direct from the definition that $\Psi_A \in \mathcal{S}_-(A)$ and $(\Psi^\sigma \circ \varphi^\sigma)_C \in \mathcal{S}_+(C)$. Consequently, by the uniqueness clause of Proposition 7.4.4,

$$(x^\delta \odot_\omega x^\tau) \odot_{\psi_B^\tau \circ \theta} x^\sigma = (u^\delta \odot u^\tau) \odot u^\sigma$$

and $a_{\sigma, \tau, \delta}((u^\delta \odot u^\tau) \odot u^\sigma) = u^\delta \odot (u^\tau \odot u^\sigma)$ as required. \square

A bicategory. At last, we have the ingredients to conclude:

Theorem 7.4.13. *There is a bicategory **TCG** with games as objects, strategies as morphisms, and positive morphisms as 2-cells.*

Proof. Identities are defined in Proposition 7.3.15, and composition in Proposition 7.3.1. Each strategy has an identity 2-cell (the identity positive isomorphism). Vertical composition of positive morphisms is composition in **ESS**, and horizontal composition is Proposition 7.4.8. Associators are defined in Proposition 7.4.1 and unitors in Proposition 7.4.2. It remains to prove the following laws:

Functoriality of horizontal composition. For functoriality, we compute

$$\begin{aligned}
& (g' \odot f') \circ (g \odot f)(x^\tau \odot x^\sigma) \\
&= (g' \odot f')(g(x^\tau) \odot_{g[x^\tau]_B \circ f[x^\sigma]_B^{-1}} f(x^\sigma)) \\
&= g'(g(x^\tau)) \odot_{g'[g(x^\tau)]_B \odot g[x^\tau]_B \odot f[x^\sigma]_B^{-1} \circ f'[f(x^\sigma)]_B^{-1}} f'(f(x^\sigma)) \\
&= (g' \circ g)(x^\tau) \odot_{(g' \circ g)[x^\tau]_B \circ (f' \circ f)[x^\sigma]_B^{-1}} (f' \circ f)(x^\sigma) \\
&= ((g' \circ g) \odot (f' \circ f))(x^\tau \odot x^\sigma)
\end{aligned}$$

using definition of $g \odot f$, Lemma 7.4.9, definition of the preservation of symmetry of $g' \circ g$ and $f' \circ f$, and definition of $(g' \circ g) \odot (f' \circ f)$. By Lemma 6.1.12, we conclude that $(g' \odot f') \circ (g \odot f) = (g' \circ g) \odot (f' \circ f)$ as required. Likewise, it is straightforward using the same lemmas that horizontal composition preserves identity positive morphisms.

Naturality of unitors. We first show the naturality of $l_\sigma : \mathfrak{C}_B \odot \sigma \cong \sigma$ in σ , i.e. that for all $\sigma, \sigma' : A \vdash B$ and $f : \sigma \Rightarrow \sigma'$ a positive morphism, we have $f \circ l_\sigma = l_{\sigma'} \circ (\mathfrak{C}_B \odot f)$, and by Corollary 7.2.12 it suffices to prove the equality for $+$ -covered configurations. By Lemma 6.4.4, $+$ -covered configurations in $\mathfrak{C}_B \odot \sigma$ have form $\mathfrak{C}_{x_B^\sigma} \odot x^\sigma \in \mathcal{C}^+(\mathfrak{C}_B \odot \sigma)$, and by definition of horizontal composition (in Proposition 7.4.8),

$$(\mathfrak{C}_B \odot f)(\mathfrak{C}_{x_B^\sigma} \odot x^\sigma) = \mathfrak{C}_{x_B^\sigma} \odot_{f[x^\sigma]_B^{-1}} f(x^\sigma),$$

but that must be $\mathfrak{C}_{f(x^\sigma)_B} \odot f(x^\sigma) \in \mathcal{C}^+(\mathfrak{C}_B \odot \sigma')$ as it is easily proved to satisfy the properties in Proposition 7.4.4 defining it uniquely.

Now, from the characterization of the action of unitors in Proposition 7.4.2,

$$l_{\sigma'}(\mathfrak{C}_{f(x^\sigma)_B} \odot f(x^\sigma)) = f(x^\sigma)$$

which is also $f(l_\sigma(\mathfrak{C}_{x_B^\sigma} \odot x^\sigma)) = f(x^\sigma)$. The proof for r_σ is symmetric.

Naturality of associators. Consider now $f : \sigma \Rightarrow \sigma' : A \vdash B$, $g : \tau \Rightarrow \tau' : B \vdash C$ and $h : \delta \Rightarrow \delta' : C \vdash D$ positive morphisms, and calculate:

$$\begin{aligned}
& (h \odot (g \odot f)) \circ a_{\sigma, \tau, \delta}((x^\delta \odot x^\tau) \odot x^\sigma) \\
&= (h \odot (g \odot f))(x^\delta \odot (x^\tau \odot x^\sigma)) \\
&= h(x^\delta) \odot_{h[x^\delta]_C \circ (g \odot f)[x^\tau \odot x^\sigma]_C^{-1}} (g \odot f)(x^\tau \odot x^\sigma) \\
&= h(x^\delta) \odot_{h[x^\delta]_C \circ (g \odot f)[x^\tau \odot x^\sigma]_C^{-1}} (g(x^\tau) \odot_{g[x^\tau]_B \circ h[x^\sigma]_B^{-1}} f(x^\sigma)) \\
&= h(x^\delta) \odot_{h[x^\delta]_C \circ g[x^\tau]_C^{-1} \circ (g \odot f)[x^\tau \odot x^\sigma]_C^{-1}} (g(x^\tau) \odot_{g[x^\tau]_B \circ f[x^\sigma]_B^{-1}} f(x^\sigma)) \\
&= a_{\sigma', \tau', \delta'}((h(x^\delta) \odot_{h[x^\delta]_C \circ g[x^\tau]_C^{-1}} g(x^\tau)) \odot_{\psi_B^\tau \circ g[x^\tau]_B \circ f[x^\sigma]_B^{-1}} f(x^\sigma)) \\
&= a_{\sigma', \tau', \delta'}((h(x^\delta) \odot_{h[x^\delta]_C \circ g[x^\tau]_C^{-1}} g(x^\tau)) \odot_{(h \odot g)[x^\delta \odot x^\tau]_B \circ f[x^\sigma]_B^{-1}} f(x^\sigma)) \\
&= a_{\sigma', \tau', \delta'}((h \odot g)(x^\tau \odot x^\sigma) \odot_{(h \odot g)[x^\delta \odot x^\tau]_B \circ f[x^\sigma]_B^{-1}} f(x^\sigma)) \\
&= a_{\sigma', \tau', \delta'}(((h \odot g) \odot f)((x^\delta \odot x^\tau) \odot x^\sigma))
\end{aligned}$$

using Lemma 7.4.9 twice, then the witness exposed in (7.7) for the fact that $g \odot f$ is a positive morphism, then Lemma 7.4.12 (using the same notations for φ and ψ), then

again the witness in (7.7) for the fact that hog is a positive morphism, concluding with two applications of Lemma 7.4.9 – which finally ends the proof. \square

Though we have constructed a bicategory **TCG**, most of this monograph relies on a simpler structure ignoring coherence laws, sufficient for many semantic purposes.

7.4.4 A \sim -category

To wrap up this chapter, we introduce a simplified higher-dimensional structure.

In order to use **TCG** for semantic purposes, we need much additional categorical structure: in particular, we must at least have a cartesian closed structure in order to interpret the simply-typed λ -calculus. Following Theorem 7.4.13, the reader may naturally expect us to aim towards constructing the *bicategorical* version of this structure. Such a construction is described by Paquet in [Paquet, 2020]. But ultimately, the explicit managing of 2-cells does not appear particularly useful for the semantic purposes of this monograph so we opted to not develop this higher-dimensional structure further.

Alternatively, we could quotient **TCG** to a category having as morphisms equivalence classes of strategies with respect to positive isomorphism (as done for instance in AJM games [Abramsky et al., 2000]). But this is actually harmful: it is methodologically better for the interpretation to yield actual concrete strategies, rather than equivalence classes (for instance, recursion is obtained via completeness properties of an ordering on strategies that does not seem to lift to equivalence classes). We also want to avoid the vagueness often encountered in switching implicitly between equivalence classes and representatives.

So we shall work neither with categories nor bicategories but with \sim -categories:

Definition 7.4.14. A (small) \sim -category C consists in a set of objects C_0 ; for each A, B , a set of morphisms $C(A, B)$ with an equivalence relation \sim ; a composition operation

$$(- \circ -) : C(B, C) \times C(A, B) \rightarrow C(A, C)$$

for all $A, B, C \in C_0$; an identity morphism $\text{id}_A \in C(A, A)$ for all $A \in C_0$, subject to:

- associativity: for all $f \in C(A, B), g \in C(B, C), h \in C(C, D)$, $(h \circ g) \circ f \sim h \circ (g \circ f)$,
- identity: for all $f \in C(A, B)$, $\text{id}_B \circ f \sim f \circ \text{id}_A \sim f$,
- congruence: for all $f \sim f' \in C(A, B)$, $g \sim g' \in C(B, C)$, $g' \circ f' \sim g \circ f$.

In other words, a \sim -category is a special case of a bicategory where any two morphisms are related by at most one (invertible) 2-cell. Any bicategory yields a \sim -category where $f \sim g$ iff f and g are related by an isomorphism.

Using \sim -categories avoids quotienting morphisms, while making all coherence laws hold vacuously. All categorical notions extend transparently to \sim -categories, with equations holding up to \sim rather than equality. For instance, functors between \sim -categories respect \sim , and preservation of identities and composition only holds up to \sim . In the

sequel, we shall use several standard categorical notions in the \sim -categorical world without explicitly spelling them out – hopefully this causes no confusion.

From now on, we regard **TCG** as a \sim -category:

Proposition 7.4.15. *There is **TCG**, a \sim -category with tcgs as objects, strategies as morphisms, and equivalence relation \approx the thin equivalence, from Corollary 7.2.9.*

7.5 History and Related Work

Genesis of TCG. I arrived as a postdoc with Glynn Winskel in Cambridge in September 2011. After participating in a first project to enrich concurrent games with winning conditions [Clairambault et al., 2012], my main task was the supposedly simple objective to extend **CG** with *symmetry* [Winskel, 2007], in particular so that concurrent games could handle replication as in AJM games [Abramsky et al., 2000]. What was initially expected to be relatively straightforward proved to be an endeavour full of surprises, for which we were soon joined by a research intern, Simon Castellan.

As discussed earlier, there are two ways to handle uniformity in (copy indices-based) game semantics. The *fat* way, and the *thin* way. The fat approach to uniformity consists in forcing strategies to be saturated: if a strategy can play one copy index, it must play non-deterministically *all* copy indices. In contrast, in the thin approach, a strategy in a given state will play *one canonically chosen* copy index. Uniformity then becomes an interactive property: any change in Opponent’s copy indices causes a subsequent change in Player’s copy indices. This *thin/fat* distinction is not about concurrent games at all, but present as soon as there are strategies with copy indices: the original AJM games are thin [Abramsky et al., 2000], but they can be made fat – the very first fat model is [Baillot et al., 1997a]. The fat approach is simpler; however, it conflates uniformity with non-determinism and also is not a compatible extension of the symmetry-free layer as the basic mechanisms (identity and composition) must be changed.

In developing concurrent games with symmetry, we found that we could not make the thin approach work as we lacked a good way to describe the interactive reindexing process involved in uniformity. So the first published version of concurrent games with symmetry is fat³ [Castellan et al., 2014]⁴. However, saturation made the model difficult to use, and I kept looking for a thin version, whose maturing was a long process.

Thin concurrent games first appeared in [Castellan et al., 2015] (the paper was mainly about parallel innocence – see Chapter 10. The core of this setting is the one presented in this chapter, though with many crucial changes and simplifications. Many significant improvements (including the very definition of a tcg!) are due to Simon Castellan, developed during his PhD thesis; others came while writing the journal version detailing thin concurrent games [Castellan et al., 2019], and others still appeared since.

³In fact, the very first version of concurrent games with symmetry is thin, and only appears in Simon Castellan’s (unpublished) internship report – however, we could only make this work adopting a coarse bisimulation-like equivalence relation on strategies, whereas we wanted an isomorphism.

⁴Not very well-known is the fact that this paper also contains the very first non-deterministic generalization of Hyland-Ong innocence!

Further comments. This chapter mainly presents constructions first introduced in [Castellan et al., 2015] and improved in [Castellan et al., 2019]. However, the latter uses results from [Castellan et al., 2017a], making for a somewhat intimidating chain of dependencies. In contrast, the present chapter is entirely self-contained and constructs thin concurrent games from the ground, with full details.

Finally, this chapter builds a bicategory, whereas the papers cited above only construct a category up to positive isomorphisms. This is not a new result: in his PhD thesis, Paquet proves that thin concurrent games form a bicategory (actually a symmetric monoidal closed bicategory [Paquet, 2020]). However, Paquet does not detail the naturality of unitors and associators; so our proof complements his.

It seems fair to say that **TCG** is the most expressive and fine-grained intensional model on the market. Much of our further work builds on it, and there is a lot left to explore. I consider this to be one of the main contributions in this line of work.

Related work. The core idea behind **TCG** is splitting symmetries into *positive* and *negative* symmetries: it is thanks to this separation that we are able to factor out Opponent's reindexings when expressing uniformity and when comparing strategies. As mentioned earlier, this idea was first developed by Melliès as his approach to uniformity in asynchronous games [Melliès, 2003]. There, he considers an *asynchronous game* with *two* groups acting on plays, one for Player reindexings and one for Opponent reindexings; uniformity is there specified as a *bi-invariance* property with respect to these group actions, analogous in spirit to Lemma 7.2.6 – note that recently, Paquet showed that Melliès' mechanism could also be applied in the setting of concurrent games [Paquet, 2023], obtaining a different, more global, notion of uniformity which differs from the present one for subtle reasons.

In any case, this polarized decomposition of symmetry is clearly a fundamental structure that deserves further exploration. One such recent exploration is a joint work with Simon Forest, in which we build a bicategory of *thin spans of groupoids* sharing much of the structure of **TCG**, but completely statically [Clairambault and Forest, 2023], *i.e.* without time or causality. There, an *object* is a groupoid with two sub-groupoids, the *positive* and the *negative* ones. Morphisms are spans of groupoids selected by a biorthogonality relation ensuring that their composition satisfies a universal property analogous to Proposition 7.3.1, but without causal compatibility. Overall, we get a proof-relevant relational model sharing much of the flavour of **TCG**, but avoiding the combinatorics of concurrent strategies and their composition.

Chapter 8

Constructing Games and Strategies

The previous chapter describes the bare bones of thin concurrent games: namely, the composition mechanism and the tools to handle symmetry. In the present chapter, we describe the constructions on games and strategies that will be exploited throughout this monograph, aiming to establish categorical structures useful for semantics.

In practice, we will build from **TCG** two radically different categorical structures.

TCG *per se* is an *unpolarized* category of games, in that games have no assigned polarity: in a game A , both players may be able to start. This is in contrast with most notions of games used in game semantics, which focus on *negative* games where only Opponent starts (for call-by-name languages), or on *positive* games where only Player starts (for call-by-value languages). A widespread urban legend has it that such non-polarized game settings suffer from a lack of associativity, due to the infamous *Blass problem* [Abramsky, 2003]. In fact, this phenomenon is very specific to Blass games, and Proposition 6.2.25 ensures that composition in **TCG** is indeed associative up to iso. But **TCG** is certainly not the first unpolarized category of games, the pioneering example is Joyal’s category of Conway games [Joyal, 1977]. As pioneered by Joyal, the expected categorical structure of such an unpolarized category of games is that of a *compact closed category*, i.e. a degenerate model of MLL where $\otimes = \wp$. Our first objective in this chapter will be to show that **TCG** is a compact closed \sim -category.

To prepare the ground for semantics, we must then switch to *polarized* games – *negative*, since the languages considered here are call-by-name. We shall equip the negative sub- \sim -category of **TCG** with (a variation of) the structure of a *Seely category* [Melliès, 2009]. Seely categories are a standard model of ILL, and thus of the simply-typed λ -calculus; and we shall use this structure as the basis for our interpretations.

8.1 A Compact closed \sim -Category

Compact closed categories are amongst the simplest categorical structures that exhibit higher-order behaviour. They are essentially those degenerated models of multiplicative linear logic where the \otimes and its dual \wp coincide – intuitively, they resemble proof nets without the correctness criterion. Like Conway games [Joyal, 1977], categories of unpolarized games are often compact closed, and this structure underlies much of the additional structure appearing in polarized categories of games, as explored by Melliès [Melliès, 2004b]. It is thus natural to start there.

Recall that a compact closed category is a symmetric monoidal category \mathcal{C} where each object A has a (necessarily unique up to iso) *dual object* A^* . This means we have

$$\begin{array}{lll} \text{the unit, } \eta_A & : & I \rightarrow A \otimes A^* \\ \text{the co-unit, } \epsilon_A & : & A^* \otimes A \rightarrow I \end{array}$$

two morphisms for all A , respectively akin to the *axiom* and *cut* of proof nets, satisfying two coherence axioms [Kelly and Laplaza, 1980]. A compact closed \sim -category is defined likewise, with all laws holding up to \sim instead of equality.

8.1.1 A bifunctor

For thin concurrent games we start with the first component, a \sim -bifunctor:

$$(- \parallel -) : \mathbf{TCG} \times \mathbf{TCG} \rightarrow \mathbf{TCG}.$$

Notice that we use \parallel for the tensor of the compact closed structure, as we shall reserve \otimes for a later construction in the context of negative games.

Definition. On tcgs, the construction \parallel is defined in Section 7.1.4. If A and B are tcgs, then $A \parallel B$ is thought of as the game formed of A and B played in parallel. Note that though the games are independent, a strategy on $A \parallel B$ may very well impose crossed immediate causal links between the two components!

On strategies, the action of \parallel is given by the following:

Proposition 8.1.1. Consider A, B, C, D tcgs, and $\sigma : A \vdash B$, $\tau : C \vdash D$ strategies. Then, there is a strategy $\sigma \parallel \tau : A \parallel C \vdash B \parallel D$, unique up to iso, s.t. there are

$$\begin{array}{lll} (- \parallel -) & : & \mathcal{C}^+(\sigma) \times \mathcal{C}^+(\tau) \simeq \mathcal{C}^+(\sigma \parallel \tau) \\ (- \parallel -) & : & \mathcal{S}^+(\sigma) \times \mathcal{S}^+(\tau) \simeq \mathcal{S}^+(\sigma \parallel \tau) \end{array}$$

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

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

Moreover, $(- \parallel -)$ preserves \approx (a property we call congruence).

Proof. Existence. As the notation suggests, $\sigma \parallel \tau$ has ess the simple parallel composition of σ and τ . Its display map is set as:

$$\begin{aligned} \partial_{\sigma \parallel \tau}(1, s) &= \begin{cases} (1, (1, a)) & \text{if } \partial_\sigma(s) = (1, a) \\ (2, (1, b)) & \text{if } \partial_\sigma(s) = (2, b) \end{cases} \\ \partial_{\sigma \parallel \tau}(2, t) &= \begin{cases} (1, (2, c)) & \text{if } \partial_\tau(t) = (1, c) \\ (2, (2, d)) & \text{if } \partial_\tau(t) = (2, d), \end{cases} \end{aligned}$$

we omit the routine verification that this yields a strategy. All additional conditions follow as in Lemma 7.1.22 along with routine verifications.

Uniqueness. Straightforward consequence of Corollary 7.2.13.

Congruence. Immediate from the definition. \square

Functoriality. Proposition 8.1.1 makes functoriality easy. First:

Lemma 8.1.2. *For any tcgs A and B , we have $\mathfrak{C}_{A \parallel B} \approx \mathfrak{C}_A \parallel \mathfrak{C}_B$.*

Proof. We calculate:

$$\begin{aligned} \mathcal{C}^+(\mathfrak{C}_A \parallel \mathfrak{C}_B) &\cong \mathcal{C}^+(\mathfrak{C}_A) \times \mathcal{C}^+(\mathfrak{C}_B) \\ &= \{x_A \parallel x_A \mid x_A \in \mathcal{C}(A)\} \times \{x_B \parallel x_B \mid x_B \in \mathcal{C}(B)\} \\ &\cong \{(x_A \parallel x_B) \parallel (x_A \parallel x_B) \mid x_A \parallel x_B \in \mathcal{C}(A \parallel B)\} \\ &= \mathcal{C}^+(\mathfrak{C}_{A \parallel B}) \end{aligned}$$

using Proposition 8.1.1, then Lemma 6.4.4, then the obvious bijection, then Lemma 6.4.4 again – all these are order-isos, which preserve display maps in the obvious sense.

The same reasoning applies to symmetries compatibly, using Proposition 8.1.1 and Lemma 7.3.14. Finally, the seeked isomorphism follows from Corollary 7.2.13. \square

For preservation of composition, consider four strategies

$$\sigma_1 : A_1 \vdash B_1, \quad \sigma_2 : A_2 \vdash B_2, \quad \tau_1 : B_1 \vdash C_1, \quad \tau_2 : B_2 \vdash C_2.$$

The only additional ingredient needed is the following lemma:

Lemma 8.1.3. *Take $x^{\sigma_1} \in \mathcal{C}^+(\sigma_1)$, $x^{\sigma_2} \in \mathcal{C}^+(\sigma_2)$, $x^{\tau_1} \in \mathcal{C}^+(\tau_1)$ and $x^{\tau_2} \in \mathcal{C}^+(\tau_2)$.*

Then, $x^{\sigma_1} \parallel x^{\sigma_2} \in \mathcal{C}^+(\sigma_1 \parallel \sigma_2)$ and $x^{\tau_1} \parallel x^{\tau_2} \in \mathcal{C}^+(\tau_1 \parallel \tau_2)$ are causally compatible iff (1) x^{σ_1} and x^{τ_1} are causally compatible; and (2) x^{σ_2} and x^{τ_2} likewise.

Proof. A routine verification, left to the reader. \square

Now, we can finally conclude:

Proposition 8.1.4. *Parallel composition extends to a \sim -bifunctor:*

$$(- \parallel -) : \mathbf{TCG} \times \mathbf{TCG} \rightarrow \mathbf{TCG}.$$

Proof. Preservation of copycat is by Lemma 8.1.2. For composition, we calculate:

$$\begin{aligned}
& \mathcal{C}^+((\tau_1 \parallel \tau_2) \odot (\sigma_1 \parallel \sigma_2)) \\
& \cong \{(x^{\tau_1} \parallel x^{\tau_2}, x^{\sigma_1} \parallel x^{\sigma_2}) \in \mathcal{C}^+(\tau_1 \parallel \tau_2) \times \mathcal{C}^+(\sigma_1 \parallel \sigma_2) \mid \text{caus. comp.}\} \\
& \cong \{(x^{\tau_1}, x^{\sigma_1}) \in \mathcal{C}^+(\tau_1) \times \mathcal{C}^+(\sigma_1) \mid \text{caus. comp.}\} \times \\
& \quad \{(x^{\tau_2}, x^{\sigma_2}) \in \mathcal{C}^+(\tau_2) \times \mathcal{C}^+(\sigma_2) \mid \text{caus. comp.}\} \\
& \cong \mathcal{C}^+(\tau_1 \odot \sigma_1) \times \mathcal{C}^+(\tau_2 \odot \sigma_2)
\end{aligned}$$

using Proposition 7.3.1, then Proposition 8.1.1 along with Lemma 8.1.3, then Proposition 7.3.1 again. All these are order-isos, preserve display maps, and the same reasoning applies to symmetries – so the desired isomorphism follows from Corollary 7.2.13. \square

8.1.2 Lifting Maps to Strategies

Besides this bifunctor, we need structural morphisms for the monoidal structure.

We introduce a method to construct them systematically. The idea is that all these morphisms are direct variations of copycat: they play copycat following a certain isomorphism of games. Accordingly, we show how to lift certain maps to strategies.

Renamings. We shall lift more than just isos: “liftable” maps are called *renamings*:

Definition 8.1.5. A *renaming* from game A to B is a function $f : |A| \rightarrow |B|$ satisfying:

$$\begin{aligned}
& \text{validity: } \forall x \in \mathcal{C}(A), f x \in \mathcal{C}(B) \\
& \text{local injectivity: } \forall a_1, a_2 \in x \in \mathcal{C}(A), f a_1 = f a_2 \Rightarrow a_1 = a_2 \\
& \text{pol-preserving: } \forall a \in A, \text{pol}_B(f a) = \text{pol}_A(a) \\
& \text{sym-preserving: } \forall \theta \in \mathcal{S}(A) (\text{resp. } \mathcal{S}_+(A), \mathcal{S}_-(A)), \\
& \quad f \theta \in \mathcal{S}(B), (\text{resp. } \mathcal{S}_+(B), \mathcal{S}_-(B)) \\
& \text{strong-receptivity: } \forall \theta \in \mathcal{S}(A), \forall f \theta \subseteq^- \varphi \in \mathcal{S}(B), \exists ! \theta' \subseteq^- \theta' \in \mathcal{S}(A) f \theta' = \varphi \\
& \text{courtesy: } \forall a_1 \rightarrow_A a_2, (\text{pol}_A(a_1) = + \vee \text{pol}_A(a_2) = -) \Rightarrow f a_1 \rightarrow_B f a_2.
\end{aligned}$$

We write **Ren** for the category of renamings, and $f : A \rightarrow B$ for a renaming.

Any isomorphism of games (*i.e.* an isomorphism preserving all structure) is a renaming. Renamings let us change the display of events of strategies, via the operation:

Definition 8.1.6. Consider $\sigma : A \vdash B$, with renamings $f : A^\perp \rightarrow A'^\perp$, $g : B \rightarrow B'$.

Then, the **redisplayed** strategy $g \cdot \sigma \cdot f$ has components the same as σ , except:

$$\partial_{g \cdot \sigma \cdot f} = (f \parallel g) \circ \partial_\sigma.$$

This is still a strategy – note the strategy itself is not really changed, only its display. Here, $f \parallel g$ refers to the monoidal structure of **ESS**, defined with $(f \parallel g)(1, a) = (1, f(a))$ and $(f \parallel g)(2, b) = (2, g(b))$. Note also that we can redisplay on one side of a strategy only, as in $g \cdot \sigma$: this means $g \cdot \sigma \cdot \text{id}_A$ – likewise, $\sigma \cdot f$ means $\text{id}_B \cdot \sigma \cdot f$.

Lifting. We motivated renamings by the *lifting* operation, which is:

Definition 8.1.7. Consider $f : A \rightarrow B$ a renaming.

Its *forward lifting* $\mathfrak{C}_f : A \vdash B$ is defined as $\mathfrak{C}_f = f \cdot \mathfrak{C}_A$.

Its *backward lifting* $\mathfrak{D}_f : B^\perp \vdash A^\perp$ is defined as $\mathfrak{D}_f = \mathfrak{C}_{A^\perp} \cdot f$.

This gives a simple definition of a strategy playing copycat alongside a certain map. This is a nice feature of concurrent games: such a construction is ubiquitous also in more traditional game semantics settings, but it is very rarely made explicit.

The fundamental property enjoyed by this construction is the **lifting lemma**:

Lemma 8.1.8. Consider $\sigma : A \vdash B$, and renamings $f : A^\perp \rightarrow A'^\perp$, $g : B \rightarrow B'$.

Then, $g \cdot \sigma \cdot f \approx \mathfrak{C}_g \odot \sigma \odot \mathfrak{D}_f$.

Proof. We show $g \cdot \sigma \approx \mathfrak{C}_g \odot \sigma$ – note $\sigma \cdot f \approx \sigma \odot \mathfrak{D}_f$ is symmetric, and the general case follows since $g \cdot \sigma \cdot f = g \cdot (\sigma \cdot f) = (g \cdot \sigma) \cdot f$ and \odot is associative up to \approx .

By definition, $\mathcal{C}^+(\mathfrak{C}_f) = \mathcal{C}^+(\mathfrak{C}_B) = \{x_B \parallel x_B \mid x_B \in \mathcal{C}(B)\}$ (Lemma 6.4.4), with

$$\partial_{\mathfrak{C}_g}(x_B \parallel x_B) = x_B \parallel g(x_B).$$

Thus a pair $(x^\sigma, x_B \parallel x_B)$ is causally compatible for σ and \mathfrak{C}_f iff it is causally compatible for σ and \mathfrak{C}_B . We may then reuse the unitor of Proposition 6.4.9 which yields $r_\sigma : \mathcal{C}^+(\sigma \odot \mathfrak{C}_f) \cong \mathcal{C}^+(\sigma)$. Computing the display maps, we have

$$\partial_{g \cdot \sigma}(x^\sigma) = g(x_B^\sigma) = g(x_B), \quad \partial_{\mathfrak{C}_g \odot \sigma}((x_B \parallel x_B) \odot x^\sigma) = g(x_B)$$

where $x_B^\sigma = x_B$ as the pair $(x^\sigma, x_B \parallel x_B)$ is matching. The same reasoning applies to symmetries using Lemma 7.3.14. By Corollary 7.2.13, we deduce the desired iso. \square

We deduce one last convenient property of lifting:

Proposition 8.1.9. There are two \sim -functors:

$$\mathfrak{C}_- : \mathbf{Ren} \rightarrow \mathbf{TCG}, \quad \mathfrak{D}_- : \mathbf{Ren}^{\text{op}} \rightarrow \mathbf{TCG},$$

such that if $f : A \rightarrow B$ is an iso, then \mathfrak{C}_f and \mathfrak{D}_{f^\perp} are inverses, and $\mathfrak{C}_{f^{-1}} \approx \mathfrak{D}_{f^\perp}$.

Proof. By Lemma 8.1.8 and a direct verification via Corollary 7.2.13. \square

Above, we write $f^\perp : A^\perp \rightarrow B^\perp$ for $f : A \rightarrow B$ defined as the same map as f . Note that if $f : A \rightarrow B$ is an iso, then $f^\perp : A^\perp \rightarrow B^\perp$ is an iso hence a renaming.

8.1.3 Compact Closed Structure

We may now directly apply this for the symmetric monoidal structure.

Symmetric monoidal structure. We need structural isomorphisms for **TCG**.

First, notice that **Ren** is a symmetric monoidal category. Indeed, for $f : A \rightarrow B$ and $f' : A' \rightarrow B'$ renamings, then $f \parallel f' : A \parallel A' \rightarrow B \parallel B'$ is a renaming. Furthermore, there are obvious isomorphisms of tcgs, so invertible renamings

$$\begin{array}{lll} \rho_A & : & A \parallel \mathbf{1} \rightarrow A \\ \lambda_A & : & \mathbf{1} \parallel A \rightarrow A \\ \alpha_{A,B,C} & : & (A \parallel B) \parallel C \rightarrow A \parallel (B \parallel C) \\ s_{A,B} & : & A \parallel B \rightarrow B \parallel A \end{array}$$

natural in A, B, C and satisfying the coherence laws of a symmetric monoidal category.

From this and the lifting operation introduced above, we prove:

Proposition 8.1.10. *Consider the \sim -bifunctor $(- \parallel -)$ along with the empty tcg $\mathbf{1}$, and*

$$\begin{array}{lll} \rho_A & : & A \parallel \mathbf{1} \rightarrow A \\ \lambda_A & : & \mathbf{1} \parallel A \rightarrow A \\ \alpha_{A,B,C} & : & (A \parallel B) \parallel C \rightarrow A \parallel (B \parallel C) \\ s_{A,B} & : & A \parallel B \rightarrow B \parallel A \end{array}$$

defined with $\rho_A = \mathfrak{c}_{\rho_A}, \lambda_A = \mathfrak{c}_{\lambda_A}, \alpha_{A,B,C} = \mathfrak{c}_{\alpha_{A,B,C}}$ and $s_{A,B} = \mathfrak{c}_{s_{A,B}}$.

This data makes **TCG** a symmetric monoidal \sim -category.

Proof. The bifunctor is constructed in Proposition 8.1.4. The coherence laws follow from the symmetric monoidal structure of **Ren** along with Proposition 8.1.9.

It remains to establish naturality, which is now relatively direct. We detail naturality of $s_{A,B}$ for illustration purposes: take $\sigma : A \vdash A'$ and $\tau : B \vdash B'$. Observe

$$s_{A',B'} \cdot (\sigma \parallel \tau) \approx (\tau \parallel \sigma) \cdot s_{A,B}^\perp \quad (8.1)$$

as follows directly by Corollary 7.2.13 applied to the obvious isomorphisms $\mathcal{C}^+(\sigma \parallel \tau) \cong \mathcal{C}^+(\tau \parallel \sigma)$ and $\mathcal{S}^+(\sigma \parallel \tau) \cong \mathcal{S}^+(\tau \parallel \sigma)$. From that we calculate:

$$\begin{aligned} s_{A',B'} \odot (\sigma \parallel \tau) &\approx s_{A',B'} \cdot (\sigma \parallel \tau) \\ &\approx (\tau \parallel \sigma) \cdot s_{A,B}^\perp \\ &\approx (\tau \parallel \sigma) \odot \mathfrak{p}_{s_{A,B}^\perp} \\ &\approx (\tau \parallel \sigma) \odot \mathfrak{c}_{s_{B,A}} \\ &\approx (\tau \parallel \sigma) \odot s_{B,A} \end{aligned}$$

using Lemma 8.1.8, equation (8.1), Lemma 8.1.8 again, Proposition 8.1.9, and definition of $s_{A,B}$. All other naturality squares follow using the same route. \square

Compact closed structure. Briefly recall that a compact closed category \mathcal{C} is a symmetric monoidal category where each object comes with a *dual* A^* and two morphisms

$$\epsilon_A : A^* \otimes A \rightarrow 1 \quad \eta_A : 1 \rightarrow A \otimes A^*,$$

the *unit* and the *co-unit*, satisfying two equations (see Theorem 8.1.12). A compact closed category is automatically symmetric monoidal closed: the monoidal closure can be defined as $A \multimap B = A^* \otimes B$. While this structure is not adequate to directly interpret a programming language, it is the fundamental structure expected of an unpolarized category of games, and it will be helpful for further constructions.

In **TCG**, the dual of a game A is as expected A^\perp . The unit and co-unit are:

Definition 8.1.11. Consider A a tcg. The **co-unit** $\epsilon_A : A^\perp \parallel A \vdash \mathbf{1}$ and **unit** $\eta_A : \mathbf{1} \vdash A \parallel A^\perp$, formed with ess respectively \mathfrak{c}_{A^\perp} and \mathfrak{c}_A with display maps respectively

$$\begin{array}{lll} \partial_{\epsilon_A} : |\mathfrak{c}_{A^\perp}| & \rightarrow & |A^\perp \parallel A \vdash \mathbf{1}| \\ (1, a) & \mapsto & (1, (1, a)) \\ (2, a) & \mapsto & (1, (2, a)), \end{array} \quad \begin{array}{lll} \partial_{\eta_A} : |\mathfrak{c}_A| & \rightarrow & |\mathbf{1} \vdash A \parallel A^\perp| \\ (1, a) & \mapsto & (2, (1, a)) \\ (2, a) & \mapsto & (2, (2, a)) \end{array}$$

are valid strategies.

Altogether, this provides the data for the following theorem:

Theorem 8.1.12. **TCG** is a compact closed \sim -category.

Proof. It remains to prove the two duality conditions:

$$\begin{aligned} \mathfrak{c}_A &\approx \rho_A \odot (\mathfrak{c}_A \parallel \epsilon_A) \odot \alpha_{A, A^\perp, A} \odot (\eta_A \parallel \mathfrak{c}_A) \odot \lambda_A^{-1} \\ \mathfrak{c}_{A^\perp} &\approx \lambda_{A^\perp} \odot (\epsilon_A \parallel \mathfrak{c}_{A^\perp}) \odot \alpha_{A^\perp, A, A^\perp}^{-1} \odot (\mathfrak{c}_{A^\perp} \parallel \eta_{A^\perp}) \odot \rho_{A^\perp}^{-1} \end{aligned}$$

which follow from Corollary 7.2.13 using Proposition 7.3.1, Lemmas 6.4.4 and 7.3.14 and the direct verification that all matching pairs involved are causally compatible. \square

As a compact closed \sim -category, **TCG** is a model of MLL and hence of the linear λ -calculus. But not more: like Conway games, **TCG** has neither products nor coproducts (as analysed by Melliès [Melliès, 2004b]). It also does not support the resource modality $!$ of Definition 7.1.26, which was defined on *negative* tcgs.

This missing structure may be recovered by switching to a *polarized* category of games: here *negative* games, where Opponent always starts (see Definition 7.1.24).

8.2 Negative Winning Games

As mentioned earlier, in order to obtain the categorical structure required to interpret programming language, one must commit to either *negative* games (which have products but no coproducts); or *positive* games (which have coproducts but no products). A pleasant feature of game semantics is that this choice reflects the evaluation order that one wishes to represent: call-by-name languages are naturally interpreted in categories of negative games, while call-by-value languages are interpreted in categories of positive games [Honda and Yoshida, 1999]. Here, we focus on negative games.

Winning games. We also introduce another condition: *winning* games and strategies. This mechanism, adapted from [Melliès, 2005, Melliès and Tabareau, 2007], has two purposes: (1) it compensates the inherently affine nature of game semantics, tightening the connections with relational-like semantics (see Section 10.4); and (2) it replaces the traditional but often bulky notion of *well-bracketing* in banning control operators like call/cc. We wire this into our ambient game semantics. We call this *winning* as it sets as common objective for both players to reach a winning position with as many Player as Opponent moves – and where intuitively speaking, all questions have an answer.

Relative Seely categories. Here, a reader familiar with categorical models of linear logic might expect us to construct a *Seely category*, a categorical model of intuitionistic linear logic [Melliès, 2009]. One can indeed construct a Seely category of thin concurrent games [Castellan and Clairambault, 2021]. But that construction is slightly unpleasant due to types such as $o \multimap (o \otimes o)$, i.e. tensors appearing on the right hand side of an arrow, which do not appear in our languages¹. Worse, this construction breaks the clean connection with the relational model. So we shall aim for a slightly weaker categorical structure called a *relative Seely category*, which slightly restricts the monoidal closure, banning in particular tensors on the right of an arrow. Like Seely categories, relative Seely categories still enjoy that the Kleisli category for $!$ is cartesian closed.

8.2.1 Relative Seely categories

Though that is a detour in our game-theoretic narrative, we must first define what we mean by relative Seely categories. Relative Seely categories model the fragment of intuitionistic linear logic with formulas those generated by the grammar

$$\begin{aligned} S, T &::= \top \mid S \& T \mid A \multimap S \\ A, B &::= 1 \mid A \otimes B \mid S \mid !S \end{aligned}$$

splitting formulas into *strict* S, T and *general* A, B formulas. Note that this covers all formulas involved in Girard’s call-by-name translation for types.

In the following definition, we make use of the standard notions of relative adjunctions and relative comonads. For the reader unfamiliar with these, the definition also contains explicit data. For full details, see Appendix A.

Definition 8.2.1. A *relative Seely category* is a symmetric monoidal category $(C, \otimes, 1)$ equipped with a full subcategory C_s together with the following data and axioms:

- C_s has finite products $(\&, \top)$ preserved by the inclusion functor $J : C_s \hookrightarrow C$.
- For every $B \in C$ there is a functor $B \multimap - : C_s \rightarrow C_s$, such that there is

$$\Lambda(-) : C(A \otimes B, S) \simeq C(A, B \multimap S).$$

a bijection natural in $A \in C$ and $S \in C_s$.

¹This triggers a duplication of the o on the left with conflict between the two copies. Monoidal closure works, but with a cumbersome technical overhead that we prefer to avoid here. Note that *products* on the right hand side of an arrow are fine; indeed we shall eventually form a cartesian closed Kleisli category.

- There is a **J**-relative comonad $! : C_s \rightarrow C$: we have, for every $S \in C_s$, an object $!S \in C$ and a **dereliction** morphism $\text{der}_S : !S \rightarrow S$, and for every $\sigma : !S \rightarrow T$, a **promotion** $\sigma^! : !S \rightarrow !T$, subject to three axioms [Altenkirch et al., 2010]:

$$\begin{aligned} \text{der}_T \circ \sigma^! &= \sigma & (\sigma : !S \rightarrow T) \\ \text{der}_S^! &= \text{id}_{!S} & (S \in C_s) \\ (\tau \circ \sigma^!)^! &= \tau^! \circ \sigma^! & (\sigma : !S \rightarrow T, \tau : T \rightarrow U), \end{aligned}$$

which make $! : C_s \rightarrow C$ a functor, via $!\sigma = (\sigma \circ \text{der}_S)^!$ for $\sigma : S \rightarrow T$.

- The functor $! : C_s \rightarrow C$ is symmetric strong monoidal $(C_s, \&, \top) \rightarrow (C, \otimes, 1)$, so

$$m_0 : 1 \rightarrow !\top \quad m_{S,T} : !S \otimes !T \rightarrow !(S \& T)$$

are natural isos, additional compatible with promotion: the diagram

$$\begin{array}{ccc} !\Gamma & \xrightarrow{\langle f, g \rangle^!} & !(S \& T) \\ \langle \text{der}, \text{der} \rangle^! \downarrow & & \downarrow m^{-1} \\ !(\Gamma \& \Gamma) & & !S \otimes !T \\ & \searrow m^{-1} \quad \nearrow f^! \otimes g^! & \\ & !\Gamma \otimes !\Gamma & \end{array}$$

commutes for all $\Gamma, S, T \in C_s$, $f \in C(!\Gamma, S)$, $g \in C(!\Gamma, T)$.

Note any Seely category is canonically a relative Seely category with $C = C_s$. For any relative Seely category, the Kleisli category associated with $!$, denoted $C_!$, is cartesian closed: it has objects those of C_s , and $C_!(S, T) = C(!S, T)$.

Lemma 8.2.2. *For a relative Seely category C , the Kleisli category $C_!$ is cartesian closed with finite products given as in C_s , and function space $S \Rightarrow T = !S \multimap T$.*

Proof. The proof is essentially as for Seely categories; see Appendix A.1. □

For **TCG** these definitions must of course be taken in \sim -categorical form: as before this means that all operations preserve \sim , and all conditions hold up to \sim ; yielding a notion of **relative Seely \sim -category**. We omit the straightforward adaptation.

We now start building the concrete relative Seely \sim -category **NTCG**.

8.2.2 Playing Board Games

Objects of **NTCG** are certain negative tcgs, equipped with an additional component κ , a so-called *payoff function* expressing the *winning conditions*. More precisely, κ records which configurations are complete and assigns responsibility of non-completeness to one of the players otherwise. The resulting object will be the main notion of game used in the rest of this monograph, so we need a distinctive name that remains “gamey”:

Boards. We now give the full definition of the objects of **NTCG**:

Definition 8.2.3. A **board** is a tcg A along with $\kappa_A : \mathcal{C}(A) \rightarrow \{-1, 0, +1\}$ a **payoff function**, such that this data satisfies the following conditions:

- invariant: for all $\theta : x \cong_A y$, we have $\kappa_A(x) = \kappa_A(y)$,
- race-free: for all $a \sim_A a'$, we have $\text{pol}_A(a) = \text{pol}_A(a')$.
- forestial: for all $a_1, a_2, a \in A$, if $a_1, a_2 \leq_A a$, then $a_1 \leq_A a_2$ or $a_2 \leq_A a_1$,
- alternating: for all $a_1, a_2 \in A$, if $a_1 \rightarrow_A a_2$, then $\text{pol}_A(a_1) \neq \text{pol}_A(a_2)$,

A **--board** is additionally negative, and must also satisfy:

$$\text{initialized: } \kappa_A(\emptyset) \geq 0.$$

Finally, a **--board** A is **strict** if $\kappa_A(\emptyset) = 1$ and all its initial moves are in pairwise conflict. It is **well-opened** if it is strict with exactly one initial move.

The payoff function κ_A assigns a value to each configuration. Configurations with payoff 0 are called complete: they correspond to *terminated* executions, which have reached an adequate stopping point where in particular all calls have adequately returned. Otherwise, κ_A assigns a responsibility for why a configuration is non-complete. If $\kappa_A(x) = -1$ then Player is responsible, otherwise it is Opponent.

Basic --boards. It seems natural to introduce first the boards for PCF base types.

Recall that in Figures 7.7, 7.8 and 7.9 we introduced tcgs \mathbf{U} , \mathbf{B} and \mathbf{N} respectively for the unit type, booleans and natural numbers. These three tcgs are turned into boards by adjoining them a payoff function which, for all three, is defined similarly by

$$\begin{aligned} \kappa(\emptyset) &= 1 \\ \kappa(\{\mathbf{q}^-\}) &= -1 \\ \kappa(\{\mathbf{q}^-, a^+\}) &= 0 \end{aligned}$$

which is exhaustive as maximal configurations have only two elements.

We keep the notations \mathbf{U} , \mathbf{B} and \mathbf{N} for these boards. They are *strict*: the payoff function essentially *forces* Opponent to ask the initial question. In these three basic boards, maximal configurations have null payoff – as mentioned earlier, we call those *complete*:

Definition 8.2.4. Consider A a board, and $x \in \mathcal{C}(A)$.

We say that x is **complete** iff $\kappa_A(x) = 0$, and write $x \in \mathcal{C}^0(A)$.

Complete configurations are intuitively those where every call is answered, in that they reflect the complete plays of traditional game semantics (see Definition 3.3.6). They shall also inform the link with relational semantics: this will be detailed in Section 10.4, but note already that $\mathcal{C}^0(\mathbf{U})$, $\mathcal{C}^0(\mathbf{B})$ and $\mathcal{C}^0(\mathbf{N})$ are in bijection with the usual relational model interpretation of the corresponding types.

The other basic **--boards** are the units. In the presence of the payoff function the empty tcg $\mathbf{1}$ splits into two units, reflecting the units of multiplicative and additive conjunctions in linear logic: the **top** \top has $\kappa_\top(\emptyset) = 1$, while the **one** $\mathbf{1}$ has $\kappa_{\mathbf{1}}(\emptyset) = 0$ – we overload the notation for the empty tcg $\mathbf{1}$, which should not cause any confusion.

\otimes	-1	0	+1
-1	-1	-1	-1
0	-1	0	+1
+1	-1	+1	+1

\wp	-1	0	+1
-1	-1	-1	+1
0	-1	0	+1
+1	+1	+1	+1

Figure 8.1: Payoff for \otimes and \wp

Dual, tensor and par. First, the **dual** (Definition 7.1.23) extends with payoff via $\kappa_{A^\perp}(x) = -\kappa_A(x)$ as expected. Of course, the dual does not preserve $--$ -boards.

Like the units, *parallel composition* splits into two:

Definition 8.2.5. Consider A and B two boards.

Their **tensor** $A \otimes B$ and their **par** $A \wp B$ are $A \parallel B$ enriched with:

$$\kappa_{A \otimes B}(x_A \parallel x_B) = \kappa_A(x_A) \otimes \kappa_B(x_B), \quad \kappa_{A \wp B}(x_A \parallel x_B) = \kappa_A(x_A) \wp \kappa_B(x_B)$$

with the operations \otimes and \wp defined on $\{-1, 0, +1\}$ in Figure 8.1.

The tensor of two $--$ -boards is still a $--$ -board, though tensor does not preserve *strict* $--$ -boards. The par also preserves $--$ -boards, but we shall not use it on $--$ -boards: if A and B are $--$ -boards, let us use $A \vdash B$ to denote the board $A^\perp \wp B$ used to define the strategies *from* A *to* B . Observe that even if A and B are $--$ -boards, $A \vdash B$ is not.

The with. We only consider the additive conjunction of linear logic: the with. But in order to define it, we must first define a new operation of ess and tcgs.

Definition 8.2.6. Let A_1 and A_2 be two tcgs.

Then, we define their **sum** $A_1 + A_2$ as comprising the components:

$$\begin{array}{llll}
\text{events:} & |A_1 \parallel A_2| & = & |A_1| + |A_2| \\
\text{causality:} & (i, a) \leq_{A_1 \parallel A_2} (j, a') & \Leftrightarrow & i = j \ \& \ a \leq_{A_i} a' \\
\text{conflict:} & (i, a) \#_{A_1 \parallel A_2} (j, a') & \Leftrightarrow & i \neq j \ \vee \ a \#_{E_i} a', \\
\text{symmetry:} & \theta \in \mathcal{S}(A_1 \parallel A_2) & \Leftrightarrow & \exists \theta_i \in \mathcal{S}(A_i), \theta = \theta_1 \parallel \theta_2, \\
\text{positive symmetries:} & \theta_1 \parallel \theta_2 \in \mathcal{S}_+(A_1 + A_2) & \Leftrightarrow & \theta_1 \in \mathcal{S}_+(A_1) \ \& \ \theta_2 \in \mathcal{S}_+(A_2) \\
\text{negative symmetries:} & \theta_1 \parallel \theta_2 \in \mathcal{S}_-(A_1 + A_2) & \Leftrightarrow & \theta_1 \in \mathcal{S}_-(A_1) \ \& \ \theta_2 \in \mathcal{S}_-(A_2).
\end{array}$$

where, necessarily, one of θ_1 or θ_2 must be empty.

Ignoring the positive and negative symmetries, this also yields an operation $+$ on plain event structures with symmetry that we shall use later on.

If A, B are tcgs and $x_A \in \mathcal{C}(A)$, we write $(1, x_A) \in \mathcal{C}(A + B)$ as a shorthand for $\{1\} \times x_A$ and likewise for $(2, x_B) = \{2\} \times x_B \in \mathcal{C}(A + B)$ for $x_B \in \mathcal{C}(B)$. Note that all configurations of $A + B$ have the form $(1, x_A)$ for $x_A \in \mathcal{C}(A)$ or $(2, x_B)$ for $x_B \in \mathcal{C}(B)$. For non-empty configurations, this decomposition is *unique*. We shall also use the corresponding notations for symmetries, with e.g. $(1, \theta_A) : (1, x_A) \cong_{A+B} (1, y_A)$ for $\theta_A : x_A \cong_A y_A$ comprising all $((1, a), (1, a'))$ for $(a, a') \in \theta_A$.

This sum operation yields the *with* operation on strict boards:

Definition 8.2.7. Consider A and B two strict $--$ -boards.

Then, their **with** $A \& B$ is the strict $--$ -board with *tcg* the sum $A + B$ and

$$\kappa_{A \& B}(1, x_A) = \kappa_A(x_A), \quad \kappa_{A \& B}(2, x_B) = \kappa_B(x_B),$$

for non-empty configurations and $\kappa_{A \& B}(\emptyset) = 1$.

As we will see, this construction will give a cartesian product in the subcategory of strict $--$ -boards. It can also be applied to non-strict $--$ -boards, but then it is not a product: if one of the A_i is not strict then the corresponding projection does not respect payoff (in the sense of Definition 8.2.14), because we have set $\kappa_{A_1 \& A_2}(\emptyset) = 1$. On the other hand having $\kappa_{A_1 \& A_2}(\emptyset) = 0$ breaks the correspondence with the relational model, since the empty configuration does not correspond in a canonical way to one of the components.

In the sequel, we should use the obvious n -ary generalization of the product, with respect to which any strict $--$ -board decomposes into a with of well-opened $--$ -boards:

Lemma 8.2.8. Consider A a strict $--$ -board. Then there is a family $(A_i)_{i \in I}$ of well-opened $--$ -boards, unique up to isomorphism, such that $A \cong \&_{i \in I} A_i$.

Proof. Straightforward. \square

We need notations for configurations of this board. Writing $\mathcal{C}^{\neq \emptyset}(A)$ (resp. $\mathcal{S}^{\neq \emptyset}(A)$) for the set of non-empty configurations (resp. symmetries) of A , we observe:

Lemma 8.2.9. Consider $(A_i)_{i \in I}$ a family of well-opened $--$ -boards. Then there are

$$\begin{aligned} \mathcal{C}^{\neq \emptyset}(\&_{i \in I} A_i) &\cong \sum_{i \in I} \mathcal{C}^{\neq \emptyset}(A_i) \\ \mathcal{S}^{\neq \emptyset}(\&_{i \in I} A_i) &\cong \sum_{i \in I} \mathcal{S}^{\neq \emptyset}(A_i) \\ \mathcal{S}_-^{\neq \emptyset}(\&_{i \in I} A_i) &\cong \sum_{i \in I} \mathcal{S}_-^{\neq \emptyset}(A_i) \\ \mathcal{S}_+^{\neq \emptyset}(\&_{i \in I} A_i) &\cong \sum_{i \in I} \mathcal{S}_+^{\neq \emptyset}(A_i) \end{aligned}$$

order-isos commuting with *dom* and *cod*.

Proof. Straightforward. \square


For $x \in \mathcal{C}^{\neq \emptyset}(A_i)$ (resp. $\theta \in \mathcal{S}^{\neq \emptyset}(A_i)$) we write $(i, x) \in \mathcal{C}^{\neq \emptyset}(\&_{i \in I} A_i)$ (resp. $(i, \theta) \in \mathcal{S}^{\neq \emptyset}(\&_{i \in I} A_i)$). The final linear connective we shall consider is the linear arrow:

Linear arrow. First, we define it in the case the rhs board is well-opened:

Definition 8.2.10. Consider A a $--$ -board and B a well-opened $--$ -board.

Then, $A \multimap B$ has *tcg* as defined in Definition 7.1.25, and payoff function:

$$\kappa_{A \multimap B}(x_A \parallel x_B) = \kappa_{A \vdash B}(x_A \parallel x_B) = \kappa_{A^\perp}(x_A) \wp \kappa_B(x_B).$$

Figure 8.2: The board $U \multimap (U \& U)$

The definition of relative Seely categories requires us to define $A \multimap B$ not only when B is well-opened (which has no particular status in the definition of relative Seely categories), but when it is strict. In that case, $A \multimap B$ may be defined directly via the decomposition of Lemma 8.2.8, as done in the following definition:

Definition 8.2.11. Consider A a $--$ -board, and B a strict $--$ -board, with $B \cong \&_{i \in I} B_i$. Then, we define $A \multimap B = \&_{i \in I} A \multimap B_i$.

The construction is illustrated in Figure 8.2.2, in which we omit payoff. The construction duplicates the argument arena, with one copy for each initial move of B . In that, the reader familiar with Hyland-Ong games – or the attentive reader of Part I) will recognize its arrow arena construction (modulo the conflict, see *e.g.* Figure 4.1).

It will be useful to have the analogue of Lemma 8.2.9 for the linear arrow:

Lemma 8.2.12. Consider A, B $--$ -boards with B strict. Then, there are:

$$\begin{aligned} \mathcal{C}^{\neq \emptyset}(A \multimap B) &\cong \mathcal{C}(A) \times \mathcal{C}^{\neq \emptyset}(B) \\ \mathcal{S}^{\neq \emptyset}(A \multimap B) &\cong \mathcal{S}(A) \times \mathcal{S}^{\neq \emptyset}(B) \\ \mathcal{S}_+^{\neq \emptyset}(A \multimap B) &\cong \mathcal{S}_+(A) \times \mathcal{S}_+^{\neq \emptyset}(B) \\ \mathcal{S}_-^{\neq \emptyset}(A \multimap B) &\cong \mathcal{S}_-(A) \times \mathcal{S}_-^{\neq \emptyset}(B) \end{aligned}$$

order-isos commuting with dom and cod.

Proof. If B is well-opened, then $\mathcal{C}^{\neq \emptyset}(A \multimap B) \cong \mathcal{C}(A) \times \mathcal{C}^{\neq \emptyset}(B)$ is obvious. Now if $B \cong \sum_{i \in I} B_i$ via Lemma 8.2.8, we compose bijections:

$$\begin{aligned} \mathcal{C}^{\neq \emptyset}(A \multimap B) &\cong \sum_{i \in I} \mathcal{C}^{\neq \emptyset}(A \multimap B_i) \\ &\cong \sum_{i \in I} \mathcal{C}(A) \times \mathcal{C}^{\neq \emptyset}(B_i) \\ &\cong \mathcal{C}(A) \times \left(\sum_{i \in I} \mathcal{C}^{\neq \emptyset}(B_i) \right) \\ &\cong \mathcal{C}(A) \times \mathcal{C}^{\neq \emptyset}(B) \end{aligned}$$

using Lemma 8.2.9, the well-opened case, distributivity, followed by Lemma 8.2.9 again. The same construction applies to symmetries in a compatible way. \square

Following this lemma, we adopt the convention that for each $x_A \in \mathcal{C}(A)$ and $x_B \in \mathcal{C}^{\neq \emptyset}(B)$, $x_A \multimap x_B \in \mathcal{C}^{\neq \emptyset}(A \multimap B)$ denotes the corresponding configuration.

Bang. The next definition enriches Definition 7.1.26 with payoff:

Definition 8.2.13. Consider A a \multimap -board.

Then, $!A$ has tcg as in Definition 7.1.26 enriched with:

$$\begin{aligned} \kappa_{!A}(\parallel_{i \in I} x_i) &= \bigotimes_{i \in I} \kappa_A(x_i) & (I \subseteq \mathbb{N}, \forall i \in I, x_i \neq \emptyset) \\ \kappa_{!A}(\emptyset) &= 0 \end{aligned}$$

where $\parallel_{i \in I} x_i = \biguplus_{i \in I} (\{i\} \times x_i)$.

So the payoff for the bang acts as a n -ary tensor for the non-empty components, while the empty configuration gets assigned payoff zero. This means that it is considered *complete*: a complete execution might not initiate any of the copies; but a copy initiated must be brought to a complete state if the whole computation is to be complete.

A pleasing consequence of this definition is that if A is strict, then the symmetry classes of complete configurations of $!A$ are in one-to-one correspondence with finite multisets of symmetry classes of complete configurations of A , informing the link with the relational model – again, see Section 10.4 for more details.

8.2.3 Negative Winning Strategies

Winning strategies. First, we define what it means for a strategy to be well-behaved with respect to payoff. Recall that this has two effects: it compensates for the inherently affine nature of strategies, and it also replaces the traditional well-bracketing condition.

Intuitively, a strategy is *winning* when its stopping points are complete, *or* the incompleteness can be attributed – via the payoff function – to Opponent:

Definition 8.2.14. Consider A a board, and $\sigma : A$ a strategy. We define conditions:

$$\begin{aligned} \text{negative:} & \text{ for all } s \in \sigma \text{ minimal for } \leq_\sigma, \text{ we have } \text{pol}_A(\partial_\sigma s) = -, \\ \text{winning:} & \text{ for all } x^\sigma \in \mathcal{C}^+(\sigma), \kappa_A(\partial_\sigma x^\sigma) \geq 0. \end{aligned}$$

Negative means that Opponent always starts, also in strategies. *Winning* forces σ to be strictly linear rather than affine, *i.e.* it *must* investigate arguments not marked with a $!$. For instance, the strategy displayed in Figure 8.3 is non-winning: as U is strict, the fact that the strategy does not investigate its argument is punished by a payoff of -1 on the left, yielding a global payoff of $-1 \wp 0 = -1$ for the maximal configuration.

This definition also ensures a measure of well-bracketing, though it may not be obvious to the reader. Conceptually, Melliès and Tabareau’s insight is that well-bracketing is a linearity constraint [Melliès and Tabareau, 2007]: in a well-bracketed completed (in the sense that the initial question receives an answer) execution, each question receives exactly one answer. Still, it might be helpful to examine a concrete example of how typically non-well-bracketed behaviour in the traditional sense is banned. Such an

$$U \multimap U$$



$$(U \multimap U) \multimap U$$




Figure 8.3: An affine non-winning strategy

Figure 8.4: A non-winning strategy

example appears in Figure 8.4: the configuration without the gray event is $+$ -covered. Its payoff is $(-1 \otimes +1) \bowtie 0 = -1$, where the -1 is the price of Player omitting to answer the “pending” left-most question. But this is actually quite subtle and one should not attempt to match this notion of well-bracketing with the traditional ones. In particular, the same causal pattern may very well appear in a winning strategy: for instance, we invite the reader to verify that the strategy in Figure 7.6 (on the left) is winning.

We obtain a \sim -category of $--$ -boards and negative winning strategies:

Proposition 8.2.15. *There is **NTCG**, a \sim -category with $--$ -boards as objects, negative winning strategies on $A \vdash B$ as morphisms from A to B , and \approx as equivalence.*

Proof. *Copycat.* *Negative* is immediate by inspection of Definition 6.4.1. For *winning*, consider $x_A \parallel x_A \in \mathcal{C}^+(\mathfrak{A}_A)$ exploiting Lemma 6.4.4. Then, $\kappa_{A \vdash A}(x_A \parallel x_A) = -p \bowtie p$ for $p = \kappa_A(x_A)$, and for all $p \in \{-1, 0, +1\}$, $-p \bowtie p \geq 0$ as required.

Composition. For *negative*, assume *w.l.o.g.* that there is $p \in \tau \odot \sigma$ minimal positive. In particular, $x^\tau \odot x^\sigma = \{p\} \in \mathcal{C}^+(\tau \odot \sigma)$. Now p cannot occur in C , as minimal events in C are negative, so it must occur in A . That means $x^\sigma \in \mathcal{C}^+(\sigma)$ is non-empty. As σ is negative, there is a minimal (necessarily) negative event in x^σ , which must occur in B . As x^σ and x^τ are matching, x^τ is non-empty as well. By *negative*, there is a minimal (necessarily) negative event in x^τ , which must occur in C . But as $\partial_{\tau \odot \sigma}(x^\tau \odot x^\sigma) = x_A^\sigma \parallel x_C^\tau$, $x^\tau \odot x^\sigma$ must have an event occurring in C , absurd.

For *winning*, consider $x^\tau \odot x^\sigma \in \mathcal{C}^+(\tau \odot \sigma)$, with $\partial_{\tau \odot \sigma}(x^\tau \odot x^\sigma) = x_A^\sigma \parallel x_C^\tau$. If $\kappa_A(x_A^\sigma) = -1$, then $\kappa_{A \vdash C}(x_A^\sigma \parallel x_C^\tau) = 1$ and we are done. Likewise if $\kappa_C(x_C^\tau) = 1$, we are done, so assume neither. Now if $\kappa_A(x_A^\sigma) = 1$, as σ is winning we must have $\kappa_B(x_B^\sigma) = \kappa_B(x_B^\tau) = 1$ as well; and as τ is winning this entails $\kappa_C(x_C^\tau) = 1$, absurd. Symmetrically, $\kappa_C(x_C^\tau) = -1$ entails $\kappa_A(x_A^\sigma) = -1$, absurd. The only case left has $\kappa_A(x_A^\sigma) = \kappa_C(x_C^\tau) = 0$, so that $\kappa_{A \vdash C}(x_A^\sigma \parallel x_C^\tau) = 0 \geq 0$ as required. \square

The properties of the equivalence relation \approx packaged in Proposition 7.4.15 are undisturbed by these now conditions and apply transparently.

8.2.4 Symmetric Monoidal Structure

We now extend the symmetric monoidal structure.

Tensor. First, we extend the tensor of strategies:

Proposition 8.2.16. *The tensor operation of Definition 8.2.5 yields a \sim -bifunctor:*

$$(- \otimes -) : \mathbf{NTCG} \times \mathbf{NTCG} \rightarrow \mathbf{NTCG}.$$

Proof. Consider $\sigma : A \vdash B, \tau : C \vdash D$, their tensor $\sigma \otimes \tau$ is defined on strategies as in Proposition 8.1.1. We must check that this preserves *negative* and *winning*.

Negative. Consider $m \in \sigma \parallel \tau$ minimal. Seeking a contradiction, assume m is positive, so $\{m\}$ is $+$ -covered. By Proposition 8.1.1, $\{m\}$ may be written as $x^\sigma \parallel x^\tau \in \mathcal{C}^+(\sigma \otimes \tau)$. As $\partial_{\sigma \otimes \tau}(m) \in A \otimes C \vdash B \otimes D$, m occurs in one of the four components. Say *w.l.o.g.* that it is in A or B , then x^σ must be a singleton. But as σ is negative, $\partial_\sigma x^\sigma$ is a singleton whose only event is negative in $A \vdash B$. By the constraints on display maps given by Proposition 8.1.1, it follows that m is negative, contradiction.

Winning. Consider $x^\sigma \parallel x^\tau \in \mathcal{C}^+(\sigma \otimes \tau)$. If $\kappa_A(x_A^\sigma) = -1$ or $\kappa_C(x_C^\tau) = -1$, then we are done. As σ and τ are winning, this directly entails that we cannot have $\kappa_B(x_B^\sigma) = -1$ or $\kappa_D(x_D^\tau) = -1$ either, so $\kappa_B(x_B^\sigma), \kappa_D(x_D^\tau) \geq 0$. Now, assume $\kappa_{A \otimes C}(x_A^\sigma \parallel x_C^\tau) = 1$. If $\kappa_A(x_A^\sigma) = 1$, then as σ is winning we have $\kappa_B(x_B^\sigma) = 1$. But as $\kappa_D(x_D^\tau) \geq 0$, $\kappa_{B \otimes D}(x_B^\sigma \parallel x_D^\tau) = 1$ and we are done. Likewise if $\kappa_C(x_C^\tau) = 1$, we may conclude. \square

Lifting. The monoidal structure shall be obtained following Section 8.1.2, by lifting maps. However, we must first define those renamings that yield winning strategies.

Definition 8.2.17. *Consider A, B two boards, and $f : A \rightarrow B$ a renaming.*

*We say that f is **winning** if for all $x_A \in \mathcal{C}(A)$, $\kappa_B(f(x_A)) \geq \kappa_A(x_A)$.*

As expected, winning renamings preserve winning strategies:

Lemma 8.2.18. *Consider winning renamings $f : A^\perp \rightarrow A'^\perp$ and $g : B \rightarrow B'$.*

Then, $g \cdot \sigma \cdot f : A' \vdash B'$ is winning.

Proof. Consider $x^\sigma \in \mathcal{C}^+(\sigma)$. If $\kappa_{A'}(f(x_A^\sigma)) = -1$ or $\kappa_{B'}(g(x_B^\sigma)) = 1$ then we are done, so assume otherwise. If $\kappa_{A'}(f(x_A^\sigma)) = 1$, then $\kappa_{A'^\perp}(f(x_A^\sigma)) = -1$, but since f is winning, we have $\kappa_{A'^\perp}(f(x_A^\sigma)) \geq \kappa_{A^\perp}(x_A^\sigma)$ so $\kappa_{A^\perp}(x_A^\sigma) = -1$ as well. Now since σ is winning, this entails $\kappa_B(x_B^\sigma) = 1$. But then $\kappa_{B'}(g(x_B^\sigma)) \geq \kappa_B(x_B^\sigma)$ as g is winning, so $\kappa_{B'}(g(x_B^\sigma)) = 1$, contradiction. Symmetrically, if $\kappa_{B'}(g(x_B^\sigma)) = -1$, this entails $\kappa_{A'}(f(x_A^\sigma)) = -1$ as f, σ and g are winning, contradiction. \square

It follows directly that we may lift winning renamings to strategies:

Corollary 8.2.19. *Consider $f : A \rightarrow B$ a winning renaming. Then,*

- (1) $\mathfrak{C}_f : A \vdash B$ is a negative, winning strategy,
- (2) $\mathfrak{D}_f : B^\perp \vdash A^\perp$ is a negative, winning strategy.

Proof. Immediate by Lemma 8.2.18. \square

And finally, we deduce the symmetric monoidal structure:

Proposition 8.2.20. *Proposition 8.2.16 along with the data of Proposition 8.1.10 makes \mathbf{NTCG} a symmetric monoidal \sim -category.*

Proof. Direct from Propositions 8.2.16, 8.1.10 and Corollary 8.2.19. \square

8.2.5 Cartesian Product of Strict $--$ -Boards

Next, we must define the cartesian structure. Recall from the definition of relative Seely categories (Definition 8.2.1) that the cartesian product is not required to be defined on the whole category, but only on a full subcategory C_s . Here, we will consider the full \sim -subcategory \mathbf{NTCG}_s with as objects the *strict* $--$ -boards.

Projections. First, we define the projections. First we need renamings:

Lemma 8.2.21. *Consider A, B strict $--$ -boards.*

The injections $\iota_A : A^\perp \rightarrow (A \& B)^\perp$, $\iota_B : B^\perp \rightarrow (A \& B)^\perp$ are winning renamings.

Proof. Immediate verification. \square

The above lets us define the projections via lifting:

$$\pi_A = \mathfrak{P}_{\iota_A} : A \& B \vdash A, \quad \pi_B = \mathfrak{P}_{\iota_B} : A \& B \vdash B$$

winning negative strategies by Corollary 8.2.19.

Pairing. As for earlier constructions, we build pairing via a universal construction:

Proposition 8.2.22. *For $--$ -boards Γ, A, B , with A, B strict, and negative winning strategies $\sigma : \Gamma \vdash A$, $\tau : \Gamma \vdash B$, there is a negative winning strategy $\langle \sigma, \tau \rangle : \Gamma \vdash A \& B$, unique up to iso, s.t. there are order-isos:*

$$\begin{aligned} \mathcal{C}^+(\sigma) + \mathcal{C}^+(\tau) &\simeq \mathcal{C}^+(\langle \sigma, \tau \rangle) \\ \mathcal{S}^+(\sigma) + \mathcal{S}^+(\tau) &\simeq \mathcal{S}^+(\langle \sigma, \tau \rangle) \end{aligned}$$

commuting with dom , cod , and such that for all $\theta^\sigma \in \mathcal{S}^+(\sigma)$ and $\theta^\tau \in \mathcal{S}^+(\tau)$, we have

$$\partial_{\langle \sigma, \tau \rangle}(\iota_\sigma(\theta^\sigma)) = \theta^\sigma_\Gamma \parallel (\theta^\sigma_A \parallel \emptyset), \quad \partial_{\langle \sigma, \tau \rangle}(\iota_\tau(\theta^\tau)) = \theta^\tau_\Gamma \parallel (\emptyset \parallel \theta^\tau_B)$$

with $\iota_\sigma : \mathcal{C}^+(\sigma) \rightarrow \mathcal{C}^+(\langle \sigma, \tau \rangle)$ and $\iota_\tau : \mathcal{C}^+(\tau) \rightarrow \mathcal{C}^+(\langle \sigma, \tau \rangle)$ the induced injections.

Moreover, $\langle -, - \rangle$ preserves \approx .

Proof. Existence. We set $\langle \sigma, \tau \rangle$ as having as ess $\sigma + \tau$, with display maps given by

$$\begin{aligned} \partial_{\langle \sigma, \tau \rangle}(\iota_\sigma(x^\sigma)) &= x^\sigma_\Gamma \vdash (1, x^\sigma_A) \\ \partial_{\langle \sigma, \tau \rangle}(\iota_\tau(x^\tau)) &= x^\tau_\Gamma \vdash (2, x^\tau_B) \end{aligned}$$

where $\partial_\sigma(x^\sigma) = x^\sigma_\Gamma \vdash x^\sigma_A$ and $\partial_\tau(x^\tau) = x^\tau_\Gamma \vdash x^\tau_B$; and likewise for symmetries. All required verifications are immediate, and this yields a negative winning strategy.

Uniqueness. Direct consequence of Corollary 7.2.13.

Congruence. Direct from the definition. \square

We need one last lemma to conclude with the cartesian product:

Lemma 8.2.23. *Consider Γ, A, B --boards with A, B strict, and $\alpha : \Gamma \vdash A \& B$. Then,*

$$\begin{aligned}\mathcal{C}^+(\pi_A \odot \alpha) &\cong \{x^\alpha \in \mathcal{C}^+(\alpha) \mid \partial_\alpha x^\alpha = x_\Gamma^\alpha \vdash (1, x_A^\alpha)\} \\ \mathcal{S}^+(\pi_A \odot \alpha) &\cong \{\theta^\alpha \in \mathcal{S}^+(\alpha) \mid \partial_\alpha \theta^\alpha = \theta_\Gamma^\alpha \vdash (1, \theta_A^\alpha)\}\end{aligned}$$

are order-isomorphisms compatible with display maps – and likewise for π_B .

Proof. A direct adaptation of the proof of Proposition 6.4.9. \square

Proposition 8.2.24. *For any two strict --boards A and B , $A \& B$ is a cartesian product of A and B in **NTCG**. Moreover, \top is terminal.*

Proof. We prove the universal property of cartesian products.

Existence. For $\sigma : \Gamma \vdash A$, $\tau : \Gamma \vdash B$ we form $\langle \sigma, \tau \rangle$ via Proposition 8.2.22. Then:

$$\begin{aligned}\mathcal{C}^+(\pi_A \odot \langle \sigma, \tau \rangle) &\cong \{x \in \mathcal{C}^+(\langle \sigma, \tau \rangle) \mid \partial x = x_\Gamma \vdash (1, x_A)\} \\ &\cong \mathcal{C}^+(\sigma)\end{aligned}$$

using Lemma 8.2.23 followed by Proposition 8.2.22, and those order-isos preserve display maps. The same reasoning applies to symmetries, so that $\pi_A \odot \langle \sigma, \tau \rangle \approx \sigma$ by Corollary 7.2.13. Symmetrically, we have $\pi_B \odot \langle \sigma, \tau \rangle \approx \tau$ as well.

Uniqueness. Consider $\alpha : \Gamma \vdash A \& B$ such that $\pi_A \odot \alpha \approx \sigma$ and $\pi_B \odot \alpha \approx \tau$. First,

$$\alpha \cong \langle \pi_A \odot \alpha, \pi_B \odot \alpha \rangle$$

follows from Lemma 8.2.23, uniqueness of Proposition 8.2.22 and immediate verifications. Finally, we have $\langle \pi_A \odot \alpha, \pi_B \odot \alpha \rangle \approx \langle \sigma, \tau \rangle$ as $\langle -, - \rangle$ preserves \approx .

Terminal. For any --board Γ , it is immediate that the empty strategy $\epsilon : \Gamma \vdash \top$ is a winning negative strategy. Moreover, given any $\alpha : \Gamma \vdash \top$ negative, any minimal event in α must be negative and hence map to an event in \top – but there are none. \square

This concludes the investigation of cartesian products in **NTCG**.

8.2.6 Relative Closure

The final piece of the linear structure is the relative closure (see Definition 8.2.1). The main operation we must define is **currying** – as for the other operations, we characterize it uniquely up to isomorphism via its configurations.

Proposition 8.2.25. *Consider Γ, A, B --boards with B strict.*

For $\sigma : \Gamma \otimes A \vdash B$, there is $\Lambda(\sigma) : \Gamma \vdash A \multimap B$, unique up to iso, s.t. there are

$$\begin{aligned}\Lambda(-) &: \mathcal{C}^+(\sigma) \cong \mathcal{C}^+(\Lambda(\sigma)) \\ \Lambda(-) &: \mathcal{S}^+(\sigma) \cong \mathcal{S}^+(\Lambda(\sigma))\end{aligned}$$

order-isos commuting with dom, cod, and such that for all θ^σ non-empty,

$$\partial_{\Lambda(\sigma)}(\Lambda(\theta^\sigma)) = \theta_\Gamma^\sigma \vdash \theta_A^\sigma \multimap \theta_B^\sigma \quad (8.2)$$

whenever $\partial_\sigma(\theta^\sigma) = \theta_\Gamma^\sigma \parallel \theta_A^\sigma \vdash \theta_B^\sigma$.

Moreover, $\Lambda(-)$ preserves \approx .

Proof. Note that (8.2) relies on Lemmas 7.1.22, 8.2.9 and 8.2.12, via which there are

$$\mathcal{C}(\Gamma \otimes A) \times \mathcal{C}^{\neq \emptyset}(B) \cong \mathcal{C}(\Gamma) \times \mathcal{C}^{\neq \emptyset}(A \multimap B) \quad (8.3)$$

$$\mathcal{S}(\Gamma \otimes A) \times \mathcal{S}^{\neq \emptyset}(B) \cong \mathcal{S}(\Gamma) \times \mathcal{S}^{\neq \emptyset}(A \multimap B) \quad (8.4)$$

order-isos preserving dom and cod – by negativity of σ , for $\theta^\sigma \in \mathcal{S}^{\neq \emptyset}(\sigma)$, $\theta_B^\sigma \neq \emptyset$.

Existence. We set $\Lambda(\sigma)$ with the same ess as σ , so in particular $\mathcal{C}^+(\sigma) = \mathcal{C}^+(\Lambda(\sigma))$ and $\mathcal{S}^+(\sigma) = \mathcal{S}^+(\Lambda(\sigma))$. It remains to redefine the display map of $\Lambda(\sigma)$; but given the equalities above and Lemma 7.1.9, (8.2) may be read as a definition.

Uniqueness. Follows from Corollary 7.2.13.

Congruence. Consider $f : \sigma \approx \tau$ a positive isomorphism. As $\Lambda(-)$ leaves the ess unchanged, it suffices to show that $f : \Lambda(\sigma) \cong \Lambda(\tau)$ is still a positive isomorphism. Now, by hypothesis, we know that for all $x^\sigma \in \mathcal{C}(\sigma)$, the following square commutes:

$$\begin{array}{ccc} x^\sigma & \xrightarrow{f} & f(x^\sigma) \\ \partial_\sigma \downarrow & & \downarrow \partial_\tau \\ x_\Gamma^\sigma \parallel x_A^\sigma \vdash x_B^\sigma & \xrightarrow{\theta_\Gamma \parallel \theta_A \vdash \theta_B} & y_\Gamma^\tau \parallel y_A^\tau \vdash y_B^\tau \end{array}$$

where $\theta_\Gamma \parallel \theta_A \vdash \theta_B : x_\Gamma^\sigma \parallel x_A^\sigma \vdash x_B^\sigma \cong_{\Gamma \otimes A \vdash B}^+ y_\Gamma^\tau \parallel y_A^\tau \vdash y_B^\tau$ is a positive symmetry.

Now, from Lemma 8.2.12 and definition of $\partial_{\Lambda(\sigma)}$, $\partial_{\Lambda(\tau)}$, the following typechecks:

$$\begin{array}{ccc} x^\sigma & \xrightarrow{f} & f(x^\sigma) \\ \partial_\sigma \downarrow & & \downarrow \partial_\tau \\ x_\Gamma^\sigma \vdash x_A^\sigma \multimap x_B^\sigma & \xrightarrow{\theta_\Gamma \vdash \theta_A \multimap \theta_B} & y_\Gamma^\tau \vdash y_A^\tau \multimap y_B^\tau \end{array}$$

and it commutes by monotonicity of all constructions involved, since maps in this square are in particular maps of event structures (regarding configurations as conflict-free event structures) and maps coinciding on configurations must be equal (Lemma 6.1.12).

Finally, from Lemmas 7.1.22, 8.2.9 and 8.2.12, $\theta_\Gamma \vdash \theta_A \multimap \theta_B$ is positive. \square

Next in the construction of the closure comes that $\Lambda(-)$ has an inverse:

Lemma 8.2.26. *For any Γ, A, B --boards with B strict, we have a bijection:*

$$\Lambda(-) : \mathbf{NTCG}(\Gamma \otimes A, B) \simeq \mathbf{NTCG}(\Gamma, A \multimap B)$$

preserving and reflecting \approx .

Proof. Given $\sigma : \Gamma \vdash A \multimap B$, $\Lambda^{-1}(\sigma) : \Gamma \otimes A \vdash B$ has same ess as σ , with

$$\partial_{\Lambda^{-1}(\sigma)}(x^\sigma) = x_\Gamma^\sigma \parallel x_A^\sigma \vdash x_B^\sigma$$

for $x^\sigma \in \mathcal{C}^{\neq \emptyset}(\sigma)$, using (8.3) and likewise for symmetries via (8.4). As in the proof of Proposition 8.2.25, via Lemma 7.1.9 this may be indeed read as a definition.

That $\Lambda^{-1}(-)$ preserves \approx is as in the proof of Proposition 8.2.25. Finally, from the definition we have $\Lambda(\Lambda^{-1}(\sigma)) = \sigma$ and $\Lambda^{-1}(\Lambda(\tau)) = \tau$, *actual equalities*, as definitions of $\Lambda(-)$ and Λ^{-1} follow opposite directions of the isomorphism (8.3). \square

As usual, **evaluation** $\mathbf{ev}_{A,B} : (A \multimap B) \otimes A \vdash B$ may be defined as

$$\mathbf{ev}_{A,B} = \Lambda^{-1}(\mathfrak{c}_{A \multimap B})$$

the uncurrying of the identity, for A, B $--$ -boards with B strict. We have:

Lemma 8.2.27. *Consider A, B $--$ -boards with B strict. Then,*

$$\begin{aligned} \mathcal{C}^{\neq \emptyset, +}(\mathbf{ev}_{A,B}) &= \{ \mathfrak{c}_{x_A \multimap x_B} \mid x_A \in \mathcal{C}(A), x_B \in \mathcal{C}^{\neq \emptyset}(A) \}, \\ \mathcal{S}^{\neq \emptyset, +}(\mathbf{ev}_{A,B}) &= \{ \mathfrak{c}_{\theta_A \multimap \theta_B} \mid \theta_A \in \mathcal{S}(A), \theta_B \in \mathcal{S}^{\neq \emptyset}(B) \} \end{aligned}$$

with $\partial_{\mathbf{ev}_{A,B}}(\mathfrak{c}_{x_A \multimap x_B}) = (x_A \multimap x_B) \parallel x_A \vdash x_B$ and idem for symmetries.

Proof. Direct from Lemmas 6.4.4, 7.3.14 and the definition of $\Lambda^{-1}(-)$. \square

Finally, uncurrying may be equivalently computed by composing with evaluation:

Lemma 8.2.28. *Consider $\sigma : \Gamma \vdash A \multimap B$ a negative winning strategy.*

Then, $\Lambda^{-1}(\sigma) \approx \mathbf{ev}_{A,B} \odot (\sigma \otimes \mathfrak{c}_A)$.

Proof. We calculate:

$$\begin{aligned} &\mathcal{C}^+(\mathbf{ev}_{A,B} \odot (\sigma \otimes \mathfrak{c}_A)) \\ &\cong \{ (\mathfrak{c}_{x_A \multimap x_B}, (x^\sigma \parallel \mathfrak{c}_{y_A})) \mid \text{matching, c.c.} \} \\ &= \{ (\mathfrak{c}_{x_A \multimap x_B}, (x^\sigma \parallel \mathfrak{c}_{y_A})) \mid \text{matching} \} \\ &= \{ (\mathfrak{c}_{x_A^\sigma \multimap x_B^\sigma}, (x^\sigma \parallel \mathfrak{c}_{x_A^\sigma})) \mid x^\sigma \in \mathcal{C}^+(\sigma) \} \\ &\cong \mathcal{C}^+(\sigma) \end{aligned}$$

using Propositions 7.3.1 and 8.1.1, Lemmas 6.4.4 and 8.2.27; a direct verification to ensure that no causal loop arises; and resolving the constraints due to *matching* by definition of the display maps. All order-isomorphisms involved are compatible with display maps, and the same reasoning compatibly applies to symmetries, so altogether we get a strong isomorphism by Corollary 7.2.13. \square

Finally, we may conclude the universal property of relative closure for **NTCG**:

Proposition 8.2.29. *Consider Γ, A, B $--$ -arenas with B strict.*

For any $\sigma : \Gamma \otimes A \vdash B$ negative and winning, there is a unique (up to \approx) $\Lambda(\sigma) : \Gamma \vdash A \multimap B$ such that $\mathbf{ev}_{A,B} \odot (\Lambda(\sigma) \otimes \mathfrak{c}_A) \approx \sigma$.

Proof. *Existence.* First, $\Lambda(\sigma)$ is obtained by Proposition 8.2.25. Now we have

$$\mathbf{ev}_{A,B} \odot (\Lambda(\sigma) \otimes \mathfrak{c}_A) \approx \Lambda^{-1}(\Lambda(\sigma)) \approx \sigma$$

by Lemma 8.2.28 followed by Lemma 8.2.26.

Uniqueness. If $\tau : \Gamma \vdash A \multimap B$ satisfies the desired property, then $\Lambda^{-1}(\tau) \approx \sigma$ by Lemma 8.2.28, so $\tau \approx \Lambda(\sigma)$ since $\Lambda(-)$ preserves \approx and by Lemma 8.2.26. \square

8.3 Non-Linear Structure

Next, we describe the non-linear structure in boards.

8.3.1 The Resource Modality

Last but not least, we must define $!$, the resource modality. This is where all the somewhat heavy handling of symmetries up to now finally strikes in.

Functorial action. On a strict $--$ -board A , the (non-strict) $--$ -board $!A$ was defined in Definition 8.2.13. We must extend it to a functor, thus first define the functorial action.

As observed below Definition 7.1.26, the construction $!(-)$ applies to plain ess, so we may form the ess $!\sigma$ which we must equip with a display map. For that, we use:

Lemma 8.3.1. *For E an ess, there is an order-isomorphism*

$$[-] : \text{Fam}(\mathcal{C}^{\neq \emptyset}(E)) \simeq \mathcal{C}(!E)$$

with $\text{Fam}(X)$ the set of families of elements of X indexed by finite subsets of \mathbb{N} .

Proof. If $(x_i)_{i \in I}$ is a family of non-empty configurations of E , we write

$$[(x_i)_{i \in I}] = \parallel_{i \in I} x_i = \biguplus_{i \in I} \{i\} \times x_i \in \mathcal{C}(!E),$$

and this decomposition is clearly unique if we insist that all x_i s are non-empty. \square

This order-iso for configurations of $!E$ is very much like similar isomorphisms we encountered for other constructions, see *e.g.* Lemma 7.1.22 for parallel composition / tensor, Lemma 8.2.9 with $\&$, and Lemma 8.2.12 for \multimap . However, unlike those, it does not extend to symmetries. This is because symmetries of $!E$ do not restrict component-wise, on the contrary, their role is precisely to span accross components. However:

Lemma 8.3.2. *Consider A a strict $--$ -board. Then, there is an order-isomorphism*

$$[-] : \text{Fam}(\mathcal{S}_+^{\neq \emptyset}(A)) \cong \mathcal{S}_+(!A)$$

compatible with dom and cod.

Proof. Direct from the definition. \square

In order to define the strategy $!\sigma$, it suffices to assign a display map:

Proposition 8.3.3. *Consider A, B strict $--$ -boards, and $\sigma : A \vdash B$ negative winning.*

Then, the ess $!\sigma$ may be equipped with a display map $\partial_{!\sigma}$ such that

$$\partial_{!\sigma}([x^i \mid i \in I]) = [x_A^i \mid i \in I] \vdash [x_B^i \mid i \in I]$$

where $\partial_{\sigma}(x^i) = x_A^i \vdash x_B^i$ for $i \in I$; making $!\sigma : !A \vdash !B$ negative winning.

Moreover, $!(-)$ preserves \approx .

Proof. Strategy. The display map defined above on configurations preserves unions and cardinal; moreover it is straightforward to extend it to symmetries. This gives a map of ess by Lemma 7.1.9, we omit the verifications that this yields a strategy.

Negative, winning. A n -ary analogue of the tensor (see Proposition 8.2.16).

Congruence. Consider $f : \sigma \approx \tau$. We construct $!f : !\sigma \approx !\tau$ simply component-wise. Clearly, it is an isomorphism, we must show that it commutes with the display to the board up to a positive symmetry. For that, consider $[x^i \mid i \in I] \in \mathcal{C}(!\sigma)$. For $i \in I$

$$\begin{array}{ccc} x^i & \xrightarrow{f} & f(x^i) \\ \partial_\sigma \downarrow & & \downarrow \partial_\tau \\ x_A^i \vdash x_B^i & \xrightarrow{\theta_A^i \vdash \theta_B^i} & y_A^i \vdash y_B^i \end{array}$$

where $\theta_A^i \in \mathcal{S}_-(A)$ and $\theta_B^i \in \mathcal{S}_+(B)$. It follows by Lemma 8.3.2 that the diagram

$$\begin{array}{ccc} [x^i \mid i \in I] & \xrightarrow{!f} & [f(x^i) \mid i \in I] \\ \partial_\sigma \downarrow & & \downarrow \partial_\tau \\ [x_A^i \mid i \in I] \vdash [x_B^i \mid i \in I] & \xrightarrow{[\theta_A^i \mid i \in I] \vdash [\theta_B^i \mid i \in I]} & [y_A^i \mid i \in I] \vdash [y_B^i \mid i \in I] \end{array}$$

typechecks, where $[\theta_A^i \mid i \in I] \in \mathcal{S}_-(!A)$ and $[\theta_B^i \mid i \in I] \in \mathcal{S}_+(!A)$. It also commutes as it typechecks for all sub-configurations of $[x^i \mid i \in I]$, all as constructions involved are monotone and preserve cardinal, by Lemma 6.1.12 this implies commutation. \square

It remains to prove that $!(-)$ is a \sim -functor:

Proposition 8.3.4. *We have a \sim -functor $!(-) : \text{NTCG} \rightarrow \text{NTCG}$.*

Proof. Preservation of copycat. We build an isomorphism $!(\mathfrak{c}_A) \approx \mathfrak{c}_{!A}$. This is done via Corollary 7.2.13, by constructing an order-isomorphism between $\mathcal{S}^+(!(\mathfrak{c}_A))$ and $\mathcal{S}^+(\mathfrak{c}_{!A})$ compatible with dom and cod and with display maps. First, notice that

$$\mathcal{S}^+(\mathfrak{c}_{!A}) \cong \mathcal{S}^+(!A) \tag{8.5}$$

by Lemma 7.3.14. From the definition and Lemma 7.3.14, $\mathcal{S}^+(!A)$ comprises

$$\theta : [x_A^i \parallel x_A^i \mid i \in I] \cong [y_A^j \parallel y_A^j \mid j \in J]$$

given by some bijection $\pi : I \simeq J$ and $\theta^i \parallel \theta^i : x_A^i \parallel x_A^i \cong_{\mathfrak{c}_A} y_A^j \parallel y_A^j$ – but this is clearly isomorphic to the data of $\pi : I \simeq J$ and a family of symmetries $\theta^i : x_A^i \cong_A y_A^j$, i.e. to $\mathcal{S}^+(!A)$. Together with (8.5), this yields a bijection $\mathcal{S}^+(\mathfrak{c}_{!A}) \simeq \mathcal{S}^+(!(\mathfrak{c}_A))$ which is easily checked to be an order-isomorphism preserving display maps as required.

Preservation of composition. We build an iso $\mathcal{S}^+(!(\tau \odot \sigma)) \cong \mathcal{S}^+(!(\tau) \odot !(\sigma))$.

By definition and Proposition 7.3.1, symmetries in $\mathcal{S}^+(!(\tau \odot \sigma))$ are given by a bijection $\pi : I \simeq J$, and for each $i \in I$, a symmetry

$$\theta_i^\tau \odot \theta_i^\sigma : x_i^\tau \odot x_i^\sigma \cong_{\tau \odot \sigma} y_{\pi(i)}^\tau \odot y_{\pi(i)}^\sigma$$

in $\mathcal{S}^+(\tau \odot \sigma)$. Then $\pi : I \simeq J$ together with the family of $\theta_i^\tau : x_i^\tau \cong_\tau y_{\pi(i)}^\tau$ yields

$$\theta^{! \tau} : [x_i^\tau \mid i \in I] \cong_{! \tau} [y_j^\tau \mid j \in J].$$

Likewise, set $I' = \{i \in I \mid \theta_i^\sigma \neq \emptyset\}$, and $\pi' : I' \simeq J'$ as the restriction of π . Then, $\pi' : I' \simeq J'$ together with the family of $\theta_i^\sigma : x_i^\sigma \cong_\sigma y_{\pi'(i)}^\sigma$ for $i \in I'$ yields

$$\theta^{! \sigma} : [x_i^\sigma \mid i \in I'] \cong_{! \sigma} [y_j^\sigma \mid j \in J'].$$

It is direct that $\theta^{! \sigma}$ and $\theta^{! \tau}$ are causally compatible, forming $\theta^{! \tau} \odot \theta^{! \sigma} \in \mathcal{S}^+(\tau \odot \sigma)$, and that the construction preserves unions and display maps.

Reciprocally, symmetries in $\mathcal{S}^+(\tau \odot \sigma)$ are given by causally compatible

$$\theta^{! \sigma} : x^{! \sigma} \cong_{! \sigma} y^{! \sigma}, \quad \theta^{! \tau} : x^{! \tau} \cong_{! \tau} y^{! \tau}$$

respectively given by $\pi^\sigma : I \simeq J$ and $\theta_i^\sigma : x_i^\sigma \cong_\sigma y_{\pi(i)}^\sigma$ for $i \in I$; and by $\pi^\tau : I' \simeq J'$ and $\theta_i^\tau : x_i^\tau \cong_\tau y_{\pi'(i)}^\tau$ for $i \in I'$. Now, for $i \in I$, θ_i^σ is non-empty. Since σ is negative, $(\theta_i^\sigma)_B$ is non-empty. So as $\theta^{! \sigma}$ and $\theta^{! \tau}$ are matching, we must have $i \in I'$ as well with θ_i^τ non-empty. This also entails that $\pi(i) = \pi'(i)$, so that $I \subseteq I', J \subseteq J'$ and $\pi \subseteq \pi'$. For $i \in I' \setminus I$, set $x_i^\sigma = \emptyset, y_i^\sigma = \emptyset$, and $\theta_i^\sigma = \emptyset$. Now it is direct that for any $i \in I$, the symmetries θ_i^σ and θ_i^τ are matching causally compatible, so that $\theta_i^\tau \odot \theta_i^\sigma \in \mathcal{S}^+(\tau \odot \sigma)$. So altogether we have a bijection $\pi : I \simeq J$, and for all $i \in I$ a symmetry

$$\theta_i^\tau \odot \theta_i^\sigma : x_i^\tau \odot x_i^\sigma \cong_{\tau \odot \sigma} y_i^\tau \odot y_i^\sigma,$$

that is, the data for a symmetry in $!(\tau \odot \sigma)$. It is a direct verification that this transformation preserves unions and display maps.

These operations are defined on symmetries but the same applies to configurations compatibly, informing an isomorphism $!(\tau \odot \sigma) \approx !\tau \odot !\sigma$ by Corollary 7.2.13. \square

Notice that we formulated the latter isomorphism directly on symmetries, whereas in many earlier proofs, we worked on configurations and only remarked that the same applies to symmetries. In most cases (*e.g.* tensor, with, composition), the construction applies to configurations and symmetries in exactly the same way. But the situation is really different for $!(-)$: this is the only construction that introduces new, non-trivial symmetries, and so treats the two levels genuinely differently.

A \sim -comonad. It remains to equip $!(-)$ with adequate structural morphisms, which as before will be obtained by lifting appropriate renamings:

Lemma 8.3.5. *Consider A a $--$ -board. Then, we have winning renamings:*

$$\begin{array}{ll} \text{der}_A : A^\perp \rightarrow (!A)^\perp & \text{dig}_A : (!A)^\perp \rightarrow (!A)^\perp \\ a \mapsto (0, a) & (i, (j, a)) \mapsto (\langle i, j \rangle, a) \end{array}$$

with $\langle -, - \rangle : \mathbb{N} \times \mathbb{N} \simeq \mathbb{N}$ a fixed bijection.

Proof. Direct verification. \square

Using these, for any $--$ -board A we may obtain strategies for the comonad structure

$$\mathbf{der}_A : !A \vdash A, \quad \mathbf{dig}_A : !A \vdash !!A$$

respectively called **dereliction** and **digging** after the following components in linear logic, simply by backward lifting: $\mathbf{der}_A = \mathfrak{D}_{\mathbf{der}_A}$ and $\mathbf{dig}_A = \mathfrak{D}_{\mathbf{dig}_A}$.

We shall prove that this makes $(!, \mathbf{der}, \mathbf{dig})$ a \sim -comonad on **NTCG**, but this requires us to prove a number of coherence laws. For the symmetric monoidal structure, recall that such coherence laws simply followed by equations on renamings, and concluding from functoriality of lifting (*i.e.* Proposition 8.1.9). However here, the coherence laws for renamings do not hold on the nose. Consider for example one of the unit laws:

$$\begin{array}{ccc} !A & & (!A)^\perp \\ \mathbf{dig}_A \downarrow & \searrow \mathfrak{C}_A & \uparrow \text{id} \\ !!A & \xrightarrow{\mathbf{der}_A} & !A \\ & & \downarrow \mathfrak{D}_{\mathbf{der}_A} \\ & & (!A)^\perp \end{array}$$

with the corresponding diagram on renamings on the right hand side, where $\mathfrak{D}_{\mathbf{der}_A}(i, a) = (i, \mathbf{der}_A(a)) = (i, (0, a))$. This diagram for renamings does not commute: the bottom-left path takes (i, a) to $(\langle i, 0 \rangle, a)$ instead of (i, a) as it should. But of course, this is no surprise: it is well-known that laws for exponentials require reindexing, and it is exactly the reason why we constructed an equivalence \approx on strategies authorizing reindexings.

Note then that the diagram above does commute up to *positive symmetry*, *i.e.*

$$\mathbf{dig}_A \circ \mathfrak{D}_{\mathbf{der}_A} \sim^+ \text{id}_{(!A)^\perp}$$

where recall that \sim^+ , defined in Definition 7.2.3, formalizes that the two maps are equal up to a *positive symmetry*, *i.e.* reindexing of Player moves.

This property will be sufficient to establish positive isomorphisms, exploiting:

Lemma 8.3.6. *Consider $f, f' : A \rightarrow B$ winning renamings such that $f \sim^+ f'$ (respectively, $g, g' : B^\perp \rightarrow A^\perp$ such that $g \sim^+ g'$). Then, $\mathfrak{C}_f \approx \mathfrak{C}_{f'}$ (resp. $\mathfrak{D}_g \approx \mathfrak{D}_{g'}$).*

Proof. The positive isos are simply identity maps between the underlying ess, and commutation with the display maps up to positive symmetry is obvious. \square

Now, we can prove the expected comonad laws.

Proposition 8.3.7. *We have $(!, \mathbf{der}, \mathbf{dig})$ a \sim -comonad on **NTCG**.*

Proof. Naturality. We show naturality of \mathbf{der}_A . Consider $\sigma : A \vdash B$. By Lemma 8.1.8, $\mathcal{S}^+(\sigma \odot \mathbf{der}_A) \cong \mathcal{S}^+(\sigma)$ with, for $\theta \in \mathcal{S}^+(\sigma)$, $\partial(\theta) = (0, \theta_A) \vdash \theta_B$ – where $(0, \theta_A)$ stands for $\{((0, a), (0, a')) \mid (a, a') \in \theta_A\}$. Likewise by Proposition 7.3.1 and Lemma 7.3.14, $\mathcal{S}^+(\mathbf{der}_B \odot !\sigma)$ comprises matching pairs of $\theta^{! \sigma} \in \mathcal{S}^+(!\sigma)$ and $\theta_B \parallel \theta_B \in \mathcal{S}^+(\mathfrak{C}_B)$, where the latter displays to $(0, \theta_B) \vdash \theta_B$. So by *matching*, $\theta^{! \sigma}$ spans only one copy of index 0, and thus must display to some $(0, \theta_A) \vdash (0, \theta_B)$. From this,

the required order-isomorphism follows easily.

Next, we show naturality of \mathbf{dig}_A . First, symmetries in $\mathcal{S}^+(\mathbf{!}\sigma \odot \mathbf{dig}_A)$ are somewhat hard to describe: by Lemma 8.1.8, they correspond to those

$$\theta : [[x^{i,j} \mid j \in J_i] \mid i \in I] \cong [[y^{k,l} \mid l \in L_k] \mid k \in K] \quad (8.6)$$

given by (1) a bijection $\pi : I \simeq K$; (2) for each $i \in I$, a bijection $\pi_i : J_i \simeq L_{\pi(i)}$; and (3) for each $i \in I, j \in J_i$, a symmetry $\theta^{i,j} : x^{i,j} \cong_{\sigma} y^{\pi(i), \pi_i(j)}$. This is displayed to

$$\begin{array}{ll} \theta_{!A} : & [x_A^{i,j} \mid \langle i, j \rangle \in X] \cong [y_A^{k,l} \mid \langle k, j \rangle \in Y] \\ \theta_B : & [[x_B^{i,j} \mid j \in J_i] \mid i \in I] \cong [[y_B^{k,l} \mid l \in L_k] \mid k \in K] \end{array}$$

where X comprises those $\langle i, j \rangle$ such that $i \in I$ and $j \in J_i$ and $x_A^{i,j}$ is non-empty, and likewise for Y . For the other side of the naturality diagram, by Proposition 7.3.1 and Lemma 7.3.14, symmetries in $\mathcal{S}^+(\mathbf{dig}_B \odot \mathbf{!}\sigma)$ correspond to pairs

$$\begin{array}{ll} \theta^{! \sigma} : & [x^m \mid m \in X] \cong [y^n \mid n \in Y] \\ \theta_{!!B} : & [[x^{i,j} \mid i \in I] \mid j \in J_i] \cong [[y^{k,l} \mid k \in K] \mid l \in L_k] \end{array}$$

which are matching, implying that X comprises those $\langle i, j \rangle$ such that $i \in I$ and $j \in J_i$; Y comprises those $\langle k, l \rangle$ such that $k \in K$ and $l \in L_k$, and $\theta^{! \sigma}$ first follows a bijection $\pi : I \simeq K$ and for each $i \in I$ a bijection $J_i \simeq K_{\pi(i)}$. This means that equivalently,

$$\theta^{! \sigma} : [[x^{\langle i, j \rangle} \mid j \in J_i] \mid i \in I] \cong [[y^{\langle k, l \rangle} \mid l \in L_k] \mid k \in K]$$

a symmetry in $\mathbf{!}\sigma$. This matches (8.6); this correspondence is an order-iso compatible with display maps so that $\mathbf{!}\sigma \odot \mathbf{dig}_A \approx \mathbf{dig}_B \odot \mathbf{!}\sigma$ as required by Corollary 7.2.13.

Coherence. First, we define renamings:

$$?der_A : (!A)^{\perp} \rightarrow (!!A)^{\perp}, \quad ?dig_A : (!!A)^{\perp} \rightarrow (!!!A)^{\perp}$$

simply copying the index for the outer $!$. Now, Proposition 8.3.4 (in particular, preservation of copycat) provides us with isomorphisms of ess

$$!\mathfrak{c}_A \cong \mathfrak{c}_{!A}, \quad !(\mathfrak{c}_{!A}) \cong \mathfrak{c}_{!!A}$$

which simply by verifying compatibility with display maps, yield positive isomorphisms

$$!(\mathfrak{w}_{der_A}) \approx \mathfrak{w}_{?der_A}, \quad !(\mathfrak{w}_{dig_A}) \approx \mathfrak{w}_{?dig_A}, \quad (8.7)$$

so that finally, noting the positive symmetries between renamings

$$\begin{array}{lll} dig_A \circ ?der_A & \sim^+ & id \\ dig_A \circ der_{!A} & \sim^+ & id \\ dig_A \circ dig_{!A} & \sim^+ & dig_A \circ ?dig_{!A} \end{array}$$

we get the coherence laws up to \approx by (8.7), Lemma 8.3.6 and Proposition 8.1.9. \square

This structure lets us define the particularly important following operation. If A and B are strict $--$ -boards and $\sigma : !A \vdash B$, then its **promotion** is

$$\sigma^! : !A \vdash !B \quad (8.8)$$

defined as $\sigma^! = !(\sigma) \odot \mathbf{dig}_A$, involving both digging and the functorial action of $!$. Promotion may be taken as a primitive, letting us state:

Corollary 8.3.8. *Consider the \sim -functor $! : \mathbf{NTCG}_s \rightarrow \mathbf{NTCG}$, together with: (1) the dereliction strategy $\mathbf{der}_A : !A \vdash A$ for strict A ; and (2) the promotion operation $\sigma^! : !A \vdash !B$ for all $\sigma : !A \vdash B$ with A and B strict.*

This yields a relative \sim -comonad, with respect to the inclusion $\mathbf{NTCG}_s \subseteq \mathbf{NTCG}$.

Proof. Straightforward. \square

But since we have a regular comonad, why emphasize this less standard relative comonad structure? Well, because as we shall see in Section 10.4, the comonad structure is not preserved by the collapse to the relational model: we shall see that the interplay between the collapse and $!$ only works well when A is strict.

Seely isomorphisms. First, we have an equality $!T = \mathbf{1}$, so the corresponding isomorphism is realized by the identity. We need one final isomorphism

$$\mathbf{mon}_{A,B} : !A \otimes !B \cong !(A \& B),$$

which we shall obtain via the following renamings:

Lemma 8.3.9. *Consider A, B two strict $--$ -boards. Then, we have winning renamings:*

$$\begin{aligned} \mathbf{mon}_{A,B} & : (!A \& B)^\perp \rightarrow (!A \otimes !B)^\perp \\ & (i, (j, m)) \mapsto (j, (i, m)) \\ \mathbf{mon}_{A,B}^* & : (!A \otimes !B)^\perp \rightarrow (!A \& B)^\perp \\ & (i, (j, m)) \mapsto (\langle i, j \rangle, (i, m)) \end{aligned}$$

satisfying $\mathbf{mon}_{A,B} \circ \mathbf{mon}_{A,B}^* \sim^+ \text{id}$ and $\mathbf{mon}_{A,B}^* \circ \mathbf{mon}_{A,B} \sim^+ \text{id}$.

Proof. Straightforward. \square

As expected, the required isomorphism $\mathbf{mon}_{A,B}$ is obtained by backward lifting, i.e.

$$\mathbf{mon}_{A,B} = \mathfrak{D}_{\mathbf{mon}_{A,B}}, \quad \mathbf{mon}_{A,B}^{-1} = \mathfrak{D}_{\mathbf{mon}_{A,B}^*},$$

inverses up to positive iso as follows from Lemmas 8.3.9, 8.3.6 and Proposition 8.1.9.

Note that these strategies cannot be defined via forward lifting, as neither $\mathbf{mon}_{A,B}$ nor $\mathbf{mon}_{A,B}^*$ would be valid renamings if typed without the $(-)^{\perp}$, both failing *receptive*.

It remains to show that the Seely isomorphism is compatible with promotion. For that, it is convenient to introduce one last map. Consider the winning renaming

$$\begin{aligned} \text{con}_A &: ((!A) \otimes (!A))^\perp \rightarrow (!A)^\perp \\ (i, (j, a)) &\mapsto (\langle i, j \rangle, a), \end{aligned}$$

called *contraction* (for A a $-$ -board), which induces $\mathbf{con}_A = \mathfrak{P}_{\text{con}_A} : !A \vdash !A \otimes !A$. Using this additional strategy, we first prove the following lemma:

Lemma 8.3.10. *Consider Γ, A, B strict $-$ -boards, and $\sigma : !\Gamma \vdash A, \tau : !\Gamma \vdash B$. Then, the following diagram commutes up to positive isomorphism:*

$$\begin{array}{ccc} !!\Gamma & \xrightarrow{!(\sigma, \tau)} & !(A \& B) \\ \text{con}_\Gamma \downarrow & & \downarrow \text{mon}_{A,B}^{-1} \\ !!\Gamma \otimes !!\Gamma & \xrightarrow{!\sigma \otimes !\tau} & !A \otimes !B \end{array}$$

Proof. By Lemma 8.1.8, the bottom-left path is $(!\sigma \otimes !\tau) \cdot \text{con}_\Gamma$, with $\text{ess } !\sigma \parallel !\tau$.

We now examine the $+$ -covered symmetries in the upper-right strategy. Via Proposition 7.3.1 and Lemma 7.3.14, those correspond to symmetries $\theta \in \mathcal{S}^+(!(\sigma + \tau))$ whose display on $!(A \& B)$ is in the image of $\text{mon}_{A,B}^*$, i.e. of the form

$$\begin{aligned} &\{(\langle 1, i \rangle, (1, a)), (\langle 1, j \rangle, (1, a')) \mid ((i, a), (j, a')) \in \theta_A\} \\ \cup &\{(\langle 2, i \rangle, (2, b)), (\langle 2, j \rangle, (2, b')) \mid ((i, b), (j, b')) \in \theta_B\} \end{aligned}$$

for $\theta_A \in \mathcal{S}(!A)$ and $\theta_B \in \mathcal{S}(!B)$; i.e. symmetries using copy indices of the form $\langle 1, - \rangle$ for A and $\langle 2, - \rangle$ for B . This means that θ also decomposes uniquely in this way, relating events of σ on indices $\langle 1, - \rangle$, and events of τ on indices $\langle 2, - \rangle$. This directly informs the required order-isomorphism with $!\sigma \parallel !\tau$; we omit further details. \square

With this lemma, we may finally prove compatibility of \mathbf{mon} with promotion:

Lemma 8.3.11. *Consider Γ, A, B strict $-$ -boards, and $\sigma : !\Gamma \vdash A, \tau : !\Gamma \vdash B$.*

Then, the following diagram commutes up to positive symmetry:

$$\begin{array}{ccc} !\Gamma & \xrightarrow{\langle \sigma, \tau \rangle^\dagger} & !(S \& T) \\ \langle \text{der}, \text{der} \rangle^\dagger \downarrow & & \downarrow m^{-1} \\ !(\Gamma \& \Gamma) & & !S \otimes !T \\ m^{-1} \swarrow & & \nearrow \sigma^\dagger \otimes \tau^\dagger \\ !\Gamma \otimes !\Gamma & & \end{array}$$

Proof. Via Lemma 8.3.10, it remains to show that the following diagram commutes

$$\begin{array}{ccccc}
 & & !!\Gamma & \xrightarrow{\text{con}_\Gamma} & !!\Gamma \otimes !!\Gamma \\
 \text{dig}_\Gamma \nearrow & & & & \nwarrow \text{dig}_\Gamma \otimes \text{dig}_\Gamma \\
 \Gamma & & & & \\
 \text{dig}_\Gamma \searrow & & !!\Gamma & \xrightarrow{! \langle \text{der}_\Gamma, \text{der}_\Gamma \rangle} & !(\Gamma \& \Gamma) \\
 & & & & \nwarrow \text{mon}_{\Gamma, \Gamma}^{-1} \\
 & & & & \Gamma \otimes \Gamma
 \end{array}$$

up to positive symmetry. We first observe that considering the winning renaming

$$\begin{aligned}
 \mu_\Gamma &: (!(\Gamma \& \Gamma))^\perp \rightarrow (!!\Gamma)^\perp \\
 (i, (j, m)) &\mapsto (i, (0, m)),
 \end{aligned}$$

we have $! \langle \text{der}_\Gamma, \text{der}_\Gamma \rangle \approx \mathfrak{P}_{\mu_\Gamma}$, as is established by a – by now routine – analysis of symmetries of $! \langle \text{der}_\Gamma, \text{der}_\Gamma \rangle$. The diagram then follows from Proposition 8.1.9, Lemma 8.3.6 and a direct computation of the composition of the renamings involved. \square

We may finally conclude the main result of this chapter:

Theorem 8.3.12. *There is \mathbf{NTCG} , a relative Seely \sim -category.*

Proof. The only components which were not detailed are naturality of $\mathbf{mon}_{A,B}$, and the coherence diagrams making $!(-)$ a strong symmetric monoidal \sim -functor. All these follow from Proposition 8.1.9 via a now routine methodology; we omit details. \square

And of course (see Appendix A.1 for the construction of the Kleisli category):

Corollary 8.3.13. *The Kleisli \sim -category \mathbf{NTCG}_1 is cartesian closed.*

In the sequel, we write \odot_1 for Kleisli composition, *i.e.* composition in \mathbf{NTCG}_1 .

8.3.2 Recursion

The last piece of structure that we shall construct in \mathbf{NTCG} is the interpretation of the *fixpoint combinator*, used to interpret recursion. Following the standard denotational semantics practice, this infinite strategy will be obtained as a least fixed point with respect to a partial order on concrete strategies.

First, we introduce this partial order:

Definition 8.3.14. *Consider A a board, and $\sigma, \tau : A$ winning strategies.*

We write $\sigma \triangleleft \tau$ if $\mathcal{C}(\sigma) \subseteq \mathcal{C}(\tau)$ – so in particular $|\sigma| \subseteq |\tau|$, and 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.

Strategies on A , ordered by \triangleleft , form a *directed complete partial order*, with respect to which all operations on strategies are easily shown to be continuous. We have:

Proposition 8.3.15. *Consider A a board, and D a directed set of strategies on A . Then D has a least fixed point for \triangleleft , written $\vee D$, and satisfying:*

$$\mathcal{E}^+(\vee D) = \bigcup_{\sigma \in D} \mathcal{E}^+(\sigma), \quad \mathcal{S}^+(\vee D) = \bigcup_{\sigma \in D} \mathcal{S}^+(\sigma).$$

Moreover, if every $\sigma \in D$ is winning, then so is $\vee D : A$.

Proof. We set $|\vee D| = \bigcup_{\sigma \in D} |\sigma|$, and take the union as well for all other components. All verifications are straightforward, and omitted. \square

For any $--$ -boards A and B this turns $\mathbf{NTCG}(A \vdash B)$ into a dcpo. We have:

Lemma 8.3.16. *The following operations on strategies are continuous:*

$$\begin{aligned} - \odot - & : \mathbf{NTCG}(B \vdash C) \times \mathbf{NTCG}(A \vdash B) \rightarrow \mathbf{NTCG}(A \vdash C) \\ - \otimes - & : \mathbf{NTCG}(A \vdash B) \times \mathbf{NTCG}(C \vdash D) \rightarrow \mathbf{NTCG}(A \otimes C \vdash B \otimes D) \\ \langle -, - \rangle & : \mathbf{NTCG}(\Gamma \vdash A) \times \mathbf{NTCG}(\Gamma \vdash B) \rightarrow \mathbf{NTCG}(\Gamma \vdash A \& B) \\ (-)^! & : \mathbf{NTCG}(!\Gamma \vdash A) \rightarrow \mathbf{NTCG}(!\Gamma \vdash !A) \end{aligned}$$

Proof. Straightforward. \square

The plan is then to follow the standard denotational semantics route, and define the recursion combinator via Kleene's fixed point theorem, as the least fixed point of a well-chosen functional continuous via the proposition above.

For that, we must deal with a minor inconvenience: Kleene's fixed point theorem applies to continuous functionals over a dcpo *with a least element*, but the dcpo $\mathbf{NTCG}(A)$ of winning strategies over a board A *does not* in general have a least element. Indeed, strategies on A cannot in general be empty, as they must have – by receptivity – events matching the negative minimal events of A . But “matching” is loose: they are free to *name* those events arbitrarily. We solve this as in [Castellan et al., 2019]: we choose one minimal $\perp_A : A$, serving as the canonical least element. Then for any $\sigma : A$, we pick an isomorphic $\sigma \approx \sigma^b : A$ s.t. $\perp_A \triangleleft \sigma^b$, obtained by renaming minimal events. We write \mathcal{D}_A for the pointed dcpo of strategies above \perp_A . For this new operation:

Lemma 8.3.17. *For any board A , the following operation is continuous:*

$$(-)^b : \mathbf{NTCG}(A) \rightarrow \mathcal{D}_A.$$

Proof. Straightforward. \square

As all operations involved are continuous, for any strict $--$ -board A we have

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

continuous, written in λ -calculus syntax relying on the constructions on strategies corresponding to the cartesian closed structure of $\mathbf{NTCG}_!$ [Lambek and Scott, 1988]. Here, we use $A \rightarrow B$ as a notation for $!A \multimap B$. Hence by Kleene’s fixpoint theorem,

$$\mathcal{Y}_A = \bigvee_{n \in \mathbb{N}} F^n(\perp) \in \mathbf{NTCG}_!(\mathsf{T}, (A \rightarrow A) \rightarrow A)$$

is a least fixpoint of F . Finally, in the presence of a strict $--$ -board Γ serving as context, we set $\mathcal{Y}_{\Gamma, A} = \mathcal{Y}_A \odot_! e_\Gamma$ where $e_\Gamma \in \mathbf{NTCG}_!(\Gamma, \mathsf{T})$ is the terminal morphism, yielding

$$\mathcal{Y}_{\Gamma, A} \in \mathbf{NTCG}_!(\Gamma, !(A \multimap A) \multimap A) \quad (8.9)$$

which will provide the interpretation of recursion throughout this monograph.

8.4 History and Related Work

Compact Closed Category. Following Joyal’s compact closed category of Conway games [Joyal, 1977], it is certainly expected that a category of unpolarized games should be compact closed – a notable exception to that is of course Blass games [Blass, 1992] which famously do not form a category. A detailed analysis of this so-called “Blass problem” in the context of Conway games may be found in [Melliès, 2004b].

For our particular technical setting of concurrent games on event structures, the compact closed structure was first described in [Castellan et al., 2017a] without symmetry, and extended to thin concurrent games in [Castellan et al., 2019].

Seely Category. Seely categories is a convenient categorical axiomatization of models of intuitionistic linear logic, in the presence of a product. Many categories of negative games form Seely categories, starting with *simple games* [Hyland, 1997] with respect to several choices of exponential modalities.

In the original presentation of thin concurrent games (first in conference version in [Castellan et al., 2015] then detailed in [Castellan et al., 2019]), we did not construct a model of intuitionistic linear logic. Indeed, we opted to construct a category called *Concurrent Hyland-Ong games*, with objects exactly the same *arenas* as in traditional Hyland-Ong games. We wanted a conservative extension of HO games: the appeal was to make it easier to adapt and extend the developments that led to their versatility and impact – and in particular, concepts such as *justifiers* and conditions on strategies such as *innocence*. Technically, this was done by building a tcg from any Hyland-Ong arena via a variant of the resource modality (inspired by [Harmer et al., 2007]) performing duplications not just at the root, but deep within the arena.

In developing probabilistic concurrent games in [Castellan et al., 2018b], we reverted to an AJM-style $!$, which made it easier to link with the weighted relational model. But then we (re-)discovered (Melliès, at least, was certainly aware of this from his work in asynchronous games) that justifiers, and in general the necessary ingredients to the Hyland-Ong machinery, were already available in concurrent games even with an AJM-style exponential! This is one of the advantages of the causal setting, as opposed to

traditional AJM games. In retrospect, this makes *Concurrent Hyland-Ong Games* misguided: they suffer from the same defect as traditional Hyland-Ong games (lacking a linear decomposition), while working in the Kleisli category of a Seely category of thin concurrent games gives you the best of both AJM and HO worlds.

Concerning *relative* Seely category: from **TCG** one can get an actual Seely category [Castellan and Clairambault, 2021]. But the general linear arrow construction is awkward, and the interest of the added expressiveness is unclear for semantic purposes. Finally, the full Seely category structure lacks a clean connection with relational semantics. *Relative* Seely categories give a good compromise for semantics of call-by-name.

Winning Strategies. Finally, the *payoff* mechanism and winning strategies we use here are an adaptation of similar mechanisms in [Melliès, 2005] (with a later refinement in [Melliès and Tabareau, 2007]) where they serve essentially the same purpose as here: they cope for the otherwise inherently *affine* nature of game semantics, making strategies more linear and identifying the *stopping* configurations of games, those corresponding points of the web in the relational model. Note that we first adapted this mechanism to concurrent games on event structures in [Clairambault and de Visme, 2020].

It might be unclear to the reader why a mechanism making strategies *linear* is appropriate for a model intended for the languages considered in this monograph, which are non-linear and include full recursion and divergence. Firstly, the term *linear* is not really adequate; we should rather say “following the linearity constraints”, which are made more liberal by using the exponential. Secondly, our notion of winningness does not force termination, it only constrains resource usage of successful terminating executions, if they exist. More precisely, our (new) restriction of winningness to *+-covered* configurations has the subtle effect to prevent from forcing Player to respond in diverging branches, say *i.e.* a diverging boolean.

Part III

Disentangling Parallelism and State

Introduction to Part III

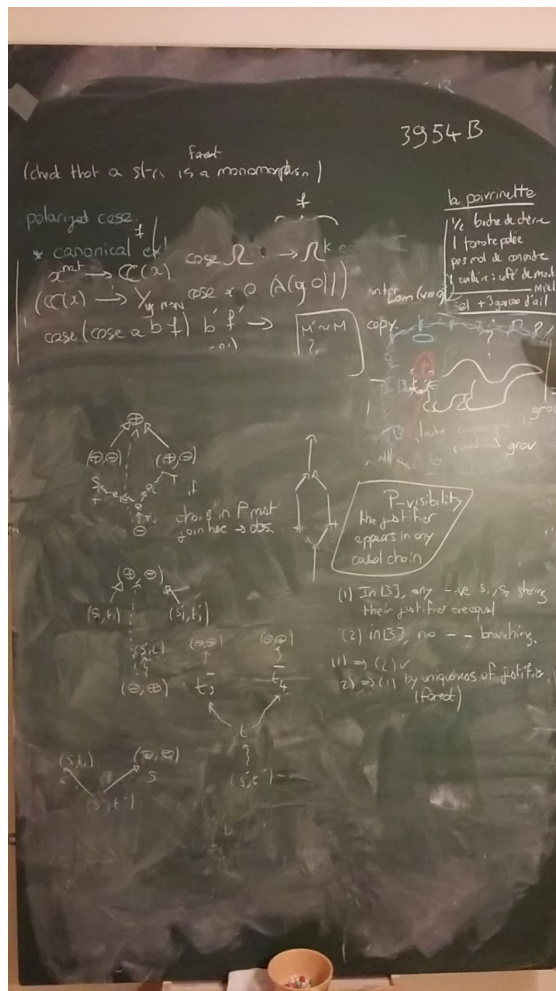


Figure 8.5: The Inception of Parallel Innocence

Now comes the core of this monograph: from **NTCG** we build a cartesian closed category $\rightarrow\text{-Strat}$ in which we *disentangle* parallelism and state, as in Section 1.7.

More precisely, we first show that $\rightarrow\text{-Strat}$ is an intensionally fully abstract denotational model for $\text{IA}_{//}$, by providing an interpretation-preserving functor

$$\mathcal{O}\text{-Unf} : \rightarrow\text{-Strat} \rightarrow \mathcal{O}\text{-Strat}$$

to non-alternating play-based strategies as described in Chapter 5. The category $\rightarrow\text{-Strat}$ is close to NTCG_1 but is not exactly that, for two reasons: firstly, because **NTCG** and $\mathcal{O}\text{-Strat}$ do not have the same objects, and there is no easy way to reconstruct an arena in the sense of Hyland-Ong game from a board – so the objects of $\rightarrow\text{-Strat}$ will remember both. Secondly, nothing in **NTCG** guarantees the logical well-bracketing used in $\mathcal{O}\text{-Strat}$, so that must be reimposed as well. This is detailed in Chapter 9.

Once $\rightarrow\text{-Strat}$ is constructed along with its interpretation of $\text{IA}_{//}$ and its unfolding to $\mathcal{O}\text{-Strat}$, we set to build our two conditions on strategies, *sequentiality* and *globularity*, respectively banning *parallelism* and *state*. We first focus on *globularity*. The heart of globularity is *parallel innocence*, a condition that removes state by banning certain causal patterns that are deemed inherently stateful. Parallel innocence can be regarded as a conservative extension of traditional Hyland-Ong innocence, in the sense that paired with sequentiality, we obtain a category (equivalent to) $\Downarrow\text{-Inn}$. As a conceptually important aside, we show that parallel innocence allows a simple functor to the relational model, compatibly with the structure of a model of intuitionistic linear logic. Finally we introduce *globularity*, which refines parallel innocence by adding a constraint with respect to Questions and Answers. This is detailed in Chapter 10.

Next, we focus on *sequentiality* – we show that in the presence of sequentiality, $\mathcal{O}\text{-Unf}$ can be refined to an *alternating* unfolding, an interpretation-preserving

$$\Downarrow\text{-Unf} : \rightarrow\text{-Seq} \rightarrow \Downarrow\text{-Strat}$$

which entails that $\rightarrow\text{-Seq}$ is intensionally fully abstract for IA . We also show that in combination with globularity, $\Downarrow\text{-Unf}$ yields an interpretation-preserving functor

$$\Downarrow\text{-Unf} : \rightarrow\text{-SeqGlob} \rightarrow \Downarrow\text{-InnWB}$$

proving $\rightarrow\text{-SeqGlob}$ intensionally fully abstract for PCF. This appears in Chapter 11.

Finally, it remains to establish that $\rightarrow\text{-Glob}$ is fully abstract for $\text{PCF}_{//}$. This is proved by an elaborate finite definability argument. To bring concurrent strategies closer to syntax, we must first show that globular strategies enjoy a so-called *meager form*, akin to the forest of *P-views* for traditional sequential innocence. This meager form provides the appropriate notion of *finiteness* or *size* of a globular strategy, on which the definability argument operates. We then show a *factorization* result, to the effect that any globular strategy factors into what is essentially a pure λ -term, and a globular strategy on a *first-order* arena. It remains to define those first-order strategies, which we do – up to the equivalence generated by the relational collapse. This is detailed in Chapter 12.

Altogether, this disentangles parallelism and state.

Chapter 9

The Causal Semantics of $\mathbf{IA}_{//}$ and its Unfolding

The purpose of this chapter is to construct $\rightarrow\text{-Strat}$ – the causal analogue of $\mathcal{O}\text{-Strat}$, the fully abstract model of $\mathbf{IA}_{//}$ – and describe the interpretation of $\mathbf{IA}_{//}$.

The outline is as follows. In Section 9.1, we will first describe the interpretation of $\mathbf{IA}_{//}$ in \mathbf{NTCG}_1 as constructed in Corollary 8.3.13. In Section 9.2, we shall refine \mathbf{NTCG}_1 into $\rightarrow\text{-Strat}$, adding a missing condition so that it unfolds to $\mathcal{O}\text{-Strat}$. In Section 9.3, we detail the unfolding to $\mathcal{O}\text{-Strat}$ and prove full abstraction for $\mathbf{IA}_{//}$.

9.1 Interpretation of $\mathbf{IA}_{//}$

We rely on the cartesian closed \sim -category \mathbf{NTCG}_1 from Part II (see Corollary 8.3.13). Thus, building on the standard interpretation of the simply-typed λ -calculus, it remains to give an interpretation to all basic datatypes and primitives of $\mathbf{IA}_{//}$.

9.1.1 Interpretation of $\mathbf{PCF}_{//}$

Interpretation of types. The cartesian closed structure of \mathbf{NTCG}_1 fixes the interpretation of the arrow constructor, which we must complete with an interpretation for the ground types \mathbb{B} , \mathbb{N} and \mathbb{U} . Earlier we have already provided $--$ -boards \mathbf{B} , \mathbf{N} and \mathbf{U} for those ground types, with for instance, for the booleans, the $--$ -board: with for instance

$$\mathbf{B} = \begin{array}{c} \mathbf{q}^- \\ \text{---} \quad \text{---} \\ \mathbf{tt}^+ \quad \text{---} \quad \mathbf{ff}^+ \end{array}$$

In the sequel, we shall refer to these as the **small** boards for ground types. The small ground boards are very natural, but in tension with the fact that in the play-based games





Figure 9.1: Configurations of basic strategies for PCF

of Part I, a question may be answered multiple times (possibly with the same answer). So as to avoid a cumbersome mismatch¹, we introduce an alternative interpretation of base types, where all possible answers are replicated. This is handled by:

Definition 9.1.1. Consider $V = \{a, b, c, \dots\}$ an at most countable set of values.

We define the strict $--$ -board $\text{ground}(V)$ as having essp:



with symmetries all bijections only changing copy indices. All symmetries are positive; negative symmetries are restricted to identities. Finally, the payoff function is:

$$\begin{aligned}
 \kappa : \mathcal{C}(\text{ground}(V)) &\rightarrow \{-1, 0, +1\} \\
 \emptyset &\mapsto +1 \\
 \{q_i^-, v_i^+\} &\mapsto 0 \\
 x &\mapsto -1 \quad \text{otherwise}
 \end{aligned}$$

Each value has countably many copies, with no conflict. Thus, when playing a value, a strategy must provide an accompanying copy index. Player may answer *several times*, either with the same or different values, but this is banned from strategies by the *winning* mechanism. However, Opponent may still answer several times (if the board appears

¹In non-alternating games, (logical) well-bracketing corrects this. The mismatch occurs only in the upcoming alternating unfolding, as in the alternating case well-bracketing is on strategies rather than on plays.

$$\begin{aligned}
\llbracket M; N \rrbracket &= \text{seq } \odot_! \langle \llbracket M \rrbracket, \llbracket N \rrbracket \rangle \\
\llbracket \text{if } M \ N_1 \ N_2 \rrbracket &= \text{if } \odot_! \langle \llbracket M \rrbracket, \llbracket N_1 \rrbracket, \llbracket N_2 \rrbracket \rangle \\
\llbracket \text{succ } M \rrbracket &= \text{succ } \odot_! \llbracket M \rrbracket \\
\llbracket \text{pred } M \rrbracket &= \text{pred } \odot_! \llbracket M \rrbracket \\
\llbracket \text{iszero } M \rrbracket &= \text{iszero } \odot_! \llbracket M \rrbracket \\
\llbracket \text{let } x = N \text{ in } M \rrbracket &= \text{let}_{\mathbf{X}, \mathbf{Y}} \odot_! \langle \llbracket N \rrbracket, \Lambda^! \langle \llbracket M \rrbracket \rangle \rangle
\end{aligned}$$





Figure 9.2: Basic interpretation clauses

Figure 9.3: Strategy for **pred**

in contravariant position) so that the model conveys information about the behaviour of programs under a context that answers several times.




We define the **large** ground boards as $\mathbf{U} = \text{ground}(\{\checkmark\})$, $\mathbf{B} = \text{ground}(\{\mathbf{tt}, \mathbf{ff}\})$ and $\mathbf{N} = \text{ground}(\mathbb{N})$ (the set of all natural numbers). As explained above, this extends to all *types* of PCF by setting $\llbracket A \rightarrow B \rrbracket = !\llbracket A \rrbracket \multimap \llbracket B \rrbracket$ using the constructions from Section 8.2.2, yielding for any type A a strict \multimap -board $\llbracket A \rrbracket$. Finally, a *context* $\Gamma = x_1 : A_1, \dots, x_n : A_n$ is interpreted as $\llbracket \Gamma \rrbracket = \&_{1 \leq i \leq n} \llbracket A_i \rrbracket$. We insist that \mathbf{IA}_{\parallel} can be interpreted with the ground types set as \mathbf{U}, \mathbf{B} and \mathbf{N} ; or as \mathbf{U}, \mathbf{B} and \mathbf{N} . We set the large boards as our default choice so as to match with non-alternating games, but we may sometimes refer to the small as well, because it allows for more succinct examples.




Interpretation of terms. A term $\Gamma \vdash M : A$ of PCF is interpreted as a strategy

$$\llbracket M \rrbracket \in \mathbf{NTCG}_!(\llbracket \Gamma \rrbracket, \llbracket A \rrbracket).$$

We invoke the cartesian closed structure of $\mathbf{NTCG}_!$ – writing $\odot_!$ for Kleisli composition, and $\Lambda^!$ for Kleisli currying. We skip the details of the interpretation of the λ -calculus combinators, which follows the standard lines of the interpretation of the simply-typed λ -calculus in a cartesian closed category [Lambek and Scott, 1988].

We must specify strategies for the term constructions that PCF adds to the simply-typed λ -calculus. For constants, we set $\llbracket \text{skip} \rrbracket : \mathbf{U}$, $\llbracket \mathbf{tt} \rrbracket : \mathbf{B}$, $\llbracket \mathbf{ff} \rrbracket : \mathbf{B}$ and $\llbracket n \rrbracket : \mathbf{N}$ as the obvious strategies replying immediately the corresponding value (with copy index 0). For all other combinators, the interpretation is in Figure 9.2. The strategies used are described in Figures 9.1, 9.3 and 9.4. The notation we employ is symbolic: the strategies comprise a branch for each instantiation of the copy index parameters i, j and the value parameters v, w, \dots . For instance, the strategy **if** has as configuration that shown in Figure 9.5 as nothing prevents Opponent from answering a question multiple times – this is in contrast with $\mathcal{G}\text{-Strat}$, where plays were always assumed to be well-bracketed. Recall also from Section 8.3.1 that $\langle -, \dots, - \rangle : \mathbb{N}^n \simeq \mathbb{N}$ denotes any bijection. Strategies specified symbolically as above carry a specific choice determining the copy indices of Player moves from those of Opponent moves. This choice does not matter up to positive isomorphism; the only thing that matters is that there is never any collision so that the display maps do satisfy the conditions of a concurrent strategy.

Figure 9.4: Typical configuration of **let**Figure 9.5: Example of a configuration of **if**Figure 9.6: Typical configuration of **plet**


 Figure 9.7: \mathbf{V}_w

 Figure 9.8: \mathbf{V}_r

 Figure 9.9: \mathbf{S}

For **let**, the strategy implements a *memoization* mechanism: it evaluates on \mathbf{X} first obtaining a value v , which it feeds to the function argument each call. The interpretation is completed with $\llbracket \mathcal{Y}_A M \rrbracket = \mathcal{Y}_{\Gamma, A} \odot_! \llbracket M \rrbracket$ for $\Gamma \vdash M : A \rightarrow A$, using (8.9).

Semantics of parallelism. This is extended to PCF_{\parallel} simply by adding the clause

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

via the strategy $\mathbf{plet}_{\mathbf{X}, \mathbf{Y}}$ with typical configurations as presented in Figure 9.6. It is interesting to compare this definition with the non-alternating plays in Figure 5.11: the diagram avoids listing all interleavings, hence staying much more faithful to the intuition behind the behaviour of this primitive.

9.1.2 Semantics of state

The interpretation of references and semaphores closely follows that in (\dagger -**Strat** and) \odot -**Strat**. For the types \mathbb{V} and \mathbb{S} , we first introduce strict $--$ -boards $\mathbf{V}_w, \mathbf{V}_r$ and \mathbf{S} as

$$\mathbf{V}_w = \bigotimes_{n \in \mathbb{N}} \mathbf{U}, \quad \mathbf{V}_r = \mathbf{N}, \quad \mathbf{S} = \mathbf{U} \& \mathbf{U}$$

with moves renamed as illustrated in Figures 9.7 (where in \mathbf{V}_w , all write requests are in pairwise conflict), 9.8 and 9.9. We set $\mathbf{V} = \mathbf{V}_w \& \mathbf{V}_r$, and $\llbracket \mathbb{V} \rrbracket = \mathbf{V}$, $\llbracket \mathbb{S} \rrbracket = \mathbf{S}$.

Stateless primitives. As in Section 4.2.2, only the interpretation of the primitives **newref** and **newsem** is effectful – other primitives only send various queries to references and semaphores, and as such are simple innocent strategies.

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

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

using the strategies of Figures 9.10, 9.11, 9.12, and 9.13. Finally, we set




Figure 9.10: assign




Figure 9.11: deref




Figure 9.12: grab




Figure 9.13: release

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

where $\llbracket M \rrbracket n$ is $\llbracket M \rrbracket$ applied to the constant strategy n (using the cartesian closed structure of \mathbf{NTCG}_1). It remains to define the semantics of **newref** and **newsem**.

Creation of reference and semaphores. For **newref** and **newsem**, again we follow in the footsteps of Section 4.2.2. However, in contrast with respect to Section 4.2.2, we must manage copy indices so as to satisfy the constraints of concurrent prestrategies.

For each $I \subseteq_f \mathbb{N}$ and $n \in \mathbb{N}$, we define languages of infinite words on $!V$ and $!S$:

$$\begin{aligned} \text{cell}_n^I &= \mathbf{r}_i^- \cdot n_0^+ \cdot \text{cell}_n^{I \uplus \{i\}} \mid \mathbf{wk}_i^- \cdot \checkmark_0^+ \cdot \text{cell}_k^{I \uplus \{i\}} & (i \notin I) \\ \text{lock}_0^I &= \mathbf{g}_i^- \cdot \checkmark_0^+ \cdot \text{lock}_1^{I \uplus \{i\}} \mid \mathbf{n}_i^- & (i \notin I) \\ \text{lock}_n^I &= \mathbf{g}_i^- \mid \mathbf{n}_i^- \cdot \checkmark_0^+ \cdot \text{lock}_0^{I \uplus \{i\}} & (i \notin I, n > 0) \end{aligned}$$



where symbols are moves in $!V$ and $!S$ respectively, separated via \cdot for readability. Note that positive moves should really include *two* copy indices: one for the outer bang, one for the value – we show only the latter, the former is always the same as the previous negative move. Here, $I \subseteq_f \mathbb{N}$ collects the copy indices already used so that each new copy index is fresh, ensuring that no single move appears twice. We set cell_n as (the prefix language of) cell_n^\emptyset and lock_n as (the prefix language of) lock_n^\emptyset . They may be viewed as event structures by setting events to be finite words, causality the prefix ordering, and setting any incomparable words to be in conflict. The two display maps

$$\begin{aligned} \partial_{\text{cell}} &: |\text{cell}| \rightarrow !V & \partial_{\text{lock}} &: |\text{lock}| \rightarrow !S \\ sa &\mapsto a & sa &\mapsto a \end{aligned}$$

only keep the last move. As configurations of cell and lock are sets of prefixes of a branch which is an alternating play in $!V$ or $!S$, we set $\mathcal{S}(\text{cell}_n)$ to comprise those bijections induced by plays $s_1 \cong_A s_2$ symmetric on A (for A being $!V$ or $!S$, see Definition 7.1.14).

Altogether, we get:

Proposition 9.1.2. *We have two concurrent prestrategies $\text{cell} : !V$ and $\text{lock} : !S$.*

Figure 9.14: Beginning of cell_0 Figure 9.15: Beginning of lock_0

We display a few early moves, of cell in Figure 9.14 and lock in Figure 9.15, with the convention that all moves in the same row are in pairwise conflict. We observe that both prestrategies fail *courtesy*: the immediate causal link $\checkmark_0^+ \rightarrow r_8^-$, for instance, would be illegal for a strategy. Note that for semaphores, in Figure 9.15, 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.

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

Proposition 9.1.3. *Consider $\sigma : !V \vdash A$ and $\tau : !S \vdash A$.*

Then, the two prestrategies $\sigma \odot \text{cell} : A$ and $\tau \odot \text{lock} : A$ are strategies.

Proof. We show it for cell , the proof for lock is the same. First, while cell is not *courteous*, it is *span-courteous* in the sense of Definition 7.3.10. Thus, by Proposition 6.2.22, $\sigma \odot \text{cell}$ is well-defined and a prestrategy. It remains to prove that $\sigma \odot \text{cell}$ is receptive and courteous. But by Proposition 6.4.12, that is the case if and only if it is invariant (up to strong iso) under composition with copycat . Thus we compute:

$$\begin{aligned} \mathfrak{C}_A \odot (\sigma \odot \text{cell}) &\cong (\mathfrak{C}_A \odot \sigma) \odot \text{cell} \\ &\cong \sigma \odot \text{cell}, \end{aligned}$$

so by Proposition 6.4.12, $\sigma \odot \text{cell}$ is receptive and courteous as required. \square

It is then easy to deduce the following fact:

Corollary 9.1.4. *Consider $\sigma \in \text{NTCG}(V, A)$ and $\tau \in \text{NTCG}(S, A)$.*

Then, $\sigma \odot \text{cell} \in \text{NTCG}(1, A)$ and $\tau \odot \text{lock} \in \text{NTCG}(1, A)$.

Proof. Again, we prove it for $\sigma \odot \text{cell}$. According to Proposition 8.2.15, we must show that it is *negative* and *winning* in the sense of Definition 8.2.14.

Negative is trivial as A is negative. For *winning*, take $x^\sigma \odot x^{\text{cell}} \in \mathcal{C}^+(\sigma \odot \text{cell})$. Then, $x^\sigma \in \mathcal{C}(\sigma)$ must be $+$ -maximal, otherwise it immediately contradicts $+$ -maximality of $x^\sigma \odot x^{\text{cell}}$ or minimality of $(x^{\text{cell}}, x^\sigma)$ in the sense of Proposition 7.3.9. Now in x^{cell} ,

maximal events (if any exists) must be positive, or contradict minimality of $(x^{\text{cell}}, x^{\sigma})$. But thus $\kappa_{!V}(x_{!V}^{\text{cell}}) = 0$, hence $\kappa_A(x_A^{\sigma}) \geq 0$ follows by σ winning. \square

We now complete the interpretation. Consider $\Gamma, x : \mathbb{V} \vdash M : \mathbb{X}$ with

$$\llbracket M \rrbracket \in \mathbf{NTCG}_!(\Gamma \& \mathbb{V}, \mathbb{X})$$

omitting some brackets. Using the cartesian closed structure of $\mathbf{NTCG}_!$, we consider

$$\Lambda_{\Gamma}^!(\llbracket M \rrbracket) \in \mathbf{NTCG}_!(! \mathbb{V}, ! \Gamma \multimap \mathbb{X})$$

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

$$\begin{aligned} \llbracket \text{newref } x := n \text{ in } M \rrbracket &= \Lambda_{\Gamma}^{!-1}(\Lambda_{\Gamma}^!(\llbracket M \rrbracket) \odot \text{cell}_n) \in \mathbf{NTCG}_!(\Gamma, \mathbb{X}) \\ \llbracket \text{newsem } x := n \text{ in } M \rrbracket &= \Lambda_{\Gamma}^{!-1}(\Lambda_{\Gamma}^!(\llbracket M \rrbracket) \odot \text{lock}_n) \in \mathbf{NTCG}_!(\Gamma, \mathbb{X}). \end{aligned}$$

concluding the interpretation of \mathbf{IA}_{\parallel} in \mathbf{NTCG} . It shall follow from this chapter that:


Theorem 9.1.5. *$\mathbf{NTCG}_!$ is adequate for \mathbf{IA}_{\parallel} .*

Proof. A consequence of the adequacy of $\mathcal{C}\text{-Strat}$ for \mathbf{IA}_{\parallel} with Theorem 9.3.9. \square

9.1.3 Discussion on the Interpretation

What does the interpretation of \mathbf{IA}_{\parallel} programs as concurrent strategies look like?

At first sight, the information on programs of \mathbf{IA}_{\parallel} that the interpretation exposes is conceptually close to that in $\mathcal{C}\text{-Strat}$: it is a representation of its interactive behaviour. But the causal information makes it more explicit: for instance, the concurrent strategy for the parallel composition operation is simply described by the diagram




which is very close to the *intention* behind the many interleavings of the corresponding non-alternating strategy. Of course that example is somewhat expected: it is the canonical example illustrating what is gained by moving from an *interleaving* to a *truly concurrent* model of concurrency. So beyond this example, let us consider various fragments of \mathbf{IA}_{\parallel} and provide examples illustrating the expressiveness of the model.

For instance, the interpretation of the following simple program:

$$\vdash \lambda f^{\mathbb{U} \rightarrow \mathbb{U}}. f (f \text{ skip}) : (\mathbb{U} \rightarrow \mathbb{U}) \rightarrow \mathbb{U} \rightarrow \mathbb{U}$$


yields the strategy whose behaviour is described, in symbolic notation, by



This shows that the traditional notion of *P-view* is really a way to recover the causal structure for sequential, purely functional programs. Concurrent games push this handle on causality way beyond purely functional sequential programs.


$$\vdash \text{newref } x \text{ in } \lambda f^{\mathbb{U} \rightarrow \mathbb{U}}. f(x := 1); \text{not (iszero !}x) : (\mathbb{U} \rightarrow \mathbb{U}) \rightarrow \mathbb{B}$$

This showcases a few important phenomena:




First of all, the example has conflict, *i.e.* non-determinism. Why is that, since the program above is in **IA**, a deterministic language? In fact, **IA** is only deterministic when observed by a sequential context: here, the model is more permissive and allows the argument function to simultaneously return and call its argument. This triggers a race in memory, which may be resolved in two incompatible ways: the write may win (leading to value **tt**), or the read may win (leading to value **ff**). Accordingly, the diagram has two maximal configurations, corresponding to the two resolutions of this race.

When drawing strategies for impure programs we avoid the symbolic representation used before, as it is very misleading. If Opponent duplicates \mathbf{q}_0^- in the diagram above, this is not going to “only” duplicate the moves that depend on it: there will now be two read requests and one write requests, leading to $3! = 6$ ways to resolve this three-party race. This will give rise to a complex pattern of conflicts too unwieldy to be pleasantly drawn in a single diagram, allowing 6 maximal configurations. One of them will be




in which by squinting appropriately, the trained eye will recognize an execution where the write triggered by \mathbf{q}_0^+ was treated first, then the read, then the write triggered by \mathbf{q}_1^+ – the only possible alternating linearization of this diagram follows this order.

Non-determinism. The example above also shows that a strategy may have two non-deterministic events with the same label. It is crucial that in this model, *conflict is not idempotent*: for instance, $\vdash M = \text{if choice then skip else skip} : \mathbb{U}$ yields a strategy



with two completely indistinguishable copies of the same computational event \checkmark_0^+ . Intuitively, there are two events because there are two derivations of $M \Downarrow \text{skip}$ – each event carries its full causal history. Keeping these events separate is sometimes important, because they may have different futures. For instance, in the strategy



the two events \mathbf{q}_0^+ look similar, but one will feed a value to the argument function while the other will not, as in the term $\vdash \lambda f. \text{if choice then } f \text{ skip else } f \perp$.

One may wonder if this very fine-grained non-deterministic branching information makes the model sound with respect to must-testing. The answer is: not directly, because non-deterministic branches not yielding any observable result are erased when performing the hiding. However, hiding can be amended to correct that: this is done in [Castellan et al., 2018a] for affine games – we shall however not follow that route here.

Observable behaviour. It should be clear from the examples above that the model is extremely intensional; it displays information that reflects the internal workings of the program and is by no means observable. A natural question is then: what parts of the concurrent strategy interpreting an IA_{\parallel} program is observable within IA_{\parallel} ? We shall provide an answer to that next, by unfolding to non-alternating game semantics.

9.2 Unfolding Causal to Non-Alternating Strategies

For the first time, we must simultaneously consider two kinds of strategies: *concurrent* strategies, as in, event structures; and non-alternating strategies – which are, of course, also concurrent! To avoid confusion, we shall say *causal* (pre)strategies for event structure based (pre)strategies, and *non-alternating* (pre)strategies otherwise.

Play-based strategies present programs by listing all their observable behaviour, presented as chronologically ordered *plays*. Causal strategies tell us more: *why* an action was performed, which actions are *independent*, also keeping track separately of non-deterministic branches. Intuitively, it is clear how to go from the former to the latter: simply forget the causal structure and retain only the induced linear orderings.

While conceptually clear, doing this concretely will require us to address a number of small but significant mismatches between the two settings.

9.2.1 A First Unfolding

Plays on games with symmetry. We start with linearizations of event structures.

Definition 9.2.1. Consider E an event structure.

A **linearization** on E is a sequence $s = s_1 \dots s_n$ which is:

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

We write $\mathcal{L}(E)$ for the set of linearizations on E .

If A is a board, then we also call $s \in \mathcal{L}(A)$ a **non-alternating play** on A , and write $\mathcal{C}\text{-Plays}(A) = \mathcal{L}(A)$. Now recall that a board also carries a notion of *symmetry* (see Section 7.1.1). As in Definition 7.1.14, this induces an equivalence on linearizations:

Definition 9.2.2. Let E be an ess and $s, t \in \mathcal{L}(E)$. We say that s and t are **symmetric**, written $s \cong_E t$, if s and t have the same length, and we have

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

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

It is straightforward to establish:

Lemma 9.2.3. Consider A a board.

Then, the relation \cong_A on $\mathcal{L}(A)$ is a length-preserving equivalence satisfying:

- preservation: if $s_1 \dots s_n \cong_A t_1 \dots t_p$, then for all $1 \leq i \leq n$, $\text{pol}_A(s_i) = \text{pol}_A(t_i)$,
- restriction: if $s_1 \dots s_n \cong_A t_1 \dots t_n$, then for all $1 \leq i \leq n$, $s_1 \dots s_i \cong_A t_1 \dots t_i$,
- simulation: if $s \cong_A t$ and $sa \in \mathcal{L}(A)$, then there is $tb \in \mathcal{L}(A)$ such that $sa \cong_A tb$.

Proof. Straightforward from Definitions 7.1.5 and 9.2.2. □

We shall rely on this to define our unfolding.

Linearizing strategies. As a first step towards the unfolding, we shall extract from any causal strategy σ on board A a subset of $\mathcal{U}\text{-Plays}(A)$ – a “non-alternating strategy on A ”, though no such formal object has been defined².

To do so, our main observation is the following:

Lemma 9.2.4. *Consider $\sigma : A$ a causal strategy on board A . Then we have:*

$$\begin{aligned} \partial_{\sigma} : \mathcal{L}(\sigma) &\rightarrow \mathcal{L}(A) = \mathcal{U}\text{-Plays}(A) \\ \varepsilon &\mapsto \varepsilon \\ sm &\mapsto \partial_{\sigma}(s)\partial_{\sigma}(m). \end{aligned}$$

a length-preserving monotone function.

Proof. It is immediate from *rule-abiding* and *locally injective* that this is well-defined. Monotonicity and preservation of length are by definition. \square

The overload of notations for ∂_{σ} should not cause any confusion.

The *pre-unfolding* of σ is then obtained simply by projecting $\mathcal{L}(\sigma)$ – all linearizations of σ – through ∂_{σ} . More precisely, we prove the following proposition:

Proposition 9.2.5. *Consider $\sigma : A$ a causal prestrategy on board A . Then, the set*

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

satisfies the following properties:

- non-empty: $\varepsilon \in \mathcal{U}\text{-Plays}(\sigma)$,
- prefix-closed: for all $t \in \mathcal{U}\text{-Plays}(\sigma)$, if $s \sqsubseteq t$, then $s \in \mathcal{U}\text{-Plays}(\sigma)$,

and furthermore, if σ is a strategy, then we additionally have:

- receptive: if $s \in \mathcal{U}\text{-Plays}(\sigma)$ and $sa^- \in \mathcal{L}(A)$, then $sa \in \mathcal{U}\text{-Plays}(\sigma)$,
- courteous: if $sabt \in \mathcal{U}\text{-Plays}(\sigma)$ and $\text{pol}(a) = +$ or $\text{pol}(b) = -$, then if $sbat \in \mathcal{L}(A)$, $sbat \in \mathcal{U}\text{-Plays}(\sigma)$.

Proof. *Non-empty, prefix-closed.* Obvious since $\varepsilon \in \mathcal{L}(\sigma)$, $\mathcal{L}(\sigma)$ is prefix-closed and ∂_{σ} is monotone and length-preserving.

Receptive. Consider $s \in \mathcal{U}\text{-Plays}(\sigma)$ and $sa^- \in \mathcal{L}(A)$. By definition, there is $u \in \mathcal{L}(\sigma)$ s.t. $s = \partial_{\sigma}(u)$. By definition of plays, $|u| \in \mathcal{C}(\sigma)$ and $\partial_{\sigma}(|u|) = |s|$ extends with a^- ; so by *receptive* there is $m \in |\sigma|$ s.t. $|u|$ extends with m and $\partial_{\sigma} m = a$. Hence, $um \in \mathcal{L}(\sigma)$ and $\partial_{\sigma}(um) = sa$, so that $sa \in \mathcal{U}\text{-Plays}(\sigma)$ as required.

Courteous. Consider $sabt \in \mathcal{U}\text{-Plays}(\sigma)$ and assume $\text{pol}(a) = +$ or $\text{pol}(b) = -$. By definition, there is $umnv \in \mathcal{L}(\sigma)$ such that $\partial_{\sigma}(u) = s$, $\partial_{\sigma}(m) = a$, $\partial_{\sigma}(n) = b$ and $\partial_{\sigma}(v) = t$. Now, of course we cannot have $n <_{\sigma} m$. If m, n are incomparable, then $unmv \in \mathcal{L}(\sigma)$ as well as needed. Otherwise, $m <_{\sigma} n$. But then, we must have $m \rightarrow_{\sigma} n$. Now, since $\text{pol}(m) = +$ or $\text{pol}(n) = -$, by *courteous* we have $a \rightarrow_A b$ as well, contradicting $sbat \in \mathcal{L}(A)$. \square

²Recall that non-alternating strategies are sets of plays with pointers on arenas, not sets of plays on boards.

Note that these properties are essentially those defining a non-alternating strategy (Definition 5.1.5), though based on $\mathcal{L}(A)$ with A a board rather than based on plays with pointers on an arena, *i.e.* as in a non-alternating version of AJM games.

It is tempting to attempt defining such a non-alternating version of AJM games, and set it as target for the unfolding of causal strategies. After all, thin concurrent games and AJM games rely on a similar mechanism for handling replication, via copy indices and symmetry, whereas \mathcal{U} -**Strat** rests on plays with pointers. However, it turns out that the AJM approach to uniformity breaks down beyond alternating deterministic strategies. An appealing alternative is suggested by Melliès’ *orbital games* [Melliès, 2003], but this has yet to be worked out. In this monograph, the non-alternating pre-unfolding above will only be an intermediary towards the unfolding to plays with pointers.

Notations. In this section, the application of the display map to plays

$$\partial_{\sigma} : \mathcal{L}(\sigma) \rightarrow \mathcal{U}\text{-Plays}(\sigma) \subseteq \mathcal{U}\text{-Plays}(A),$$

for a causal $\sigma : A$, is a central concept – it is thus worth introducing some specialized notations (inspired from our notations on configurations) to reason on such situations.

As usual, plays are ranged over via s, t, u, v , etc; while single moves are ranged over by a, b, m, n etc. For plays in a game, we often add a tag in subscript as in $s_A, t_A \in \mathcal{U}\text{-Plays}(A)$. For linearizations in a causal strategy, we add a similar tag as a superscript, with *e.g.* $s^{\sigma}, t^{\sigma} \in \mathcal{L}(\sigma)$. In that case, we write $s_A^{\sigma} \in \partial_{\sigma}(s^{\sigma}) \in \mathcal{U}\text{-Plays}(A)$.

9.2.2 Mixed Boards and Pointifixion

The above already gives a notion of unfolding of a causal strategy $\sigma : A$ on a board A as a set of non-alternating plays on A . However, the non-alternating game semantics of Chapter 5 is based on *plays with pointers*. How shall we link the two?

A causal strategy plays on a *board*, whereas non-alternating plays operate on *arenas*, which represent types without duplications. In principle, the arena may be obtained by quotienting the board, but it is not clear how to formalize this elegantly. So instead, we interpret types with a structure that stores *both* the arena and the board, along with the links between the two. Besides connecting the two structures, this will serve as a description of the “concrete” boards, that arise as the interpretation of types.

Mixed boards. This is captured by the notion of *mixed boards*:

Definition 9.2.6. A *mixed board* is $(A, \underline{A}, \text{lbl}_A)$ with A a strict board, \underline{A} an arena, and

$$\text{lbl}_A : |A| \rightarrow |\underline{A}|$$

a *label* function preserving polarities, satisfying the following requirements:

- rigid: lbl_A preserves and reflects minimality, and preserves \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_A ,



Figure 9.17: The strict board for $\mathbb{U} \rightarrow \mathbb{U}$ 

Figure 9.18: The arena

with additional components postponed until Section 12.1.1 (Definition 12.1.1).

Mixed boards make explicit how multiple moves in the board, kept apart by copy indices, match the same in the arena. Figures 9.17 and 9.18 together form a representative example, giving the mixed board interpreting the type $\mathbb{U} \rightarrow \mathbb{U}$, where the label function is clear. *Transparent* gives a concrete account of symmetries: as they must preserve the partial order and the label, they only exchange between interchangeable copies.

Constructions on mixed boards. All the basic arenas and boards used in the interpretation of ground types may be paired to yield mixed boards: for instance,

$$(\mathbf{X}, \underline{\mathbf{X}}, \text{lbl}_{\mathbf{X}})$$

pairs the “large” boards for $\mathbb{U}, \mathbb{B}, \mathbb{N}$ from Section 9.1 with the basic arenas from Section 3.2.1, via the obvious labelling map – the mixed boards for references and semaphores are likewise straightforward. So as to lift this to the arrow type, we first need:

Lemma 9.2.7. *Consider A a mixed board. We extend lbl_A to*

$$\text{lbl}_A : |\mathbf{!}A| \rightarrow |\underline{A}|$$

by setting $\text{lbl}_A(i, a) = \text{lbl}_A(a)$. Then, we still have:

- rigid: lbl_A preserves minimality, and preserves immediate causality \rightarrow ,
- transparent: for any $x, y \in \mathcal{C}(\mathbf{!}A)$ and bijection $\theta : x \simeq y$,
then $\theta \in \mathcal{S}(\mathbf{!}A)$ iff θ is an order-iso preserving lbl_A .

Proof. Rigid is obvious; we show transparent. If. Consider $x, y \in \mathcal{C}(\mathbf{!}A)$ and a bijection $\theta : x \simeq y$, assume θ is an order-iso preserving lbl_A . Decompose x and y as

$$x = [x_i \mid i \in I] \quad y = [y_j \mid j \in J]$$

via Lemma 8.3.1 – so that in particular for all $i \in I$, $x_i \neq \emptyset$ and for all $j \in J$, $y_j \neq \emptyset$. If I and J are empty then $\theta = \emptyset \in \mathcal{S}(\mathbf{!}A)$. Otherwise, fix $i \in I$. As A is strict, x_i has exactly one initial move (i, a) , take $\theta(i, a) = (j, b)$. Set $\pi(i) = j$; done for all $i \in I$ this yields $\pi : I \simeq J$, a bijection as θ induces a bijection between the initial moves of x and y . Now since θ is an order-iso, for all $(i, a') \in x$, we have $\theta(i, a') = (\pi(i), b')$ for

some $b' \in x_{\pi(i)}$. It is direct that θ restricts to an order-iso $\theta_i : x_i \simeq y_{\pi(i)}$ which still preserves labels. This concludes the proof that $\theta \in \mathcal{S}(!A)$ as required.

Only if. Straightforward. \square

Via this lemma, all the remaining constructions on arenas extend transparently to mixed boards, pairing the arena with the matching board construction. In particular,

$$\begin{aligned} (A, \underline{A}, \text{lbl}_A) \& (B, \underline{B}, \text{lbl}_B) &= (A \& B, \underline{A \& B}, \text{lbl}_{A \& B}) \\ (A, \underline{A}, \text{lbl}_A) \Rightarrow (B, \underline{B}, \text{lbl}_B) &= (!A \multimap B, \underline{A \Rightarrow B}, \text{lbl}_{A \Rightarrow B}) \end{aligned}$$

where the labels for mixed board constructions are given by

$$\begin{aligned} \text{lbl}_{A_1 \& A_2}(i, a) &= (i, \text{lbl}_{A_i}(a)) \\ \text{lbl}_{A \Rightarrow B}(2, b) &= (2, \text{lbl}_B(b)) \\ \text{lbl}_{A \Rightarrow B}(1, (i, a)) &= (1, \text{lbl}_A(a)) && (B \text{ well-opened}) \\ \text{lbl}_{A \Rightarrow B}(1, (b, (i, a))) &= (1, (b, \text{lbl}_A(a))) && (B = \&_{i \in I} B_i \text{ with } B_i \text{ well-opened}); \end{aligned}$$

we omit the immediate verifications that these satisfy the conditions for mixed boards.

Overall, this yields an interpretation of each \mathbf{IA}_{\parallel} type (or context) A as a mixed board $\llbracket A \rrbracket$. As a convention, we shall still let A, B, C etc range over mixed boards, and silently coerce them to the board components (as in “let $a \in A$ ”, etc). The arena component will be referred to explicitly via $\underline{A}, \underline{B}$, etc.

Pointifixion. Mixed boards are the adequate setting to link causal strategies with plays with pointers. First, any play on the board A yields a play with pointer on the arena \underline{A} :

Lemma 9.2.8. *Consider A a mixed board. Then, the function*

$$(\hat{\cdot}) : \mathcal{O}\text{-Plays}(!A) \rightarrow \mathcal{O}\text{-PrePlays}(\underline{A})$$

sending $s = s_1 \dots s_n \in \mathcal{O}\text{-Plays}(A)$ to \mathfrak{s} with components,

$$\begin{aligned} \text{moves:} & \text{ the sequence } t = t_1 \dots t_n \text{ with } t_i = \text{lbl}(s_i), \\ \text{pointers:} & \text{ for } 1 \leq i < j \leq n, t_j \text{ points to } t_i \text{ iff } s_i \multimap_A t_i. \end{aligned}$$

is well-defined.

Proof. For $s = s_1 \dots s_n \in \mathcal{O}\text{-Plays}(!A)$, \mathfrak{s} is $t_1 \dots t_n$, where t_j points to t_i iff $s_i \multimap_{!A} s_j$. We must show that every non-initial move has a pointer: if t_j is non-initial, then s_j is non-initial since lbl preserves minimality. So there is $a \multimap_{!A} s_j$. But since $s \in \mathcal{O}\text{-Plays}(!A)$, $\{s_1, \dots, s_j\}$ is a configuration thus down-closed, so $a = s_i$ for some $1 \leq i < j$. It follows that t_j has a pointer to t_i as required. \square

This is illustrated in Figure 9.19: on the lhs, we show a non-alternating play on $!(!(\mathbf{U} \multimap \mathbf{U}) \multimap \mathbf{B})$ – in this figure, each move carries a sequence of copy indices, from the outermost $!$ to the innermost³. On the rhs, we show the corresponding play

³This is different to the convention in most of this monograph where only the index for the innermost $!$ is shown, the others being recovered from the immediate dependency.




Figure 9.19: Recovering pointers

with pointers: it is obtained from the diagram on the lhs by making the immediate dependency from the game an explicit component of the play, then removing copy indices.

Crucially, this construction is invariant under symmetry:

Proposition 9.2.9. *Consider A a mixed board. Then, $(\hat{\cdot})$ yields a function*


$$(\hat{\cdot}) : \mathcal{G}\text{-Plays}(!A)/\cong_{!A} \rightarrow \mathcal{G}\text{-PrePlays}(\underline{A}),$$

injective and preserving length and prefix.

Proof. First, it preserves $\cong_{!A}$: if $s = s_1 \dots s_n \cong_{!A} t_1 \dots t_n = t$, then $\theta = \{(s_i, t_i) \mid 1 \leq i \leq n\} \in \mathcal{S}(!A)$; but by *transparent* (via Lemma 9.2.8), $\text{lbl}(s_i) = \text{lbl}(t_i)$ for all $1 \leq i \leq n$. Furthermore, by Lemma 7.1.6, θ is an order-isomorphism, so $s_i \rightarrow_A s_j$ iff $t_i \rightarrow_A t_j$ for all $1 \leq i < j \leq n$, hence \mathfrak{S} and $\hat{\mathfrak{T}}$ have the same pointers. Finally, the construction clearly preserves length and prefix as required.

For injectivity, take $s, t \in \mathcal{G}\text{-Plays}(!A)$ with $\mathfrak{S} = \hat{\mathfrak{T}}$, writing $s = s_1 \dots s_n$ and $t = t_1 \dots t_n$. Then, for all $1 \leq i \leq n$, $\theta_{s,t}^i = \{(s_j, t_j) \mid 1 \leq j \leq i\}$ is a bijection between configurations of A . But since $\mathfrak{S} = \hat{\mathfrak{T}}$, $\theta_{s,t}^i$ is actually an order-iso preserving lbl , so by *transparent* (via Lemma 9.2.8) we have $\theta_i \in \mathcal{S}(!A)$. Hence, $s \cong_{!A} t$. \square

This proposition makes explicit the folklore idea that *plays with pointers* are concrete representatives for equivalence classes of plays with copy indices. Drawing inspiration from [Danos et al., 1996], we call \mathfrak{S} the **pointifixion** of s – note however that in [Danos et al., 1996] the pointers must be *reconstructed* from the copy indices, whereas here, they are already in s as the immediate causal links from the board.

Figure 9.20: A configuration of $\sigma : !(\mathbf{U} \multimap \mathbf{U}) \multimap \mathbf{B}$

Note that though pointifixion was defined as

$$(\hat{\cdot}) : \mathfrak{U}\text{-Plays}(!A)/\cong_{!A} \rightarrow \mathfrak{U}\text{-PrePlays}(\underline{A}),$$

the same definition may be transparently applied to obtain, for all A ,

$$(\hat{\cdot}) : \mathfrak{U}\text{-Plays}(A)/\cong_A \rightarrow \mathfrak{U}\text{-PrePlays}(\underline{A}),$$

used without further details. Additionally, extensions of plays lift along these functions:

Lemma 9.2.10. *Consider A a mixed board. We have the following properties:*

- (1) *For all $ta \in \mathfrak{U}\text{-PrePlays}(\underline{A})$, $s \in \mathfrak{U}\text{-Plays}(A)$ such that $t = \mathfrak{s}$, there is $sa' \in \mathfrak{U}\text{-Plays}(A)$ such that $(\hat{\cdot})(sa') = ta$;*
- (2) *For all $ta \in \mathfrak{U}\text{-PrePlays}(\underline{A})$, $s \in \mathfrak{U}\text{-Plays}(!A)$ such that $t = \mathfrak{s}$, there is $sa' \in \mathfrak{U}\text{-Plays}(!A)$ such that $(\hat{\cdot})(sa') = ta$.*

Sketch. The proof exploits that since we have interpreted ground types duplicating answers, for any extension of $t \in \mathfrak{U}\text{-PrePlays}(\underline{A})$ one can find a “fresh” move in $!A$, with a new copy index, matching the extension. However, formalizing this requires the full definition of mixed boards which we prefer to postpone to Section 12.1.1. \square

Proposition 9.2.9 and Lemma 9.2.10 entail that plays with pointers on \underline{A} may be regarded *exactly* as non-alternating plays on the $---$ board $!A$ considered up to symmetry:

Corollary 9.2.11. *Consider A a mixed board. Then, $(\hat{\cdot})$ yields a bijection*

$$(\hat{\cdot}) : \mathfrak{U}\text{-Plays}(!A)/\cong_{!A} \simeq \mathfrak{U}\text{-PrePlays}(\underline{A}),$$

preserving length and prefix.

The above development concerns *non-alternating pre-plays with pointers* (Definition 5.1.1); recall that the *non-alternating plays with pointers* from which non-alternating game semantics are built are also (logically) well-bracketed (Definition 5.1.3).




Figure 9.21: Constructing a play in the unfolding

9.2.3 Unfoldings of a Causal Strategy

Next, we describe how to *unfold* a causal strategy to a non-alternating strategy.

More precisely, we shall provide *three* different unfoldings:

- (1) $\mathcal{U}\text{-Unf}_\bullet$ unfolds $\sigma : A$ to a non-alternating thread-strategy on \underline{A} ,
- (2) $\mathcal{U}\text{-Unf}_!$ unfolds $\sigma : !A$ to a non-alternating strategy on \underline{A} ,
- (3) $\mathcal{U}\text{-Unf}$ unfolds $\sigma : !A \vdash B$ to a non-alternating strategy on $\underline{A} \Rightarrow \underline{B}$,

for all mixed boards A and B .

Unfolding to a thread-strategy. For A a mixed board, we may now *unfold* a causal strategy $\sigma : A$ to a set $\mathcal{U}\text{-Unf}_\bullet(\sigma)$ of threads. Plays in $\mathcal{U}\text{-Unf}_\bullet(\sigma)$ are obtained from σ in three steps. First (1) form the *linearizations of configurations* of σ , i.e. $\mathcal{L}(\sigma)$ (Definition 9.2.1). Then (2), project those to $\mathcal{U}\text{-Plays}(A)$ by applying ∂_σ event by event. Then finally, (3) apply pointifixion. More formally, the definition reads as:

Definition 9.2.12. Consider A a mixed board, and $\sigma : A$ a causal strategy.

The *unfolding* of σ is the set $\mathcal{U}\text{-Unf}_\bullet(\sigma) = \{(\cdot)(s_A^\sigma) \in \mathcal{U}\text{-Plays}(\underline{A}) \mid s^\sigma \in \mathcal{L}(\sigma)\}$.

Observe that not all linearizations of σ contribute to the unfolding, only those which by display and pointifixion yield a (logically) well-bracketed play.

We illustrate this in Figures 9.20 and 9.21. Figure 9.20 displays a configuration of a causal σ on $!(U \multimap U) \multimap B$, first with our usual notational convention, then explicitly separating the internal events of σ from their display in the game. In Figure 9.21, we start with the sequence $abcdef \in \mathcal{L}(\sigma)$; then follow the two steps of the construction of the corresponding play of the unfolding: first by displaying it on A , then pointifixion. In Figure 9.21, we insist that though pointers appear in the intermediary step, they are not part of the structure of plays and only represent immediate causal dependency in the game. On the right hand side, pointers are now explicitly part of the play.

As expected, this unfolding yields a valid non-alternating thread-strategy on \underline{A} :

Proposition 9.2.13. *Consider $\sigma : A$ a causal strategy.*

Then, $\mathcal{U}\text{-Unf}_{\bullet}(\sigma) : \underline{A}$ is a thread-strategy in the sense of Definition 5.1.7.

Proof. Clearly it is a set of threads. *Non-empty* and *prefix-closed* are immediate.

Receptive: consider $t = (\hat{\cdot})(s_A^\sigma)$ for $s^\sigma \in \mathcal{L}(\sigma)$, with $tb^- \in \mathcal{U}\text{-Plays}_{\bullet}(\underline{A})$. By Lemma 9.2.10, there is $s_A^\sigma a^- \in \mathcal{U}\text{-Plays}(A)$ such that $(\hat{\cdot})(s_A^\sigma a^-) = tb^-$. Now, by Proposition 9.2.5, there is $s^\sigma m \in \mathcal{L}(\sigma)$ displaying to $s_A^\sigma a^-$, and *receptive* follows.

Courteous: consider $uabv = (\hat{\cdot})(u_A^\sigma a' b' v_A^\sigma)$ where a is positive or b negative, and assume $ubav \in \mathcal{U}\text{-Plays}_{\bullet}(\underline{A})$. It follows that $u_A^\sigma b' a' v_A^\sigma \in \mathcal{U}\text{-Plays}(A)$: if we had $a' \rightarrow_A b'$, this would mean that b points on a by definition of pointifixion. Hence, $u^\sigma b' a' v_A^\sigma \in \mathcal{U}\text{-Plays}(\sigma)$ by Proposition 9.2.5, whose pointifixion is $ubav$ as required. \square

Unfolding to a non-alternating strategy. Above, we have defined the unfolding of $\sigma : A$ as a thread-strategy $\mathcal{U}\text{-Unf}_{\bullet}(\sigma)$. Next we unfold $\sigma : !A$ to a full strategy:

Definition 9.2.14. *Consider A a mixed board, and $\sigma : !A$ a causal strategy.*

The unfolding of σ is the set $\mathcal{U}\text{-Unf}_{\dagger}(\sigma) = \{(\hat{\cdot})(s_{!A}^\sigma) \mid s^\sigma \in \mathcal{L}(\sigma)\}$.

With respect to Definition 9.2.12, the only difference is that we start with a causal strategy on $!A$ rather than A . The immediate consequence is that the unfolding process will give us *plays*, rather than merely threads. As before, we have:

Proposition 9.2.15. *Consider A a mixed board and $\sigma : !A$ a causal strategy.*

Then, $\mathcal{U}\text{-Unf}_{\dagger}(\sigma) : \underline{A}$ is a non-alternating strategy as in Definition 5.1.5.

Proof. Same proof as for Proposition 9.2.13. \square

If $\sigma : A$ is a causal strategy, we may unfold it as a thread-strategy with $\mathcal{U}\text{-Unf}_{\bullet}(\sigma) : \underline{A}$; or we may unfold its promotion as a full non-alternating strategy with $\mathcal{U}\text{-Unf}_{\dagger}(\sigma^!) : \underline{A}$. In the sequel we shall need to relate these two unfoldings: in particular, $\mathcal{U}\text{-Unf}_{\dagger}(\sigma^!)$ is a single-threaded strategy whose generating thread-strategy is $\mathcal{U}\text{-Unf}_{\bullet}(\sigma)$:

Lemma 9.2.16. *Consider causal $\sigma : A$, and $\sigma^! : !A$ its promotion.*

Then, $(\mathcal{U}\text{-Unf}_{\bullet}(\sigma))^! = \mathcal{U}\text{-Unf}_{\dagger}(\sigma^!)$.

Proof. For \subseteq , we show by induction on $n \in \mathbb{N}$ that for all $n \in \mathbb{N}$, $\mathcal{U}\text{-Unf}_{\bullet}(\sigma)^{(n)}$ as defined below Definition 5.1.7, is included in $\mathcal{U}\text{-Unf}_{\dagger}(\sigma^!)$. This boils down to the fact that whenever $s \in \mathcal{U}\text{-Unf}_{\dagger}(\sigma^!)$ and $t \in \mathcal{U}\text{-Unf}_{\bullet}(\sigma)$, then $s \sqcup t \subseteq \mathcal{U}\text{-Unf}_{\dagger}(\sigma^!)$.

Consider $u \in s \sqcup t$. We have $s = \hat{u}$ and $t = \hat{v}$ for $u = \partial_{\sigma^!}(u')$ and $v = \partial_{\sigma}(v')$ with

$$u' \in \mathcal{U}\text{-Plays}(\sigma^!), \quad v' \in \mathcal{U}\text{-Plays}(\sigma).$$

Now, u' visits only finitely many copies for the outermost $!$, so considering

$$w' = (i, v'_1) \dots (i, v'_n)$$

for $v' = v'_1 \dots v'_n$, with $i \in \mathbb{N}$ not visited by u' , we get $w' \in \mathcal{U}\text{-Plays}(\sigma^!)$ satisfying that writing $w = \partial_{\sigma^!}(w')$, we have $\hat{w} = \hat{v} = t$. But necessarily, u' and w' are disjoint; so we may form $z' \in \mathcal{U}\text{-Plays}(\sigma^!)$ the interleaving of u' and w' following the same scheduling

as $u \in s \sqcup t$, and by construction $u = \mathcal{Z}$ for $z = \partial_{\sigma^!}(z')$.

For \supseteq , consider $s \in \mathcal{O}\text{-Unf}_!(\sigma^!)$. By definition, $s = \hat{u}$ and $u = \partial_{\sigma^!}(u')$ such that $u' \in \mathcal{O}\text{-Plays}(\sigma^!)$. By definition, u' visits a finite set I of copy indices; each $i \in I$ yields a restriction $u'_i \in \mathcal{O}\text{-Plays}(\sigma)$. For each $i \in I$, we write $u_i = \partial_{\sigma}(u'_i)$ and $s_i = \hat{u}_i$ – by construction, $s_i \in \mathcal{O}\text{-Unf}_!(\sigma)$. Finally, s is an interleaving of all the s_i 's. \square

Unfolding from mixed board A to mixed board B . We wish to relate the Kleisli category $\mathbf{NTCG}_!$ to the cartesian closed category $\mathcal{O}\text{-Strat}$ – this induces two complications in the unfolding, departing from the simple case described above.

Firstly, the way the two hom-sets are constructed differs slightly. Indeed, a morphism σ from A to B in $\mathbf{NTCG}_!$ is a causal strategy on $!A \vdash B$, i.e. $(!A)^\perp \parallel B$ where $!A, B$ are causally independent. In contrast, non-alternating strategies from \underline{A} to \underline{B} have plays in $\underline{A} \Rightarrow \underline{B}$, where the arrow construction imposes a dependency between (the initial moves of) the two components. We must thus consider the currying of σ :

$$\Lambda(\sigma) : !A \multimap B.$$

Secondly, the configurations of a causal strategy $\sigma : !A \multimap B$ have at most one initial move in B – in contrast with non-alternating plays on $\underline{A} \Rightarrow \underline{B}$, in which arbitrarily many initial moves in \underline{B} may occur. This is solved by unfolding not $\Lambda(\sigma)$ but

$$\Lambda(\sigma)^! : !(A \multimap B),$$

its promotion, which will be the basis of the unfolding. This prompts:

Definition 9.2.17. Consider $\sigma : !A \vdash B$ a causal strategy. Its *unfolding* is:

$$\mathcal{O}\text{-Unf}(\sigma) = \mathcal{O}\text{-Unf}_!(\Lambda(\sigma)^!) : \underline{A} \Rightarrow \underline{B}.$$

It is a consequence of Proposition 9.2.15 that this is well-defined.

9.3 $\rightarrow\text{-Strat}$ and Full Abstraction

We aim to prove that the unfolding above is as expected compatible with all operations used in the interpretation of \mathbf{IA}_{\parallel} , which shall bridge the two interpretations.

But we shall soon see that this will require us to add one further condition refining $\mathbf{NTCG}_!$ to $\rightarrow\text{-Strat}$, dealing with the fact that non-alternating plays with pointers are (logically) well-bracketed while arbitrary linearizations might not be.

9.3.1 Unfolding the Categorical Structure

We examine how unfolding preserves the categorical structure of $\mathbf{NTCG}_!$.

Unfolding of copycat. First, we unfold the identities of \mathbf{NTCG}_1 . Here, not much is happening conceptually; however technically the respective definitions of the causal and non-alternating copycat strongly differ: while the causal copycat comes from an explicit, concrete description (Definition 6.4.1), the non-alternating copycat is in a sense the “free asynchronous completion” of the alternating copycat (Definition 5.2.5). This indirect definition makes it quite unwieldy to reason with.

Proposition 9.3.1. *Consider A a mixed board.*

Then, $\mathbf{der}_A : !A \vdash A$ unfolds to the non-alternating copycat strategy, $\mathcal{O}\text{-}\mathfrak{c}_A$.

Proof. Recall from Definition 5.2.5 and Lemma 5.1.10 that the plays of $\mathcal{O}\text{-}\mathfrak{c}_A$ are those obtained from $\Downarrow\text{-}\mathfrak{c}_A$ by *receptivity* and *courtesy* in the sense of Definition 5.1.5.

By Lemma 9.2.16, it suffices to prove the equality

$$\mathcal{O}\text{-}\mathbf{Unf}_*(\Lambda(\mathbf{der}_A)) = \mathcal{O}\text{-}\mathfrak{c}_A^*,$$

where $\mathcal{O}\text{-}\mathfrak{c}_A^*$ denotes the set of *threads* of $\mathcal{O}\text{-}\mathfrak{c}_A$.

For \subseteq consider $s \in \mathcal{O}\text{-}\mathbf{Unf}_*(\Lambda(\mathbf{der}_A))$, meaning that there are $s = \hat{\tau}$ with $t = \partial_{\Lambda(\mathbf{der}_A)}(u)$ for some $u \in \mathcal{L}(\mathfrak{c}_A)$. Now, one can show that there is $v \in \Downarrow\text{-}\mathcal{L}(\mathfrak{c}_A)$ and w^- a sequence of moves of \mathfrak{c}_A negative in $A \vdash A$, such that

$$vw^- \in \mathcal{O}\text{-}\mathbf{Plays}(\mathfrak{c}_A)$$

has the same moves as u , and s.t. u may be obtained from vw^- by *courtesy*, such that (the pointifixon of) all intermediate plays are logically well-bracketed. Considering

$$t' = \partial_{\Lambda(\mathbf{der}_A)}(v), \quad s' = \hat{\tau}', \quad t'' = \partial_{\Lambda(\mathbf{der}_A)}(vw^-), \quad s'' = \hat{\tau}'',$$

it is immediate that $t' \in \Downarrow\text{-}\mathfrak{c}_A$. Thus, $t' \in \mathcal{O}\text{-}\mathfrak{c}_A$, and so $t'' \in \mathcal{O}\text{-}\mathfrak{c}_A$ by receptivity. But then, by construction s may be obtained from s'' by *courtesy*, thus $s \in \mathcal{O}\text{-}\mathfrak{c}_A$.

For \supseteq , consider $s \in \mathcal{O}\text{-}\mathfrak{c}_A$ a thread. By definition, s is obtained by *courtesy* and *receptivity* from some $s' \in \Downarrow\text{-}\mathfrak{c}_A$. Writing $s' \in \mathcal{O}\text{-}\mathbf{Plays}(\underline{A}_1 \Rightarrow \underline{A}_2)$, for all positive-ending $t' \sqsubseteq s'$, $t' \upharpoonright \underline{A}_1 = t' \upharpoonright \underline{A}_2$. By Corollary 9.2.11, there is $u \in \Downarrow\text{-}\mathbf{Plays}(A)$ s.t. $t' \upharpoonright \underline{A}_1 = \hat{u}$. We may then define $v \in \Downarrow\text{-}\mathcal{L}(\mathfrak{c}_A)$ by interleaving two copies of u in the obvious way, s.t. $v \upharpoonright A_1 = v \upharpoonright A_2 = u$. From the definition, it is direct that writing $w = \partial_{\Lambda(\mathbf{der}_A)}(v)$, $s' = \hat{w}$. By Proposition 9.2.5, applying to w instances of *courtesy* and *receptivity* as for $s' \rightsquigarrow s$, we get $z \in \mathcal{O}\text{-}\mathbf{Plays}(\mathfrak{c}_A)$ s.t. $s = \hat{\circ}(\partial_{\Lambda(\mathbf{der}_A)}(z))$. \square

The definition of $\mathcal{O}\text{-}\mathfrak{c}_A$ via saturation in Definition 5.2.5 is unpleasant: to prove that a play is in $\mathcal{O}\text{-}\mathfrak{c}_A$ we must provide a sequence of applications of *courtesy* and *receptivity* from a play of $\Downarrow\text{-}\mathfrak{c}_A$. It seems hard to avoid while staying in the context of pointer games. In contrast, \mathbf{der}_A offers via unfolding a concrete definition of $\mathcal{O}\text{-}\mathfrak{c}_A$.

Well-bracketing. Unfolding preserves the identity; but composition is *not* preserved: for $\sigma \in \mathbf{NTCG}_1(A, B)$ and $\tau \in \mathbf{NTCG}_1(B, C)$, we may not have

$$\mathcal{O}\text{-}\mathbf{Unf}(\tau \odot_1 \sigma) \subseteq \mathcal{O}\text{-}\mathbf{Unf}(\tau) \odot \mathcal{O}\text{-}\mathbf{Unf}(\sigma).$$

The issue is that unfolding only keeps (logically) well-bracketed plays, but well-bracketed plays in the unfolding of $\tau \odot \sigma$ may very well originate in pre-plays in the unfoldings of σ and τ that fail well-bracketing. In order to obtain this inclusion, we shall have to put an additional restriction on causal strategies:

Definition 9.3.2. Consider $\sigma : !A \vdash B$ a causal strategy. It is **well-bracketed** iff for all $sm^+ \in \mathcal{L}(\sigma)$, if $(\hat{\cdot})(\partial_{\Lambda(\sigma)}(s))$ is (logically) well-bracketed, then so is $(\hat{\cdot})(\partial_{\Lambda(\sigma)}(sm))$.

Alternatively, σ is well-bracketed iff the unfolding only “cuts at Opponent moves”. We skip the routine but lengthy verification of the following fact:

Theorem 9.3.3. There is $\rightarrow\text{-Strat}$, a cartesian closed \sim -category with objects mixed boards and morphisms those well-bracketed $\sigma \in \mathbf{NTCG}_1(A, B)$, which supports the interpretation of \mathbf{IA}_{\parallel} . Moreover there is an obvious forgetful \sim -functor

$$\rightarrow\text{-Strat} \rightarrow \mathbf{NTCG}_1$$

sending a mixed board to the underlying $-$ -board, leaving morphisms unchanged.

Preservation of composition. The models $\rightarrow\text{-Strat}$ and $\mathcal{O}\text{-Strat}$ differ fundamentally in that one is causal while the other is based on interleavings, but that in itself is not hard to manage, via the mechanism of linearizations.

The real challenge is that the compositional mechanisms of these two models follow significantly different principles: (1) $\rightarrow\text{-Strat}$ is constructed via the Kleisli category of the linear model \mathbf{NTCG} , while $\mathcal{O}\text{-Strat}$ is inherently non-linear; (2) $\rightarrow\text{-Strat}$ uses explicit copy indices, while $\mathcal{O}\text{-Strat}$ uses pointers; and (3) the hom-set of $\rightarrow\text{-Strat}$ is the board $!A \vdash B$ with no causal dependency between A and B , whereas $\mathcal{O}\text{-Strat}$ uses the arrow construction $\underline{A} \Rightarrow \underline{B}$ (which seems unavoidable for a pointer-based model). None of these are particularly challenging, but the three are deeply intertwined and must be managed simultaneously in the compatibility of unfolding with composition⁴.

We address the two directions of the unfolding separately. First:

Lemma 9.3.4. Consider $\sigma \in \rightarrow\text{-Strat}(A, B)$ and $\tau \in \rightarrow\text{-Strat}(B, C)$.

Then, we have the inclusion $\mathcal{O}\text{-Unf}_*(\Lambda(\tau \odot \sigma^!)) \subseteq \mathcal{O}\text{-Unf}(\tau) \odot \mathcal{O}\text{-Unf}(\sigma)$.

Proof. For $s \in \mathcal{O}\text{-Unf}_*(\Lambda(\tau \odot \sigma^!))$, by definition $s = (\hat{\cdot})(\partial_{\Lambda(\tau \odot \sigma^!)}(s'))$ for some $s' \in \mathcal{O}\text{-Plays}(\tau \odot \sigma^!)$. There is $u' \in \mathcal{O}\text{-Plays}(\tau \otimes \sigma^!)$ with visible restriction s' , and

$$u'_{\sigma^!} \in \mathcal{O}\text{-Plays}(\sigma^!) = \mathcal{O}\text{-Plays}(!\sigma), \quad u'_{\tau} \in \mathcal{O}\text{-Plays}(\tau).$$

are projections obtained applying move-by-move the (partial) projections of moves of the interaction to moves of the compound strategies as defined just above Lemma 6.2.14.

Now, u' is displayed to $v' \in \mathcal{O}\text{-Plays}(!A \parallel !B \parallel C)$ via $\partial_{\tau \otimes \sigma}$. There is an implicit dependency between moves minimal in the game: moves minimal in $!B$ depend on the minimal move in C . Likewise, there is a unique way to assign moves minimal in $!A$ to

⁴Ideally, one would deal with them separately by introducing a separate model: a non-alternating play-based games model based on copy indices rather than pointers – but we leave this out of scope.

those minimal in $!B$ in a way compatible with $\partial_{! \Lambda(\sigma)}(u'_{\sigma!}) \in \mathcal{O}\text{-Plays}(!A \multimap B)$. We form $w' \in \mathcal{O}\text{-I}(\underline{A}, \underline{B}, \underline{C})$ by first applying labels to v' as in Proposition 9.2.9, extended with this assignment. Now from the constructions, direct verifications entail

$$w' \upharpoonright \underline{A}, \underline{B} \in \mathcal{O}\text{-Unf}_!(\Lambda(\sigma)^!), \quad w' \upharpoonright \underline{B}, \underline{C} \in \mathcal{O}\text{-Unf}_!(\Lambda(\tau)^!), \quad w' \upharpoonright \underline{A}, \underline{C} = s,$$

thus w' is a witness for $s \in \mathcal{O}\text{-Unf}(\tau) \odot \mathcal{O}\text{-Unf}(\sigma)$ as required. \square

These “direct verifications” include, in particular, that the restrictions of w' are (logically) well-bracketed. That w' is itself well-bracketed is established by immediate induction, using that σ and τ are well-bracketed. Now, we focus on the other inclusion:

Lemma 9.3.5. *Consider $\sigma \in \mathbf{NTCG}_!(A, B)$, $\tau \in \mathbf{NTCG}_!(B, C)$.*

Then, threads in $\mathcal{O}\text{-Unf}(\tau) \odot \mathcal{O}\text{-Unf}(\sigma)$ are in $\mathcal{O}\text{-Unf}_!(\Lambda(\tau \odot \sigma^!))$.

Proof. Consider $s \in \mathcal{O}\text{-Unf}(\tau) \odot \mathcal{O}\text{-Unf}(\sigma)$ a thread, so there is $u \in \mathcal{O}\text{-I}(\underline{A}, \underline{B}, \underline{C})$ s.t. $u \upharpoonright \underline{A}, \underline{B} \in \mathcal{O}\text{-Unf}(\sigma)$, $u \upharpoonright \underline{B}, \underline{C} \in \mathcal{O}\text{-Unf}(\tau)$, and $u \upharpoonright \underline{A}, \underline{C} = s$. By hypothesis,

$$u \upharpoonright \underline{A}, \underline{B} = \hat{\cdot}(\partial_{\Lambda(\sigma)^!}(v^\sigma)), \quad u \upharpoonright \underline{B}, \underline{C} = \hat{\cdot}(\partial_{\Lambda(\tau)^!}(v^\tau))$$

for $v^\sigma \in \mathcal{O}\text{-Plays}(!\sigma) = \mathcal{O}\text{-Plays}(\sigma^!)$ and $v^\tau \in \mathcal{O}\text{-Plays}(!\tau)$; but as v^σ has one initial move, we consider $v^\tau \in \mathcal{O}\text{-Plays}(\tau)$ with $u \upharpoonright \underline{B}, \underline{C} = \hat{\cdot}(\partial_{\Lambda(\tau)}(v^\tau))$. We also have:

$$w^\sigma = \partial_{\sigma^!}(v^\sigma) \in \mathcal{O}\text{-Plays}(!A \vdash !B), \quad w^\tau = \partial_\tau(v^\tau) \in \mathcal{O}\text{-Plays}(!B \vdash C),$$

and since $u \upharpoonright \underline{A}, \underline{B}$ and $u \upharpoonright \underline{B}, \underline{C}$ coincide on \underline{B} , it is a direct verification that

$$\hat{\cdot}(w^\sigma \upharpoonright !B) = u \upharpoonright \underline{B} = \hat{\cdot}(w^\tau \upharpoonright !B),$$

so that $w^\sigma \upharpoonright !B \cong_{!B} w^\tau \upharpoonright !B$ by Proposition 9.2.9. Now, writing configurations

$$x^\sigma = |v^\sigma|, \quad x^\tau = |v^\tau|,$$

this informs a symmetry $\theta : x_{!B}^\sigma \cong_{!B} x_{!B}^\tau$. The difficulty is that this symmetry might not be the identity, preventing us from synchronizing x^σ and x^τ .

So additionally, we note that u informs a linearization of the composite bijection

$$\varphi[x^\sigma, \theta, x^\tau] : x^\sigma \parallel x_C^\tau \xrightarrow{\partial_\sigma \parallel x_C^\tau} x_{!A}^\sigma \parallel x_{!B}^\sigma \parallel x_C^\tau \xrightarrow{x_{!A}^\sigma \parallel \theta \parallel x_C^\tau} x_{!A}^\sigma \parallel x_{!B}^\tau \parallel x_C^\tau \xrightarrow{\partial_\tau^{-1} \parallel x_C^\tau} x_{!A}^\sigma \parallel x^\tau,$$

witnessing the fact that the triple $(x^\sigma, \theta, x^\tau)$ is causally compatible in the sense of Definition 7.4.3. We write $r \in \mathcal{O}\text{-Plays}(\varphi[x^\sigma, \theta, x^\tau])$ for this linearization; it projects to $\partial(r) \in \mathcal{O}\text{-Plays}(!A \parallel !B \parallel C)$ in the obvious way, and we have

$$\hat{\cdot}(\partial(r)) = \bar{u} \in \mathcal{O}\text{-Plays}(\underline{A} \times \underline{B} \times \underline{C})$$

where \bar{u} is u without the pointers of initial moves in \underline{A} and \underline{B} .

Thanks to causal compatibility, Proposition 7.4.4 provides us with $y^\sigma \in \mathcal{C}(\sigma^!)$ and $y^\tau \in \mathcal{C}(\tau)$ causally compatible configurations, along with symmetries

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

such that we have $\varphi_{!A}^{\sigma} \in \mathcal{S}_-(!A)$, $\varphi_C^{\tau} \in \mathcal{S}_+(C)$, and the following diagram commutes

$$\begin{array}{ccccccc}
 x^{\sigma} \parallel x_C^{\tau} & \xrightarrow{\partial_{\sigma!} \parallel x_C^{\tau}} & x_{!A}^{\sigma} \parallel x_{!B}^{\sigma} \parallel x_C^{\tau} & \xrightarrow{x_{!A}^{\sigma} \parallel \theta \parallel x_C^{\tau}} & x_{!A}^{\sigma} \parallel x_{!B}^{\tau} \parallel x_C^{\tau} & \xrightarrow{x_{!A}^{\sigma} \parallel \partial_{\tau}^{-1}} & x_{!A}^{\sigma} \parallel x^{\tau} \\
 \searrow \varphi^{\sigma} \parallel \varphi_C^{\tau} & & \searrow \varphi_{!A}^{\sigma} \parallel \varphi_{!B}^{\sigma} \parallel \varphi_C^{\tau} & & \searrow \varphi_{!A}^{\sigma} \parallel \varphi_{!B}^{\tau} \parallel \varphi_C^{\tau} & & \searrow \varphi_{!A}^{\sigma} \parallel \varphi^{\tau} \\
 y^{\sigma} \parallel y_C & \xrightarrow{\partial_{\sigma} \parallel y_C} & y_{!A} \parallel y_{!B} \parallel y_C & \xrightarrow{y_{!A} \parallel \partial_{\tau}^{-1}} & y_{!A} \parallel y^{\tau} & &
 \end{array}$$

following which the linearization $r \in \mathcal{U}\text{-Plays}(\varphi[x^{\sigma}, \theta, x^{\tau}])$ transports to a linearization $r' \in \mathcal{U}\text{-Plays}(\varphi[y^{\sigma}, y^{\tau}])$ satisfying $\partial(r) \cong_{!A \parallel !B \parallel C} \partial(r')$, so that $\hat{(\cdot)}(\partial(r')) = \bar{u}$.

Now by Proposition 6.2.7, $\varphi[y^{\tau} \otimes y^{\sigma}]$ is order-isomorphic to $y^{\tau} \otimes y^{\sigma} \in \mathcal{C}(\tau \otimes \sigma^{\dagger})$. Accordingly, r' transports to $r'' \in \mathcal{U}\text{-Plays}(y^{\tau} \otimes y^{\sigma})$ such that $\hat{(\cdot)}(\partial_{\tau \otimes \sigma^{\dagger}}(r'')) = \bar{u}$, which projects to $z \in \mathcal{U}\text{-Plays}(\tau \otimes \sigma^{\dagger})$ such that $\hat{(\cdot)}(\partial_{\tau \otimes \sigma^{\dagger}}(z)) = \bar{s}$. But s is a thread so that the pointers of initial moves in A are forced; accordingly $\hat{(\cdot)}(\partial_{\Lambda(\tau \otimes \sigma^{\dagger})}(z)) = s$. \square

Note that this inclusion does not require well-bracketing. Altogether:

Corollary 9.3.6. *Consider $\sigma \in \rightarrow\text{-Strat}(A, B)$ and $\tau \in \rightarrow\text{-Strat}(B, C)$.*

Then, $\mathcal{U}\text{-Unf}(\tau \odot_! \sigma) = \mathcal{U}\text{-Unf}(\tau) \odot \mathcal{U}\text{-Unf}(\sigma)$.

Proof. For the first inclusion, we compute:

$$\begin{aligned}
 \mathcal{U}\text{-Unf}(\tau \odot_! \sigma) &= \mathcal{U}\text{-Unf}(\tau \odot \sigma^{\dagger}) \\
 &= \mathcal{U}\text{-Unf}_!(\Lambda(\tau \odot \sigma^{\dagger})^{\dagger}) \\
 &= \mathcal{U}\text{-Unf}_*(\Lambda(\tau \odot \sigma^{\dagger}))^{\dagger} \\
 &\subseteq (\mathcal{U}\text{-Unf}(\tau) \odot \mathcal{U}\text{-Unf}(\sigma))^{\dagger} \\
 &= \mathcal{U}\text{-Unf}(\tau) \odot \mathcal{U}\text{-Unf}(\sigma)
 \end{aligned}$$

using definition of Kleisli composition, definition of $\mathcal{U}\text{-Unf}$, Lemma 9.2.16, Lemma 9.3.4. Finally, we use that by Lemma 9.2.16, $\mathcal{U}\text{-Unf}(\sigma)$ and $\mathcal{U}\text{-Unf}(\tau)$ are single-threaded strategies so that their composition is, so it is stable under interleavings.

The other inclusion holds on threads by Lemma 9.3.5 and extends to all plays by Lemma 9.2.16 since the non-alternating strategies involved are single-threaded. \square

Now, we may finally state:

Theorem 9.3.7. *Unfolding yields a \sim -functor $\mathcal{U}\text{-Unf} : \rightarrow\text{-Strat} \rightarrow \mathcal{U}\text{-Strat}$.*

Proof. For any $\sigma \in \rightarrow\text{-Strat}(A, B)$, $\mathcal{U}\text{-Unf}(\sigma)$ is a non-alternating strategy on $\underline{A} \Rightarrow \underline{B}$, single-threaded by Lemma 9.2.16, thus $\mathcal{U}\text{-Unf}(\sigma) \in \mathcal{U}\text{-Strat}(A, B)$. The only thing remaining to check is that $\mathcal{U}\text{-Unf}$ preserves equivalence, *i.e.*

$$\mathcal{U}\text{-Unf}(\sigma) = \mathcal{U}\text{-Unf}(\tau)$$

for all positive iso $\varphi : \sigma \approx \tau \in \rightarrow\text{-Strat}(A, B)$. To prove that, consider $s \in \mathcal{U}\text{-Unf}(\sigma)$, *i.e.* there must be $s^{\sigma} \in \mathcal{L}(\sigma)$ such that $\hat{(\cdot)}(\partial_{\Lambda(\sigma)}(s^{\sigma})) = s$. But then

$$\varphi(s^{\sigma}) \in \mathcal{L}(\tau)$$

and from the commutation in the game up to symmetry, it follows immediately that

$$\partial_{\Lambda(\sigma)^!}(s^\sigma) \cong_{!(A \multimap B)} \partial_{\Lambda(\tau)^!}(\varphi(s^\sigma))$$

which entails that their pointfixion is the same by Proposition 9.2.9. Thus we have found $\varphi(s^\sigma) \in \mathcal{L}(\tau)$ such that $\hat{\cdot}(\partial_{\Lambda(\tau)^!}(\varphi(s^\sigma))) = s$, so that $s \in \mathcal{U}\text{-Unf}(\tau)$. \square

9.3.2 Unfolding the Interpretation of \mathbf{IA}_{\parallel}

Next, we show that $\mathcal{U}\text{-Unf}$ preserves the interpretation of \mathbf{IA}_{\parallel} .

Cartesian closed \sim -functor. Recall that on objects, $\mathcal{U}\text{-Unf}(A, \underline{A}, \text{lbl}_A) = \underline{A}$ is a simple forgetful operation. By definition of constructions on mixed boards it follows that $\mathcal{U}\text{-Unf}(-)$ preserves on the nose the interpretation of ground types, products and arrows, and the terminal object. In order to preserve the cartesian closed structure, from the universal property it suffices to preserve projections and the evaluation morphism. For projections, a direct variation of Proposition 9.3.1 shows that

$$\mathcal{U}\text{-Unf}(\pi_A^!) = \pi_{\underline{A}}, \quad \mathcal{U}\text{-Unf}(\pi_B^!) = \pi_{\underline{B}},$$

so that the cartesian structure is preserved in the strict sense. For the evaluation,

$$\mathcal{U}\text{-Unf}(\Lambda^!(\sigma)) = \Lambda(\mathcal{U}\text{-Unf}(\sigma))$$

follows from a direct verification for $\sigma : !(\Gamma \& A) \vdash B$ with Γ, A and B mixed boards: currying is defined on both sides with essentially the same relabeling. It follows that $\mathcal{U}\text{-Unf}$ preserves Λ^{-1} as well, and thus the evaluation (which is $\Lambda^{-1}(\text{id}_{A \Rightarrow B})$).

Overall, we get:

Theorem 9.3.8. *We have a cartesian closed \sim -functor:*

$$\mathcal{U}\text{-Unf}(-) : \multimap\text{-}\mathbf{Strat} \rightarrow \mathcal{U}\text{-}\mathbf{Strat}$$

It is also straightforward that unfolding is continuous with respect to the depo structure of hom-sets, so that $\mathcal{U}\text{-Unf}(\mathcal{Y}_A) = \mathcal{Y}_{\underline{A}}$ for any mixed board A .

Unfolding \mathbf{PIA}_{\parallel} primitives. For all the basic primitives of \mathbf{PIA}_{\parallel} (constants, conditional, sequential composition, operation, sequential and parallel let bindings, reference and semaphore queries, reference and semaphore creation, and bad variable and semaphores), it follows by inspection that they are preserved by unfolding.

Altogether, we get:

Theorem 9.3.9. *Consider $\Gamma \vdash M : A$ any term of \mathbf{IA}_{\parallel} . Then,*

$$\mathcal{U}\text{-Unf}(\llbracket M \rrbracket_{\multimap\text{-}\mathbf{Strat}}) = \llbracket M \rrbracket_{\mathcal{U}\text{-}\mathbf{Strat}}.$$

Along with intentional full abstraction of $\multimap\text{-}\mathbf{Strat}$:

Theorem 9.3.10. *The interpretation of $\mathbf{IA}_{//}$ in $\rightarrow\text{-Strat}$ is intensionally fully abstract.*

Proof. Consider $\vdash M, N : A$ observationally equivalent and assume $\llbracket M \rrbracket_{\rightarrow\text{-Strat}}$ and $\llbracket N \rrbracket_{\rightarrow\text{-Strat}}$ are not, i.e. there is $\alpha \in \rightarrow\text{-Strat}(\llbracket A \rrbracket, \mathbf{U})$ which distinguishes them. Then,

$$\mathcal{O}\text{-Unf}(\alpha) \in \mathcal{O}\text{-Strat}(\llbracket A \rrbracket, \mathbf{U})$$

which, by unfolding, distinguishes $\llbracket M \rrbracket_{\mathcal{O}\text{-Strat}}$ and $\llbracket N \rrbracket_{\mathcal{O}\text{-Strat}}$. But since $\mathcal{O}\text{-Strat}$ is intensionally fully abstract for $\mathbf{IA}_{//}$, this contradicts the obs. equivalence $M \simeq N$. \square

In fact, from Theorems 5.3.2 and 9.3.9 we may immediately deduce that

$$M \simeq N \quad \Leftrightarrow \quad \text{comp}(\mathcal{O}\text{-Unf}(\llbracket M \rrbracket_{\rightarrow\text{-Strat}})) = \text{comp}(\mathcal{O}\text{-Unf}(\llbracket N \rrbracket_{\rightarrow\text{-Strat}})).$$

9.4 History and Related Work

This interpretation of higher-order concurrency into concurrent games (and its non-alternating unfolding) was first introduced in [Castellan and Clairambault, 2016], although for an affine language. The interpretation of non-affine $\mathbf{IA}_{//}$ in full concurrent games was first sketched in [Castellan et al., 2019], though without any adequacy or full abstraction result. The overall results presented here, with the unfolding to $\mathcal{O}\text{-Strat}$, were first presented in [Castellan and Clairambault, 2021].

Overall, this forms the first “truly concurrent” model of a higher-order concurrent language with shared memory. Around the same time this was developed, a truly concurrent model of the asynchronous π -calculus was proposed [Sakayori and Tsukada, 2017].

Chapter 10

Parallel Innocence

As we have seen, $\rightarrow\text{-Strat}$ (and already **NTCG**) is a wide semantic realm, with strategies witnessing a wealth of computational phenomena: they reflect higher-order concurrent programs, accessing shared state or semaphores. In this permissive universe, can we capture which causal patterns are definable via programs without state, *i.e.* with only pure parallel higher-order programming? Answering that is the purpose of this chapter, under the form of a condition on strategies called *parallel innocence*.

The outline is as follows. First, in Section 10.1 we introduce the definition of parallel innocence, motivating the different ingredients in sequence – in particular we introduce *visibility*, a first approximation of parallel innocence that already enjoys a number of its properties. Sections 10.2 and 10.3 are devoted to the stability under composition of parallel innocence: first visibility, and then innocence. In Section 10.4 we show that visible strategies admit an interpretation-preserving collapse to the relational model. Finally, in Section 10.5 we refine parallel innocence into *globularity*, that further constrains the causal shape of strategies accounting for Questions and Answers. This is less fundamental, but will be helpful later on for proving finite definability.

Remark. In this chapter, the interpretation is used mostly to provide examples. For succinctness, we shall refer in examples to the interpretation using *small* ground boards rather than *large* ground boards (see the beginning of Section 9.1.1). All formal statements on the interpretation hold for both, based on either small or large ground boards.

Besides, though we eventually must consider parallel innocence on the cartesian closed \sim -category $\rightarrow\text{-Strat}$, this development is completely independent from mixed boards and (logical) well-bracketing – thus, we prefer to develop the theory of parallel innocence for the linear structure, *i.e.* based on the relative Seely \sim -category **NTCG**. Of course, all the constructions given here will seamlessly apply to $\rightarrow\text{-Strat}$ as well.

10.1 Defining Parallel Innocence

We now embark on our quest to understand the causal shape of pure parallel programs.

10.1.1 Causal Determinism

In PCF_{\parallel} , the *parallel let* lets us write programs like

$$\Gamma \vdash \text{let } \begin{pmatrix} x & = & f \text{ skip} \\ y & = & g \text{ skip} \end{pmatrix} \text{ in } x + y : \mathbb{Y}$$

where the execution of $f \text{ skip}$ and $g \text{ skip}$ are in parallel. The *execution* of this program is non-deterministic: the scheduler must run one of $f \text{ skip}$ and $g \text{ skip}$ first. Nevertheless, one may argue that in PCF_{\parallel} , this order does not matter. Intuitively, we have

$$f \text{ skip} \cdot g \text{ skip} \sim g \text{ skip} \cdot f \text{ skip}$$

if $f \text{ skip}$ and $g \text{ skip}$ perform no side effect, in the sense that the environment has no way to distinguish between the two: whatever the order, the result of the computation will be the same. We say that this program is *causally deterministic*. This is unlike the behaviour of a true non-deterministic primitive such as **choice**.

Of course, IA_{\parallel} is a truly non-deterministic language (we have seen in Section 2.2.3 how to define **choice** via a race in the memory). However, this true non-determinism emerges from the combination of shared memory and parallelism. Hence, if *parallel innocence* is to remove shared memory, it must also reinstate causal determinism¹.

Definition. While *causal determinism* may be hard to express in traditional, play-based game semantics, it is straightforward to do so for causal strategies: it simply amounts to asking that Player does not impose any minimal conflicts.

Definition 10.1.1. Consider $\sigma : A \vdash B$ a causal strategy.

We say σ is *causally deterministic* if for all $s \sim_{\sigma} s'$ in σ , $\text{pol}(s) = \text{pol}(s') = -$.

This rejects a truly non-deterministic strategy such as in Figure 10.1, but accepts strategies such as that in Figure 10.2 where the conflict originates in Opponent moves. Keep in mind that by Lemma 6.1.17, such an immediate conflict between negative events in the strategy *must* correspond to an immediate conflict in the game.

Copycat strategies. We show that copycat strategies are deterministic. First:

Proposition 10.1.2. For any $-$ -board A , $\mathfrak{c}_A \in \text{NTCG}(A, A)$ is causally deterministic.

Proof. Consider $m_1 \sim_{\mathfrak{c}_A} m_2$ in minimal conflict. In particular, $m_1 \#_{\mathfrak{c}_A} m_2$ meaning that by definition, there is $m'_1 \leq_{\mathfrak{c}_A} m_1$ and $m'_2 \leq_{\mathfrak{c}_A} m_2$ such that

$$m'_1 \#_{A \vdash A} m'_2$$

¹The status of determinism in concurrent games is unsatisfactory – see Section 14.2.2 for a discussion.




Figure 10.1: A non-deterministic strategy




Figure 10.2: Deterministic copycat

hence there are $m_1'' \leq_{A \vdash A} m_1'$ and $m_2'' \leq_{A \vdash A} m_2'$ such that $m_1'' \sim_{A \vdash A} m_2''$. Now, by Lemma 6.1.8, $m_1'' \leq_{\mathfrak{C}_A} m_1'$ and $m_2'' \leq_{\mathfrak{C}_A} m_2'$, so that $m_1'' \leq_{\mathfrak{C}_A} m_1$ and $m_2'' \leq_{\mathfrak{C}_A} m_2$; and clearly $m_1'' \#_{\mathfrak{C}_A} m_2''$ as they conflict in the game. Thus by minimality, $m_1'' = m_1$ and $m_2'' = m_2$; hence we have shown that for any $m_1 \sim_{\mathfrak{C}_A} m_2$, $m_1 \sim_{A \vdash A} m_2$ as well.

Now assume, seeking a contradiction, that $\text{pol}_{A \vdash A}(m_1) = +$. By condition *race-free* of Definition 8.2.3, it follows that $\text{pol}_{A \vdash A}(m_2) = +$ as well. But then there are unique

$$m_1' \rightarrow_{\mathfrak{C}_A} m_1, \quad m_2' \rightarrow_{\mathfrak{C}_A} m_2$$

both negative and corresponding to the same event in A . So, $m_1' \sim_{A \vdash A} m_2'$, contradicting the minimality of $m_1 \sim_{\mathfrak{C}_A} m_2$. Hence, $\text{pol}_{A \vdash A}(m_1) = \text{pol}_{A \vdash A}(m_2) = -$. \square

For the first time, we have used the condition *race-free* from the definition of boards. Without this condition, a minimal conflict in the game of mixed polarity, as in

$$\oplus \quad \sim \quad \ominus,$$

causes copycat to be non-deterministic (as the reader may check).

It is straightforward that this applies to all structural morphisms involved in the categorical structure of **NTCG**, obtained via lifting of renamings – as the renaming operation $g \cdot \sigma \cdot f$ from Section 8.1.2 does not change the strategy, only its display map.

Composition. Next, we show preservation under composition. First, we must be able to attribute the responsibility of a minimal conflict to one of the two players.

Lemma 10.1.3. *Consider $\sigma : A \vdash B$ and $\tau : B \vdash C$ causal strategies.*

Then, for each $p \sim_{\tau \otimes \sigma} p'$, one of the two propositions hold:

- (1) p_σ, p'_σ are defined and $p_\sigma \sim_\sigma p'_\sigma$,
- (2) p_τ, p'_τ are defined and $p_\tau \sim_\tau p'_\tau$.

Proof. An equivalent statement of the hypothesis is that we have

$$x^\tau \otimes x^\sigma \xrightarrow{p} y^\tau \otimes y^\sigma, \quad x^\tau \otimes x^\sigma \xrightarrow{p'} z^\tau \otimes z^\sigma$$

with $y^\tau \otimes y^\sigma, z^\tau \otimes z^\sigma$ incompatible. By Lemma 6.2.12, this means that we have

$$\varphi[x^\sigma, x^\tau] \stackrel{(m_1, m_2)}{\dashv\vdash} \varphi[y^\sigma, y^\tau], \quad \varphi[x^\sigma, x^\tau] \stackrel{(m'_1, m'_2)}{\dashv\vdash} \varphi[z^\sigma, z^\tau]$$

with $\varphi[y^\sigma, y^\tau] \cup \varphi[z^\sigma, z^\tau]$ not a secured bijection. This may be because it is not a bijection – say $m_1 = m'_1$ while $m_2 \neq m'_2$. But then $m_2 = (2, t)$ and $m'_2 = (2, t')$ with $\partial_\tau(t) = \partial_\tau(t')$ and thus $t \#_\tau t'$ by local injectivity. Overall, $p_\tau = t$ and $p'_\tau = t'$ are defined and conflicting. Or, this may be because $m_1 \#_{\sigma \parallel C} m'_1$ or $m_2 \#_{A \parallel \tau} m'_2$, say the former. If $m_1 = (1, s)$ and $m'_1 = (1, s')$, then $p_\sigma = s$ and $p'_\sigma = s'$ are defined and conflicting. Otherwise, $m_1 = (2, c)$ and $m_2 = (2, c')$ with $c \#_C c'$. But then, necessarily $m_2 = (2, t)$ and $m'_2 = (2, t')$ with $t \#_\tau t'$, thus $p_\tau = t$ and $p'_\tau = t'$ defined and conflicting.

It remains to prove that the resulting conflict is minimal, but that is obvious: if $p_\sigma \#_\sigma p'_\sigma$, all their strict dependency lie in x^σ – and symmetrically if $p_\tau \#_\tau p'_\tau$. \square

It remains to link minimal conflicts in the composition with the interaction:

Lemma 10.1.4. *Consider $\sigma : A \vdash B$ and $\tau : B \vdash C$ causal strategies.*

If $p \sim_{\tau \odot \sigma} p'$, then we have a diagram in $\tau \otimes \sigma$:

$$\begin{array}{c} q_1 \longrightarrow q_2 \longrightarrow \dots \longrightarrow q_m \longrightarrow p \\ \left. \begin{array}{c} \vdots \\ \vdots \end{array} \right\} \\ q'_1 \longrightarrow q'_2 \longrightarrow \dots \longrightarrow q'_n \longrightarrow p' \end{array}$$

with all q_i 's synchronized (if $m > 0$) and q'_i 's synchronized (if $n > 0$).

Proof. Obvious by definition of $\tau \odot \sigma$. \square

Proposition 10.1.5. *Consider $\sigma : A \vdash B$ and $\tau : B \vdash C$ causally deterministic.*

Then, $\tau \odot \sigma : A \vdash C$ is causally deterministic.

Proof. Consider $p \sim_{\tau \odot \sigma} p'$ a minimal conflict in $\tau \odot \sigma$, and its originating immediate conflict in $\tau \otimes \sigma$ as guaranteed by Lemma 10.1.4. First, assume that $p \sim_{\tau \otimes \sigma} p'$, i.e. $m = n = 0$. By Lemma 10.1.3, this corresponds to an immediate conflict for σ or τ – assume the former, so that p_σ and p'_σ are defined and $p_\sigma \sim_\sigma p'_\sigma$. But then since σ is causally deterministic, those are negative moves, hence negative in $\tau \odot \sigma$ as well.

Otherwise, we have $m > 0$ or $n > 0$, say the former so that the immediate conflict originates in $q \sim_{\tau \otimes \sigma} q'$ with q synchronized. By Lemma 10.1.3, this corresponds to an immediate conflict for σ or τ – assume the former, so that q_σ and q'_σ are defined and $q_\sigma \sim_\sigma q'_\sigma$. Since σ is causally deterministic, q_σ and q'_σ must be negative and by Lemma 6.1.17, their display to the game are conflicting. But this implies that q_τ and q'_τ are defined and conflicting as well. Additionally we must have $q_\tau \sim_\tau q'_\tau$, or we would get a contradiction with the minimality of $q \sim_{\tau \otimes \sigma} q'$. But since τ is causally deterministic this means that q_τ and q'_τ are negative, contradiction. \square




Figure 10.3: A causal strategy





Figure 10.4: Two configurations of Figure 10.3

Full structure. For remaining operations on strategies (tensor, pairing, currying, promotion), it is immediate by construction that they preserve causal determinism. So:

Theorem 10.1.6. *The relative Seely \sim -category **NTCG** admits a lluf relative Seely sub- \sim -category **NTCG-Det** with morphisms restricted to causally deterministic strategies.*

Additionally, the least upper bound of a chain of causally deterministic strategies is still causally deterministic, hence the fixpoint combinator is causally deterministic. It also follows by direct inspection that the strategies used in the interpretation of PCF_{\parallel} are all causally deterministic; hence for any $\Gamma \vdash M : A$ in PCF_{\parallel} , it follows that

$$\llbracket M \rrbracket \in \text{NTCG-Det}_!(\llbracket \Gamma \rrbracket, \llbracket A \rrbracket).$$

But causal determinism is not sufficient to characterise pure parallel computation.

10.1.2 Pre-innocence

So, what causal shapes are distinctive of pure parallel computation?

Pure parallel programs may spawn parallel threads, which remain independent in the absence of interference. Once they 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 10.6. The slogan is:

“Player may merge threads than he himself has spawned”.

In contrast, both diagrams of Figure 10.4 bear signs of interference. In the first, the answer 1^+ depends on 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:

“Player may not merge threads spawned by Opponent”.




Figure 10.5: Maximal grounded causal chains of a causal strategy

To define parallel innocence, our first step is to introduce a formal notion of “thread”:

Definition 10.1.7. Consider A a board, 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, written $\rho = \rho_1 \rightarrow_{\sigma} \dots \rightarrow_{\sigma} \rho_n \in \text{gcc}(\sigma)$ with the causal ordering 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 10.5 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 “only merging threads forked by Player”.

Definition 10.1.8. For A a board, a causally deterministic $\sigma : A$ is **pre-innocent** iff

pre-innocent: If $m^+ \in \sigma$ and $\rho_1 \rightarrow m, \rho_2 \rightarrow m \in \text{gcc}(\sigma)$ are distinct, then $\min(\rho_1) = \min(\rho_2)$ and their least distinct moves are positive.

The strategy of Figure 10.6 is pre-innocent. In contrast, that of Figure 10.3 is not – both configurations of Figure 10.4 fail pre-innocence. For instance, the second and third gccs of Figure 10.5 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.

It will follow later on that the *sequential pre-innocent* causal strategies exactly match the standard alternating innocent strategies of Definition 3.2.9; see Section 11.3.

However, for non-sequential strategies, pre-innocence is still incomplete.

10.1.3 Visibility

The problem arises as non-stability of pre-innocence under composition.

$$!(U \ \& \ U) \vdash U$$




Figure 10.6: A typical pre-innocent strat.

$$!(U \ \& \ U) \parallel U$$





Figure 10.7: Partiality of views

$$(U \multimap U) \multimap N$$



∈

$$\text{gcc} \left((U \multimap U) \multimap N \right)$$

Figure 10.8: A gcc of a non-visible strategy, losing its pointer

A counter-example appears in Figure 10.9, which we postpone to the section devoted to compositionality of innocence. However, we can already 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 10.8 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 10.1.9. A causal strategy $\sigma : A$ is **visible** if it is:

- pointed: for any $s \in \sigma$, there is a unique $\text{init}(s) \leq_\sigma s$ minimal in σ ,
- valid-gccs: for all $\rho \in \text{gcc}(\sigma)$, $\partial_\sigma(\rho) \in \mathcal{C}(A)$.

We smuggle in *pointed*, which ensures that separate initial Opponent moves explore causally independent parts of the strategy. It is necessary for our proof of composition of visibility, but it is clearly less important than *valid-gccs*, mainly because it is already implied by pre-innocence (and also, automatic for $\sigma : A \vdash B$ with B strict). Hence, we shall tend in the future to brush over *pointed*, and identify *visible* and *valid-gccs*.

Valid-gccs means that every move in ρ points within ρ . This phrasing highlights the analogy with Definition 3.2.8, *i.e.* “Player always points in the P-view”. It is indeed this

analogy that inspired the name². 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 10.3) may not be visible. Visibility is very restrictive, it is not clear what sensible primitive would satisfy visibility but not pre-innocence.

The following lemma captures how a gcc may be regarded as a standalone thread.

Lemma 10.1.10. *Consider A a mixed board, and $\sigma : A$ a visible causal strategy.*

If $\rho = \rho_1 \rightarrow_{\sigma} \dots \rightarrow_{\sigma} \rho_n \in \text{gcc}(\sigma)$, $\hat{(\cdot)}(\partial_{\sigma}(\rho)) = \hat{(\cdot)}(\partial_{\sigma}(\rho_1) \dots \partial_{\sigma}(\rho_n))$ is a P-view.

Proof. By Lemma 6.1.16, $\partial_{\sigma}(\rho)$ is alternating. By visibility, its prefixes are configurations of A . So, $\partial_{\sigma}(\rho) \in \Downarrow\text{-Plays}(A)$. By Lemma 6.1.16 again, 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. The statement follows by definition of pointifixion. \square

We may now define *parallel innocent* causal strategies, or just *innocent* for short.

Definition 10.1.11. *Consider $\sigma : A$ causally deterministic on board A .*

*It is **parallel innocent** if it is pre-innocent and **visible**.*

A standard innocent strategy as in Section 3.2.3, under its “causal” presentation, is a forest of P-views, 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.

10.2 Composition of Visibility

First, we establish compositionality of visibility.

10.2.1 Justifiers in causal strategies

We introduce some machinery on *justifiers*. If $\sigma : A$ is a causal strategy on A some $--$ board, then as for plays, the immediate causality in A endows moves in σ with a notion of *justifier*. This extends to $\sigma : A \vdash B$ where A and B are $--$ boards:

Definition 10.2.1. *Consider A and B $--$ boards, and $\sigma : A \vdash B$ pointed.*

*We define the **justifiers** in σ by setting, for all $m, m' \in \sigma$,*

$$\begin{aligned} j(m) &= m' && \text{if } \partial_{\sigma}(m') \rightarrow_{A \vdash B} \partial_{\sigma}(m), \\ j(m) &= \text{init}(m) && \text{if } \partial_{\sigma}(m) \text{ minimal in } A, \end{aligned}$$

and undefined otherwise.

²This, plus as in traditional game semantics, visibility is a prerequisite for a working notion of innocence.

Since A and B are *forestial* and σ is pointed, $j(-)$ is well-defined as a partial function on σ , only left undefined for $m \in \sigma$ minimal in σ . 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 pointed causal strategy $\sigma : A \vdash B$ is visible iff for all $\rho \in \text{gcc}(\sigma)$, for all $m \in \rho$, $j(m) \in \rho$ as well: all gccs are closed under justifiers. We mention in passing this lemma:

Lemma 10.2.2. *Consider A, B --boards and $\sigma : A \vdash B$ a pointed causal strategy.*

Then, for any non-initial $m \in \sigma$, we have $j(m) <_{\sigma} m$. Moreover, if $\text{pol}_{\sigma}(m) = -$, then $j(m) \rightarrow_{\sigma} m$ is its (unique) immediate predecessor in σ .

Proof. As a map of es, ∂_{σ} locally reflects causality (Lemma 6.1.8), so $j(m) <_{\sigma} m$ if the first clause of Definition 10.2.1 applies; for the other we clearly have $\text{init}(m) <_{\sigma} m$.

The second fact is simply by Lemma 6.1.16. \square

10.2.2 Justifiers in interactions

So as to prove visibility stable under composition, we first extend justifiers to *interactions* – consider A, B and C three --boards, and $\sigma : A \vdash B$ and $\tau : B \vdash C$ pointed.

Definition 10.2.3. *We define $j : |\tau \otimes \sigma| \rightarrow |\tau \otimes \sigma|$ as $j(m) = m'$ if:*

- (1) $\partial_{\tau \otimes \sigma}(m') \rightarrow_{A \parallel B \parallel C} \partial_{\tau \otimes \sigma}(m)$, or
- (2) $\partial_{\tau \otimes \sigma}(m)$ is minimal in A and $m'_{\sigma} = \text{init}(m_{\sigma})$, or
- (3) $\partial_{\tau \otimes \sigma}(m)$ is minimal in B and $m'_{\tau} = \text{init}(m_{\tau})$,

*and undefined otherwise. We say that m' is the **justifier** of m in $\tau \otimes \sigma$.*

This leaves $j(m)$ 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 $m \in \tau \otimes \sigma$, if m_{σ} is defined then $j(m)_{\sigma}$ is defined and equal to $j(m_{\sigma})$, and likewise for τ .

10.2.3 Views of gccs

We introduce the main technical device on visible causal interactions.

We use polarities in interactions as in Lemma 6.2.15, and annotate events accordingly. We also write e.g. $a^{-, \prime}$ to indicate that a has polarity $-$ or \prime . If $\rho \in \text{gcc}(\tau \otimes \sigma)$ 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 10.2.4. *Take $\sigma : A \vdash B$ and $\tau : B \vdash C$ pointed with A, B, C --boards.*

If $\rho \in \text{gcc}(\tau \otimes \sigma)$ ends in σ , we (partially) define $\ulcorner \rho \urcorner^{\sigma} \in \text{gcc}(\sigma)$ by:

$$\begin{aligned} \ulcorner \rho_0 \rightarrow \dots \rightarrow \rho_n \rightarrow \rho_{n+1}^{\ell} \urcorner^{\sigma} &= \ulcorner \rho_0 \rightarrow \dots \rightarrow \rho_n \urcorner^{\sigma} \cup \{(\rho_{n+1})_{\sigma}\}, \\ \ulcorner \rho_0 \rightarrow \dots \rightarrow \rho_i \rightarrow \dots \rightarrow \rho_{n+1}^{\prime} \urcorner^{\sigma} &= \ulcorner \rho_0 \rightarrow \dots \rightarrow \rho_i \urcorner^{\sigma} \cup \{(\rho_{n+1})_{\sigma}\} && \text{if } j(\rho_{n+1}) = \rho_i \text{ in } A \text{ or } B, \\ \ulcorner \rho_0 \rightarrow \dots \rightarrow \rho_i \rightarrow \dots \rightarrow \rho_{n+1}^{\prime} \urcorner^{\sigma} &= \{(\rho_{n+1})_{\sigma}\} && \text{if } \rho_{n+1} \text{ minimal in } B, \end{aligned}$$

undefined otherwise. If $\rho \in \text{gcc}(\tau \otimes \sigma)$ ends in τ , we (partially) define $\ulcorner \rho \urcorner^{\tau} \in \text{gcc}(\tau)$:

$$\begin{aligned} \ulcorner \rho_0 \rightarrow \dots \rightarrow \rho_n \rightarrow \rho_{n+1}^{\prime} \urcorner^{\tau} &= \ulcorner \rho_0 \rightarrow \dots \rightarrow \rho_n \urcorner^{\tau} \cup \{(\rho_{n+1})_{\tau}\}, \\ \ulcorner \rho_0 \rightarrow \dots \rightarrow \rho_i \rightarrow \dots \rightarrow \rho_{n+1}^{\ell} \urcorner^{\tau} &= \ulcorner \rho_0 \rightarrow \dots \rightarrow \rho_i \urcorner^{\tau} \cup \{(\rho_{n+1})_{\tau}\} && \text{if } j(\rho_{n+1}) = \rho_i; \end{aligned}$$

when defined, $\ulcorner \rho \urcorner^\sigma \in \text{gcc}(\sigma)$ is the σ -view of ρ and $\ulcorner \rho \urcorner^\tau \in \text{gcc}(\tau)$ the τ -view of ρ .

These definitions closely follow Definition 3.2.7. 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 6.2.15, 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 10.7, is not well-defined: when attempting to compute $\ulcorner \mathbf{q}^- \mathbf{q}_0^+ \urcorner^\tau$, none of the clauses apply as $j(\mathbf{q}_0^+) = \mathbf{q}_1^+$ is outside ρ . The bulk of the proof of stability of visibility under composition, is to show that this cannot happen:

Proposition 10.2.5. *Let $\sigma : A \vdash B$ and $\tau : B \vdash C$ be visible causal strategies.*

Then, the views of gccs of $\tau \otimes \sigma$ as in Definition 10.2.4 are always well-defined.

Proof. We prove by induction on ρ that, for all prefixes of ρ ,

- (1) if ρ ends in σ , then $\ulcorner \rho \urcorner^\sigma \in \text{gcc}(\sigma)$ is well-defined,
- (2) if ρ ends in τ , then $\ulcorner \rho \urcorner^\tau \in \text{gcc}(\tau)$ is well-defined.

Assume ρ ends in τ . If the last move has polarity $-$, then either it is initial and there is nothing to prove, or by Lemma 6.2.15 its justifier is its predecessor in ρ , so $\ulcorner \rho \urcorner^\tau \in \text{gcc}(\tau)$ follows by induction hypothesis (in that case ρ does not end in σ).

If the last move has pol. $+$, write $\rho = \rho' \rightarrow m \rightarrow n^+$. By Lemma 6.2.15, $m_\tau \rightarrow_\tau n_\tau$, so m is in τ . By induction hypothesis, $\kappa = \ulcorner \rho' \urcorner^\tau \rightarrow m^\tau \in \text{gcc}(\tau)$, so

$$\kappa \rightarrow n_\tau = \ulcorner \rho \urcorner^\tau \in \text{gcc}(\tau)$$

as well. But if ρ ends in σ and τ (i.e. in B), we must further prove that $\ulcorner \rho \urcorner^\sigma \in \text{gcc}(\sigma)$. In that case, since $\ulcorner \rho \urcorner^\tau \in \text{gcc}(\tau)$ and τ is visible, it follows that $j(n_\tau) \in \ulcorner \rho \urcorner^\tau$, but this entails $j(n) \in \rho$. Hence the second clause of Definition 10.2.4 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 10.2.6. *Let $\sigma : A \vdash B$ and $\tau : B \vdash C$ be visible causal strategies.*

Then, $\tau \odot \sigma : A \vdash C$ is also visible.

Proof. First, we show that $\tau \odot \sigma$ is pointed. there are distinct $q, q' \in \tau \odot \sigma$ minimal such that $q \leq_{\tau \odot \sigma} p$ and $q' \leq_{\tau \odot \sigma} p$. It is a direct consequence of courtesy (via Lemma 6.1.16) that then, q, q' must still be minimal in $\tau \otimes \sigma$. Then there must be some r in $\tau \otimes \sigma$ minimal such that there is $m \rightarrow_{\tau \otimes \sigma} r$ and $n \rightarrow_{\tau \otimes \sigma} r$ with $[m]_{\tau \otimes \sigma}$ and $[n]_{\tau \otimes \sigma}$ disjoint. Now, we distinguish cases depending on the polarity of r . If it is ℓ , then by Lemma 6.2.15, m_σ and n_σ are defined and $m_\sigma \rightarrow_\sigma r_\sigma$ and $n_\sigma \rightarrow_\sigma r_\sigma$. But since σ is pointed, $[m_\sigma]_\sigma$ and $[n_\sigma]_\sigma$ cannot be disjoint, contradiction. The case for $+$ is symmetric.

Now, we show visibility. For that, we prove by induction on ρ that for all $\rho \in \text{gcc}(\tau \otimes \sigma)$, $\partial_{\tau \otimes \sigma}(\rho) \in \mathcal{C}(A \parallel B \parallel C)$. If ρ is empty it is clear; take $\rho \rightarrow m \in \text{gcc}(\tau \otimes \sigma)$. By induction hypothesis, $\partial_{\tau \otimes \sigma}(\rho) \in \mathcal{C}(A \parallel B \parallel C)$, we only need that the justifier of m is in ρ . We reason by cases on the polarity of m : if it is ℓ , then by Proposition 10.2.5 $\ulcorner \rho \urcorner^\sigma \rightarrow m^\sigma \in \text{gcc}(\sigma)$. But since σ is visible, the justifier of m_σ appears in $\ulcorner \rho \urcorner^\sigma$, hence




Figure 10.9: Failure of preservation of pre-innocence under composition

the justifier of m appears in 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) $\rho_{\otimes} \in \text{gcc}(\tau \otimes \sigma)$ s.t. ρ_{\odot} comprises exactly those events of ρ_{\otimes} in A or C . By the observation above, $\partial_{\tau \otimes \sigma}(\rho_{\otimes}) \in \mathcal{C}(A \parallel B \parallel C)$, hence $\partial_{\tau \odot \sigma}(\rho_{\odot}) \in \mathcal{C}(A \parallel C)$. \square

10.2.4 The Relative Seely \sim -Category NTCG-Vis

In this monograph, we are particularly interested in visibility as a means to ensure that innocence is preserved under composition. However, it is noteworthy that visibility indeed supports all the constructions involved in the relative Seely category structure.

Theorem 10.2.7. *There is NTCG-Vis, a relative Seely sub- \sim -category of NTCG, with morphisms restricted to visible causal strategies.*

Proof. As causality in boards is forestial, so is the causality of copycat – it follows immediately that copycat is visible, and this immediately extends to lifted strategies. By Proposition 10.2.6, visibility is stable under composition. For all other operations (tensor, pairing, currying, promotions), it is immediate that visibility is preserved. \square

This extends to the fixpoint operation and – by inspection – to all primitive strategies involved in the interpretation of PCF_{\parallel} , so that any term $\Gamma \vdash M : A$ of PCF_{\parallel} yields a visible causal strategy $\llbracket M \rrbracket \in \text{NTCG-Vis}_i(\llbracket \Gamma \rrbracket, \llbracket A \rrbracket)$. In contrast, the interpretation of IA_{\parallel} does *not* satisfy visibility: the non-visible behaviour from Figure 10.8 comes from

$$\vdash \text{newref } x \text{ in } \lambda f^{\mathbb{U} \rightarrow \mathbb{U}}. f(x := \mathbf{tt}); !r : (\mathbb{U} \rightarrow \mathbb{U}) \rightarrow \mathbb{B}, \quad (10.1)$$

a term of IA . Though visibility does not match tightly any of our programming languages, it will play an important role in this work. But next, we show that in the presence of visibility, innocence becomes stable under composition.

10.3 Composition of Innocence

We start by showing “what could go wrong”. In Figure 10.9, we show a counter-

$$U \multimap U \multimap U \vdash (U \multimap U \multimap U) \multimap U$$





Figure 10.10: Illegal causal merge

Figure 10.11: Merging paths in G

example to the stability under composition of pre-innocence without visibility, with the corresponding interaction appearing as Figure 10.10. 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 justification relations from Definition 10.2.1 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 $U \multimap U \multimap U$. 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 $\tau : B \vdash C$.

10.3.1 The “forking lemma”

Taking a closer look at Figure 10.10, 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 $\rho = \rho_1 \multimap \dots \multimap \rho_n$ is a gcc and $1 \leq i \leq n$, $\rho_{\leq i}$ is $\rho_1 \multimap \dots \multimap \rho_i$. 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; but if they do diverge, their remainder must be independent.

Lemma 10.3.1 (Forking lemma). *Consider $\sigma : A \vdash B$ and $\tau : B \vdash C$ visible strategies. Let $\rho, \kappa \in \text{gcc}(\tau \otimes \sigma)$ be forking gccs ending in σ , s.t. $\ulcorner \rho \urcorner^\sigma \cap \ulcorner \kappa \urcorner^\sigma \neq \emptyset$ and $\text{gce}(\ulcorner \rho \urcorner^\sigma, \ulcorner \kappa \urcorner^\sigma)$ negative (the least distinct events, if any, are positive).*

Then, $\text{gce}(\ulcorner \rho \urcorner^\sigma, \ulcorner \kappa \urcorner^\sigma) = \text{gce}(\rho, \kappa)_\sigma$. 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:

$$\begin{aligned} O\text{-edges} &= \{(m_1, m_2^\ell) \mid j(m_1) = m_2\} \\ P\text{-edges} &= \{(m_1^\ell, m_2) \mid m_2 \rightarrow_{\tau \otimes \sigma} m_1\} \end{aligned}$$

where $(-)^{\ell}$ 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 $\ulcorner \rho \urcorner^\sigma \cap \ulcorner \kappa \urcorner^\sigma \neq \emptyset$, these paths must intersect – Figure 10.11 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 10.11), then $\text{gce}(\ulcorner \rho \urcorner^\sigma, \ulcorner \kappa \urcorner^\sigma)$ is positive, which contradicts the hypothesis. For the latter (as in the top of Figure 10.11), 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: $\text{gce}(\rho, \kappa)$. \square

This provides the core argument for the compositionality of pre-innocence: 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.


10.3.2 Stability of 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 10.3.2. *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 ensure this by changing m and ρ^i , keeping the same least distinct events. Likewise, since ρ^1, ρ^2 are distinct, *w.l.o.g.* 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 gccs 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 6.2.15 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.

The polarity of n is positive for σ , so we compute the σ -views:

$$\ulcorner \kappa_{\leq n}^1 \urcorner \rightarrow n^{\neg\sigma}, \quad \ulcorner \kappa_{\leq p}^2 \urcorner \rightarrow n^{\neg\sigma} \in \text{gcc}(\sigma),$$

respectively $\ulcorner \kappa_{\leq n}^1 \urcorner^{\neg\sigma} \rightarrow n_\sigma$ and $\ulcorner \kappa_{\leq p}^2 \urcorner^{\neg\sigma} \rightarrow n_\sigma$, with $\ulcorner \kappa_{\leq n}^1 \urcorner^{\neg\sigma}$ and $\ulcorner \kappa_{\leq p}^2 \urcorner^{\neg\sigma}$ distinct as they respectively contain κ_n^1 and κ_p^2 . Since σ is pointed, $\ulcorner \kappa_{\leq n}^1 \urcorner^{\neg\sigma} \cap \ulcorner \kappa_{\leq p}^2 \urcorner^{\neg\sigma} \neq \emptyset$, so $\kappa^1 \cap \kappa^2 \neq \emptyset$ as well. Call m' the gce 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. Thus the σ -views

$$\ulcorner \kappa_{\leq n}^1 \urcorner^{\neg\sigma}, \quad \ulcorner \kappa_{\leq p}^2 \urcorner^{\neg\sigma} \in \text{gcc}(\sigma)$$

are *forking*: they coincide on a prefix and disjoint after. Indeed, since ξ^1 and ξ^2 are disjoint, any common event is in $\kappa_i^1 \rightarrow \dots \rightarrow \kappa_i^1$, before which the σ -views coincide.

Since σ is pre-innocent, the least distinct moves m'_1 and m'_2 of $\ulcorner \kappa_{\leq n}^1 \urcorner^{\neg\sigma}$ and $\ulcorner \kappa_{\leq p}^2 \urcorner^{\neg\sigma}$ are positive. Thus their common immediate predecessor is negative – but it is also their greatest common event, since $\ulcorner \kappa_{\leq n}^1 \urcorner^{\neg\sigma}$ and $\ulcorner \kappa_{\leq p}^2 \urcorner^{\neg\sigma}$ are forking. Thus by Lemma 10.3.1,

$$\text{gce}(\ulcorner \kappa_{\leq n}^1 \urcorner^{\neg\sigma}, \ulcorner \kappa_{\leq p}^2 \urcorner^{\neg\sigma}) = \text{gce}(\kappa_{\leq n}^1, \kappa_{\leq p}^2)_\sigma = m'_\sigma,$$

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

10.3.3 The Relative Seely Category NTCG-Inn

As for visibility, we have:

Theorem 10.3.3. *There is NTCG-Inn, a relative Seely sub- \sim -category of NTCG-Det, with morphisms restricted to parallel innocent causal strategies.*

Proof. For composition, this is via Propositions 10.1.5, 10.2.6 and 10.3.2. Parallel innocence of copycat is immediate, and that extends to lifted strategies. For all other operations on strategies, preservation of innocence is direct. \square

Again this extends to the fixpoint operation and to all primitive strategies involved in the interpretation of $\text{PCF}_{//}$, so that any term $\Gamma \vdash M : A$ of $\text{PCF}_{//}$ yields

$$\llbracket M \rrbracket \in \mathbf{NTCG}\text{-}\mathbf{Inn}_!(\llbracket \Gamma \rrbracket, \llbracket A \rrbracket)$$

a parallel innocent strategy.

Note that though we bundled causal determinism with parallel innocence in this monograph, they are actually independent in the sense that stability under composition of parallel innocence does not need causal determinism. However, without determinism, parallel innocence as presented here is insufficient to obtain a definability result for a non-deterministic extension of $\text{PCF}_{//}$. We have not yet been able to give a working notion of non-deterministic parallel innocence – see Section 14.2.1 for a discussion.

10.3.4 Complement: the “Bang Lemma”

Within the relative Seely \sim -category $\mathbf{NTCG}\text{-}\mathbf{Inn}$, one can prove the “bang lemma” from AJM games [Abramsky et al., 2000] – it is an important technical lemma expressing that essentially, all innocent strategies on $!A \vdash !B$ are the promotion of some strategy on $!A \vdash B$. Note that this does not hold without innocence: a non-innocent strategy on $!B$ may very well impose constraints across its different copies. In AJM games, the bang lemma depends on *history-freeness*; it does not exist in HO games in the absence of a linear decomposition. Fix A and B two strict boards, and $\sigma : !A \vdash !B$ innocent.

A minimal $m \in \sigma$ displays to $\partial_\sigma(m) = (2, (i, b))$ for some $i \in \mathbb{N}$ – we say that m is **in copy i** . Being innocent, σ is *pointed* and every event has a unique minimal cause; so we say that arbitrary $m \in \sigma$ is **in copy i** if its unique minimal cause is. Let us write

$$|\sigma_i| = \{m \in \sigma \mid m \text{ is in copy } i\}$$

for the set of events of σ in copy i . By construction, for $i, j \in \mathbb{N}$ distinct, $|\sigma_i|$ and $|\sigma_j|$ are disjoint, and each inherit from σ the structure of an ess. Furthermore, $\sigma_i : !A \vdash B$ with the display map ∂_{σ_i} defined in the obvious way. Moreover, we have:

Lemma 10.3.4. *As event structures with symmetries, we have $\sigma \cong \coprod_{i \in \mathbb{N}} \sigma_i$.*

Proof. The non-trivial point is that no immediate conflict in σ can span across copies: if $s \sim_\sigma s'$ then $\text{pol}(s) = \text{pol}(s') = -$ by *causal determinism*, and $\partial_\sigma(s) \sim \partial_\sigma(s')$ as well by Lemma 6.1.17. By definition of bang, this entails that $\partial_\sigma(s)$ and $\partial_\sigma(s')$ are in a common copy of B , or of A . If it is B then we are done. If it is A then from polarity reasons they cannot be minimal in A , and since A is strict this entails that they are above (for $\leq_{!A}$) the *same* minimal event in A . So s and s' are above a common move in σ , which entails that they are in the same copy. \square

The key argument, in the bang lemma, is that these copies are interchangeable:

Lemma 10.3.5. *For any $i, j \in \mathbb{N}$, we have $\sigma_i \approx \sigma_j$.*

Proof. We exploit Lemma 7.1.9 and build an iso between the domains of configurations, compatible with symmetry. Consider $x_i \in \mathcal{C}(\sigma_i)$, displaying to

$$\partial_\sigma(x_i) = x_A \parallel \{i\} \times x_B.$$

By definition of the board constructions, there is a symmetry exchanging i and j :

$$\theta_{i,j}^- : x_A \parallel \{i\} \times x_B \cong_{!A \vdash !B}^- x_A \parallel \{j\} \times x_B,$$

and our isomorphism will transport x_i along this negative symmetry. Indeed by Lemma 7.2.6, there are unique $x_j \in \mathcal{C}(\sigma)$ and $\psi : x_i \cong_\sigma x_j$ s.t. $\partial_\sigma \psi = \theta_{i,j}^-$ 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 by definition of the positive symmetries of the bang. 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 7.2.6; 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 10.3.6. *Consider A and B strict boards, and $\sigma, \tau : !A \vdash !B$ parallel innocent. If $\sigma_0 \approx \tau_0$, then $\sigma \approx \tau$.*

Proof. By Lemma 10.3.5, for $i \in \mathbb{N}$, $\sigma_i \approx \sigma_0 \approx \tau_0 \approx \tau_i$, and thus we have

$$\parallel_{i \in \mathbb{N}} \sigma_i \approx \parallel_{i \in \mathbb{N}} \tau_i$$

lifted to an isomorphism of ess $\sigma \cong \tau$ by Lemma 10.3.4. It is then straightforward that this isomorphism is compatible with display maps up to positive symmetries. \square

We may finally deduce the bang lemma:

Lemma 10.3.7. *For A, B strict boards and $\sigma \in \text{NTCG-Inn}(!A, !B)$ parallel innocent,*

$$(\text{der}_B \odot \sigma)^! \approx \sigma.$$

Proof. From Proposition 7.3.1, it is direct to establish $\text{der}_B \odot \sigma \approx \sigma_0$. Therefore, by Lemma 10.3.6, for any $\sigma, \tau : !A \vdash !B$, if $\text{der}_B \odot \sigma \approx \text{der}_B \odot \tau$, then $\sigma \approx \tau$. Now as

$$\text{der}_B \odot (\text{der}_B \odot \sigma)^\dagger \approx \text{der}_B \odot \sigma$$

by the laws of relative Seely categories, the lemma follows. \square




Figure 10.12: A witness and its display

10.4 Relational Collapse

Most consequences of parallel innocence will be presented in Chapter 12, when we focus on finite definability for those and intensional full abstraction for $\text{PCF}_{//}$. But we explore now a crucial consequence of parallel innocence (more precisely, of visibility): it allows a simple, transparent translation to the relational model. This will be useful technically later on, but we also regard this as an important contribution in its own right.

10.4.1 Stopping Positions and Witnesses

Intuitively, a causal strategy records observable computational events, along with the causal dependencies between them, but we may also regard it as presenting all positions reached in the sense of the relational model, enriching them with causal information.

Forgetting causality. Recall that a causal strategy $\sigma : A$ on $--$ -board A is specified by an internal event structure σ , along with a *display map*

$$\partial_\sigma : \sigma \rightarrow A.$$

A configuration $x_A \in \mathcal{C}(A)$ is **reached by** σ iff there is $x^\sigma \in \mathcal{C}(\sigma)$, its **witness**, such that $\partial_\sigma x^\sigma = x_A$. Intuitively, and as a first approximation, the relational model only remembers which configurations $x_A \in \mathcal{C}(A)$ have a witness (though we shall see below that the relational model only tracks certain configurations considered as *completed* executions, and that *symmetry* entails that points in the relational models should be certain symmetry classes rather than mere configurations).

As an example, we show in Figure 10.12 a configuration of the strategy σ for

$$\vdash \lambda f^{\mathbb{U} \rightarrow \mathbb{U}}. f \text{ skip}; f \text{ skip}; tt : (\mathbb{U} \rightarrow \mathbb{U}) \rightarrow \mathbb{B};$$

where the configuration appears on the left hand side. This diagram follows our long-standing convention to draw $x^\sigma \in \mathcal{C}(\sigma)$ with events displayed as their image via ∂_σ ,

with the diagram showing immediate causality both in σ (drawn as \rightarrow) and in A (drawn as dotted lines). Thus, its display via ∂_σ is obtained diagrammatically simply as the erasure of the \rightarrow relation, moving towards the right hand side of the diagram.

Our claim is that the diagram on the right hand side of Figure 10.12 corresponds (up to symmetry, see below) to an element in the interpretation of the type $(\mathbb{U} \rightarrow \mathbb{U}) \rightarrow \mathbb{B}$ in the relational model: that point resulting in \mathbf{tt} , where the argument is called twice, and where in turn each call of the argument calls *its* argument twice – written

$$\alpha = ([([\checkmark, \checkmark], \checkmark), ([\checkmark, \checkmark], \checkmark)], \mathbf{tt})$$

in standard relational model notation, but (ignoring copy indices) the reader should be able to see how the same information is conveyed on the right hand side of Figure 10.12: each “atom” in α corresponds to a pair of a question and an answer. Then, moving from concurrent games to the relational model essentially consists in forgetting the relation \rightarrow in all configurations – so that we may even regard concurrent games as the relational model enriched with causal information. We now aim to make this more formal.

Positions. We wish to extract from a board a set matching the relational interpretation. To that end, we start by examining the most constraining constructor, which is

$$!X = \mathcal{M}_f(X),$$

the *bang* of X a set, defined as the set of *finite multisets*. In contrast, the *configurations* of a $--$ -board A correspond to the set $\text{Fam}(A)$ of *families* of non-empty configurations indexed by finite sets of natural numbers (see Proposition 7.1.3). As $\mathcal{M}_f(X)$ may be presented as a quotient of finitely indexed families of elements of X , it is natural to define the positions of a $--$ -board A as configurations up to symmetry:

Definition 10.4.1. Consider A a board.

The set $\mathcal{Pos}(A)$ of *positions* of A is simply defined as the quotient set $\mathcal{C}(A)/\cong_A$.

By convention, we use symbols x, y, z , etc to range over positions – *i.e.* the same letters as for configurations, but with a different font.

As an example, there are *four* positions on the $--$ -board B of Section 8.2.2: the (symmetry classes corresponding to the) configurations \emptyset , $\{\mathbf{q}^-\}$, $\{\mathbf{q}^-, \mathbf{tt}^+\}$, and $\{\mathbf{q}^-, \mathbf{ff}^+\}$. This shows that *positions* do not quite match the relational model just yet, as the relational interpretation of \mathbb{B} only has *two* elements, for the two complete computations.

Stopping positions of boards. Fortunately, boards come equipped with exactly the structure needed to eliminate partial computations: the *payoff function*.

Recall from Definition 8.2.3 that boards come equipped with a *payoff function* $\kappa_A : \mathcal{C}(A) \rightarrow \{-1, 0, +1\}$, which was motivated precisely by the need to capture which configurations are complete, and to designate a responsible for non-complete configurations. By condition *invariant*, the payoff function automatically extends to

$$\kappa_A : \mathcal{Pos}(A) \rightarrow \{-1, 0, +1\}$$

a payoff function for positions. This lets us define:




Figure 10.13: Non-deterministic witnesses




Figure 10.14: Symmetric witnesses

Definition 10.4.2. Consider A a board. Its set of *stopping positions* is the set

$$\mathbf{(A)} = \{x_A \in \mathcal{Pos}(A) \mid \kappa_A(x_A) = 0\}$$

This is indeed the right definition, and we shall see in Section 10.4.4 that it makes the various constructions on boards agree with their relational counterparts.

Stopping positions of strategies. To a board A , the relational collapse will associate a set $\mathbf{(A)}$. If $\sigma : A$ is a strategy on A , we should assign it a subset of $\mathbf{(A)}$. We have seen with Proposition 7.3.1 that causal strategies almost compose relationally already with respect to $+$ -covered configurations, so the next definition is almost self-evident:

Definition 10.4.3. Consider A a board and $\sigma : A$ a causal strategy. Then,

$$\mathbf{(\sigma)} = \{x_A \in \mathbf{(A)} \mid \exists x^\sigma \in \mathcal{C}^+(\sigma), \partial_\sigma x^\sigma \in x_A\}$$

is its *relational collapse*.

For $x^\sigma \in \mathcal{C}^+(\sigma)$ such that $\partial_\sigma x^\sigma \in x_A \in \mathbf{(A)}$, we say that x^σ is a **witness** for x_A ; the relational collapse only keeps stopping positions for which there is a witness. In general, the witness is not unique. We show two examples: in Figure 10.13, the (symmetry class of) $\{q^-, tt^+\}$ has two witnesses, illustrating the two distinct non-deterministic branches realizing the value tt . But the number of witnesses does not necessarily match the number of non-deterministic branches leading to an observable result: in Figure 10.14, there is an *infinite* number of witnesses for the stopping position corresponding to value tt – even though all those witnesses are symmetric with each other³.

10.4.2 Composition and Deadlocks

Now, take A and B two $--$ -boards, and let us consider the case of a strategy *from* A *to* B . Of course the definition above applies, instantiated on the board $A \vdash B$. But in order to get a relation from $\mathbf{(A)}$ to $\mathbf{(B)}$, we shall make use of the following fact:

Lemma 10.4.4. Consider A, B two $--$ -boards. Then, there is a bijection

$$s_{A,B}^\vdash : \mathbf{(A \vdash B)} \simeq \mathbf{(A)} \times \mathbf{(B)}$$

³It is a good intuition that witnesses *should* somehow count non-deterministic branches leading to an observable result, but it requires a much finer definition of witnesses – see Section 13.3.1.

Proof. We know from Lemma 7.1.22 that configurations of $A \vdash B$ all have the form $x_A \vdash x_B$ where $x_A \in \mathcal{C}(A)$ and $x_B \in \mathcal{C}(B)$; this extends to symmetry classes since symmetries enjoy the same unique decomposition. By Definition 8.2.5, $\mathbf{x}_A \vdash \mathbf{x}_B$ has null payoff exactly when both \mathbf{x}_A and \mathbf{x}_B have (recall that $A \vdash B = A^\perp \wp B$ by definition, below Definition 8.2.5). The bijection follows this decomposition. \square

We use this bijection implicitly when defining the collapse of $\sigma : A \vdash B$ as

$$(\sigma) = \{(\mathbf{x}_A, \mathbf{x}_B) \in (A) \times (B) \mid \exists x^\sigma \in \mathcal{C}^+(\sigma), \partial_\sigma x^\sigma \in \mathbf{x}_A \vdash \mathbf{x}_B\}$$

a morphism in $\mathbf{Rel}((A), (B))$ – now, does this preserve the categorical structure?

Copycat. We start the analysis for copycat, which is straightforward:

Proposition 10.4.5. *Consider A a $--$ -board. Then, we have the equality*

$$(\mathfrak{c}_A) = \text{id}_{(A)} \in \mathbf{Rel}((A), (A)).$$

Proof. \subseteq . Consider $\mathbf{x}_A \vdash \mathbf{y}_A \in (\mathfrak{c}_A)$. By definition it has a witness in \mathfrak{c}_A , which by Lemma 6.4.4 has form $x_A \parallel x_A \in \mathcal{C}^+(\mathfrak{c}_A)$ with $x_A \vdash x_A \in \mathbf{x}_A \vdash \mathbf{y}_A$, so $\mathbf{x}_A = \mathbf{y}_A$.

\supseteq . Consider $(\mathbf{x}_A, \mathbf{x}_A) \in \text{id}_{(A)}$ and take any representative $x_A \in \mathbf{x}_A$. Again by Lemma 6.4.4, $x_A \parallel x_A \in \mathcal{C}^+(\mathfrak{c}_A)$, and it is immediate that it is a witness for $(\mathbf{x}_A, \mathbf{x}_A)$. \square

Indeed from Lemma 6.4.4, and with respect to $+$ -covered configurations, copycat already acts like the relational identity! In the light of Proposition 7.3.1, we may hope that the situation for composition will just as simple.

Composition: oplax preservation. For one of the two directions, it is simple:

Lemma 10.4.6. *Consider $--$ -boards A, B, C , and $\sigma : A \vdash B, \tau : B \vdash C$.*

Then, we have $(\tau \odot \sigma) \subseteq (\tau) \circ (\sigma)$.

Proof. Consider $(\mathbf{x}_A, \mathbf{x}_C) \in (\tau \odot \sigma)$. By definition – and inlining Proposition 7.3.1 – there is $x^\tau \odot x^\sigma \in \mathcal{C}^+(\tau \odot \sigma)$ s.t. $\partial_{\tau \odot \sigma}(x^\tau \odot x^\sigma) = x_A^\sigma \vdash x_C^\tau \in \mathbf{x}_A \vdash \mathbf{x}_C$. Thus

$$\partial_\sigma x^\sigma = x_A^\sigma \vdash x_B^\sigma \in \mathbf{x}_A \vdash \mathbf{x}_B, \quad \partial_\tau x^\tau = x_B^\tau \vdash x_C^\tau \in \mathbf{x}_B \vdash \mathbf{x}_C$$

where $x^\sigma \in \mathcal{C}^+(\sigma)$ and $x^\tau \in \mathcal{C}^+(\tau)$ and for \mathbf{x}_B the symmetry class of $x_B^\sigma = x_B^\tau = \mathbf{x}_B$.

The only thing we must check is that $\mathbf{x}_B \in (B)$, i.e. $\kappa_B(\mathbf{x}_B) = 0$. But if we had $\kappa_B(\mathbf{x}_B) = -1$, we would have $\kappa_{A \vdash B}(x_A^\sigma \vdash x_B) = -1$, contradicting that σ is winning. Likewise, $\kappa_B(\mathbf{x}_B) = 1$ would contradict that τ is winning. Hence, $\kappa_B(\mathbf{x}_B) = 0$. \square

This is not by chance: the mechanism of payoff and winning strategies has been set up specifically with this collapse in mind. Now, examine the other direction.




Figure 10.15: Deadlocking composition of causal strategies

Composition: non lax preservation. In the other direction, consider $\sigma : A \vdash B, \tau : B \vdash C$, and $(x_A, x_C) \in (\tau) \circ (\sigma)$. This means that there is $x_B \in (B)$ a stopping position such that $(x_A, x_B) \in (\sigma)$ and $(x_B, x_C) \in (\tau)$; meaning that there are two witnesses

$$x^\sigma \in \mathcal{C}^+(\sigma), \quad x^\tau \in \mathcal{C}^+(\tau)$$

such that $x_A^\sigma \in x_A, x_C^\tau \in x_C$ and $x_B^\sigma, x_B^\tau \in x_B$ – thus there is $\theta_B : x_B^\sigma \cong_B x_B^\tau$. So x^σ and x^τ may not match on B on the nose, but only up to symmetry. This is an issue, but one that we are prepared to handle just as we already did in the proof of Lemma 9.3.5, using Proposition 7.4.4, reindexing both x^σ and x^τ to find a true synchronization. Thus let us ignore that for now, and pretend that we are lucky enough to have $x_B^\sigma = x_B^\tau$.

To sum up, relational composition provides us with a *witness position* that both strategies agree (up to symmetry) is *reachable*. In contrast, 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 a *compatible 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 6.2.2. Figure 10.15 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 the relational model is not lax functorial.

10.4.3 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 10.4.7. *Take A, B –boards, $\sigma : A \vdash B$ visible, and $m, m' \in \sigma$ s.t. $m <_\sigma m'$. Then, $j(m')$ 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 , $\partial_\sigma(j(m'))$ is minimal in B , so $j(m')$ is minimal for \leq_σ by courtesy. But since σ is

pointed, $j(m')$ is initial in ρ , obviously comparable with m as ρ is totally ordered.

Else, by visibility $j(m') \in \rho$. But ρ is totally ordered, so $m, j(m')$ 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 $x^\tau \in \mathcal{C}(\tau)$ s.t., writing $\partial_\sigma x^\sigma = x_A^\sigma \parallel x_B^\sigma$ 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*. Recall from Definition 6.2.2 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 constraints from σ and τ :

$$(l, r) \triangleleft (l', r') \quad \Leftrightarrow \quad l <_{\sigma \parallel C} l' \quad \vee \quad r <_{A \parallel \tau} r'.$$

We saw in Definition 6.2.2 and Proposition 7.3.6 that (x^σ, x^τ) corresponds to a configuration of the interaction $\tau \otimes \sigma$ 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.

Proof sketch. Before giving the formal proof, we showcase the reasoning on a simplified case. The basic idea is by contradiction: starting with a putative deadlock, we repeatedly push it down the causal dependency of the board, until it reaches a minimal event – but those cannot appear in a cycle.

Consider a **simple deadlock** in φ , given by $p_1 = (l_1, r_1)$ and $p_2 = (l_2, r_2) \in \varphi$ s.t.

$$l_1 <_{\sigma \parallel C} l_2, \quad r_2 <_{A \parallel \tau} r_1,$$

an immediate incompatibility between p_1 and p_2 . In other words we have $p_1 \triangleleft p_2$ and $p_2 \triangleleft p_1$; we use $p_1 \triangleleft_\sigma p_2$ and $p_2 \triangleleft_\tau p_1$ to indicate the origin of the causal constraint. We apply the same conventions for polarity of elements of φ as for Lemma 6.2.15.

The first observation (skipped here) is that *w.l.o.g.*, the polarities are as in


$$\begin{array}{ccc} & \triangleleft_\sigma & \\ p_1' & \xrightarrow{\quad} & p_2^\ell \\ & \triangleleft_\tau & \end{array}$$

where both occur in B but not minimal in B – so we may take $j(p_i) = (j(l_i), j(r_i))$. By Lemma 10.4.7, l_1 and $j(l_2)$ are comparable for σ ; while r_2 and $j(r_1)$ are comparable for τ . If $p_1 \triangleleft_\sigma j(p_2)$ or $p_2 \triangleleft_\tau j(p_1)$, then we respectively have one of the cycles:

$$\begin{array}{ccc} p_1' & \xrightarrow{\triangleleft_\sigma} j(p_2)' & \xrightarrow{\triangleleft_\tau} p_2^\ell \\ & \searrow & \swarrow \\ & \triangleleft_\tau & \triangleleft_\sigma \end{array} \quad \text{or} \quad \begin{array}{ccc} p_1' & \xrightarrow{\triangleleft_\sigma} j(p_1)^\ell & \xrightarrow{\triangleleft_\tau} p_2^\ell \\ & \searrow & \swarrow \\ & \triangleleft_\tau & \triangleleft_\sigma \end{array}$$

so simple deadlocks between p_1 and $j(p_2)$; or between $j(p_1)$ and p_2 . Those simple deadlocks are smaller, in the sense that their cumulative *depth* in B has decreased – justifying the inductive reasoning. The case $p_1 = j(p_2)$ or $p_2 = j(p_1)$ is easily discarded.

The last case has $j(p_2) \triangleleft_\sigma p_1$ and $j(p_1) \triangleleft_\tau p_2$. But p_1 has polarity \neq , so by Lemma 10.2.2 the only immediate dependency in σ of l_1^- is $j(l_1)$. So $j(p_2) \triangleleft_\sigma p_1$ factors as $j(p_2) \triangleleft_\sigma j(p_1) \triangleleft_\sigma p_1$. Symmetrically $j(p_1) \triangleleft_\tau j(p_2)$, so we have:



closer to the root of the board. Repeating this we eventually hit an impossible simple deadlock with a minimal event in B , finally exposing the contradiction.

Detailed proof. The proof of the deadlock-free lemma is the same in essence, but challenging in form. Firstly, cycles in \triangleleft in Definition 6.2.2 may have arbitrary length. Secondly, in relational composition strategies synchronize on *symmetry classes* rather than configurations; so we must account for synchronization through symmetry.

Lemma 10.4.8. Consider A, B, C --boards, $\sigma : A \vdash B$ 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 acyclic; i.e. $(x^\sigma, \theta, x^\tau)$ causally compatible.

Proof. We use polarities ℓ, \neq or $-$ for elements of φ (i.e. pairs (l, r)) as above. 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-iso, $(\partial_\sigma \parallel C)(l)$ is minimal in B iff $(A \parallel \partial_\tau)(r)$ is. If not, then $j(l)$ and $j(r)$ also match up to θ and $(j(l), j(r))$ must be in φ as well – we write it $j(l, r)$. 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 $r_i < r_{i+1} < \dots < r_{j-1} < r_j$, 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 cycle has *depth*

$$d = \sum_{1 \leq i \leq n} \text{depth}(l_i, r_i),$$

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).

Next, we 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. 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, $\text{pol}(p_{2k}) = \tau$ and $\text{pol}(p_{2k+1}) = \ell$ so that polarity in the cycle is alternating as well. Indeed, assume *e.g.*

$$p_{2k+1} \triangleleft p_{2k+2}^{\ell} \triangleleft p_{2k+3}$$

with $p_{2k+1} \triangleleft_{\tau} p_{2k+2}$ and $p_{2k+2} \triangleleft_{\sigma} p_{2k+3}$. Then, $r_{2k+1} <_{A \parallel \tau} r_{2k+2}^-$. Negative r_{2k+2}^- cannot be minimal in B , so by Lemma 10.2.2, it has a unique predecessor $j(r_{2k+2}) \rightarrow_{A \parallel \tau} r_{2k+2}$, so $r_{2k+1} <_{A \parallel \tau} r_{2k+2}^-$ factors as $r_{2k+1} <_{A \parallel \tau} j(r_{2k+2}) \rightarrow_{A \parallel \tau} r_{2k+2}^-$. Accordingly, $p_{2k+1} \triangleleft_{\tau} j(p_{2k+2}) \triangleleft_{\tau} p_{2k+2}$ – but dependencies in the game are respected by both strategies, so $j(p_{2k+2}) \triangleleft_{\sigma} p_{2k+2}$. So $j(p_{2k+2}) \triangleleft_{\sigma} p_{2k+3}$, and we can replace the cycle fragment with

$$p_{2k+1} \triangleleft j(p_{2k+2}) \triangleleft p_{2k+3}$$

which is still in B , has the same length but strictly smaller depth. The symmetric argument applies for σ , so any p_{2k+1} has polarity ℓ and any p_{2k+2} has polarity τ .

We show the cycle cannot have an event minimal in B . Seeking a contradiction, if

$$p_{2k+1}^{\ell} \triangleleft_{\tau} p_{2k+2}^{\tau} \triangleleft_{\sigma} p_{2k+3}^{\ell}$$

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 $p_{2k+1}^{\ell} \triangleleft_{\tau} p_{2k+2}^{\ell} \triangleleft_{\tau} p_{2k+3}^{\ell}$ and again the cycle can be shortened by transitivity.

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 $p_{2k}^{\tau} \triangleleft_{\sigma} p_{2k+1}^{\ell}$, writing $p = (l, r) = j(p_{2k+1})$, $l = j(l_{2k+1})$ as well. From Lemma 10.4.7, we know that $l = j(l_{2k+1})$ is comparable with l_{2k} in $\sigma \parallel C$ (by visibility of σ). If $j(l_{2k+1}) = l_{2k}$, 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 $j(l_{2k+1}) \neq l_{2k}$. Similarly, $j(r_{2k+2})$ is comparable with r_{2k+1} in $A \parallel \tau$, but distinct.

Assume that we have $p_{2k} \triangleleft_{\sigma} j(p_{2k+1})$ for some k . Since $j(p_{2k+1}) \triangleleft_{\tau} p_{2k+1} \triangleleft_{\tau} p_{2k+2}$ we can replace the cycle fragment $p_{2k} \triangleleft p_{2k+1} \triangleleft p_{2k+2}$ with the cycle fragment

$$p_{2k} \triangleleft j(p_{2k+1}) \triangleleft p_{2k+2}$$

which has the same length but smaller depth, absurd. So, $j(p_{2k+1}) \triangleleft_{\sigma} p_{2k}$ for all k (symmetrically, $j(p_{2k+2}) \triangleleft_{\tau} p_{2k+1}$ for all k). In particular, $j(l_{2k+1}) <_{\sigma \parallel C} l_{2k}^-$ but by

Lemma 10.2.2, l_{2k}^- has a unique immediate predecessor $j(l_{2k})$. So, $j(p_{2k+1}) \triangleleft_\sigma j(p_{2k})$ for all k ; and likewise $j(p_{2k+2}) \triangleleft_\tau j(p_{2k+1})$ for all k . So we can replace the full cycle with

$$j(p_n) \triangleleft j(p_{n-1}) \triangleleft \dots \triangleleft j(p_1) \triangleleft j(p_n)$$

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

The deadlock-free lemma 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 under-approximation of innocence.

As a direct consequence of the deadlock-free lemma, we obtain:

Corollary 10.4.9. *Consider $--$ -boards A, B, C , and $\sigma : A \vdash B, \tau : B \vdash C$ visible.*

Then, $(\tau \odot \sigma) = (\tau) \circ (\sigma)$.

Proof. \subseteq . Follows from Lemma 10.4.6.

\supseteq . Consider $(x_A, x_B) \in (\sigma)$ and $(x_B, x_C) \in (\tau)$, with witnesses $x^\sigma \in \mathcal{C}^+(\sigma)$ 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, 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 by Lemma 10.4.8. Thus by Proposition 7.4.4, there are reindexings $y^\sigma \cong_\sigma x^\sigma$ 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 $(x_A, x_C) \in (\tau \odot \sigma)$ as required. \square

10.4.4 A Relative Seely \sim -Functor

Next, we show that the relational collapse preserves the rest of the structure – it should form a *relative Seely \sim -functor*. Relative Seely (\sim)-functors are simply (\sim)-functors equipped with isomorphisms expressing the preservation of the additional structure, satisfying certain coherence conditions. Just like relative Seely categories, relative Seely (\sim)-functors are not a standard notion. However, they are certainly an unsurprising one, so in this monograph we opted to relegate it to the appendix (Appendix A.2).

To construct our concrete relative Seely \sim -functor, we start with:

Proposition 10.4.10. *We have a \sim -functor $(-) : \text{NTCG-Vis} \rightarrow \text{Rel}$.*

Proof. Preservation of copycat is by Proposition 10.4.5, and preservation of composition is by Proposition 10.4.9. It remains to show that $(-)$ preserves \approx .

Hence, consider $\sigma, \sigma' : A \vdash B$ with $f : \sigma \approx \sigma'$, and take $(x_A, x_B) \in (\sigma)$. By definition, there is a witness $x^\sigma \in \mathcal{C}^+(\sigma)$ such that $\partial_\sigma x^\sigma \in x_A \vdash x_B$. But then,

$$f[x^\sigma]_A \vdash f[x^\sigma]_B : x_A^\sigma \vdash x_B^\sigma \cong_{A \vdash B} f(x^\sigma)_A \vdash f(x^\sigma)_B$$

so that $f(x^\sigma) \in x_A \vdash x_B$ as required. The other inclusion holds by symmetry. \square

Next, we examine preservation of the full categorical structure.

Preservation of the monoidal structure. Following Definition A.2.1, we require

$$\begin{array}{lcl} s_{A,B}^\otimes & : & (A) \times (B) \cong (A \otimes B) \\ s^1 & : & 1 \cong (1) \end{array}$$

isomorphisms in **Rel**, natural in A and B – those follow from the bijections of Lemma 7.1.22 in the obvious way, which is indeed compatible with payoff.

Proposition 10.4.11. *We have $((-), s^\otimes, s^1)$ a symmetric monoidal \sim -functor.*

Proof. All verifications are straightforward. For naturality, we must show that

$$\begin{array}{ccc} (A) \times (B) & \xrightarrow{s_{A,B}^\otimes} & (A \otimes B) \\ (\sigma) \times (\tau) \downarrow & & \downarrow (\sigma \otimes \tau) \\ (A') \times (B') & \xrightarrow{s_{A',B'}^\otimes} & (A' \otimes B') \end{array}$$

commutes for $\sigma : A \vdash A'$ and $\tau : B \vdash B'$, which is direct by Proposition 8.1.1. \square

Preservation of !. Next we deal with the exponential structure, which is the most challenging. Inspecting the requirements for relative Seely functors we must first show that $(-)$ preserves strict objects – but this is trivial: as a proper Seely category, **Rel** is regarded as a relative Seely category with all objects strict. Then we must provide

$$s_S^! : \mathcal{M}_f(!S) \cong (!S)$$

for every strict S , subject to coherence conditions. We first construct the bijection:

Proposition 10.4.12. *For any S strict, there is a bijection $s_S^! : \mathcal{M}_f(!S) \cong (!S)$.*

Proof. First, recall from Lemma 8.3.1 that we have a bijection

$$\text{Fam}(\mathcal{C}^{\neq \emptyset}(S)) \simeq \mathcal{C}(!S)$$

following which configurations of null payoff of $!S$ correspond to families of non-empty configurations of null payoff of S . But S is strict, thus its configurations of null payoff are non-empty. Hence the bijection above may be directly refined into another bijection

$$\text{Fam}(\llbracket S \rrbracket) \simeq \llbracket !S \rrbracket$$

where $\llbracket A \rrbracket$ is the set of configurations of null payoff. Inlining this bijection, we may define the inverse of $s_S^!$ on representatives directly as follows:

$$\begin{array}{lcl} (s_S^!)^{-1} & : & \llbracket !S \rrbracket \rightarrow \mathcal{M}_f(!S) \\ [x_i^S \mid i \in I] & \mapsto & [x_i^S \mid i \in I] \end{array}$$

where x_i^S denotes the symmetry class of x_i^S , $[x_i^S \mid i \in I]$ (with a bold font) is our notation for families, and $[x_i^S \mid i \in I]$ (with non-bold font) is multiset comprehension.

Invariant. Consider $\theta : [x_i^S \mid i \in I] \cong_{!S} [y_j^S \mid j \in J]$ a symmetry, which by definition is given by a bijection $\pi : I \simeq J$ and a family $(\theta_i)_{i \in I}$ where $\theta_i : x_i^S \cong_S y_{\pi(i)}^S$ for all $i \in I$. But then, $x_i^S = y_{\pi(i)}^S$, from which, as required, we deduce that

$$[x_i^S \mid i \in I] = [y_{\pi(i)}^S \mid i \in I] = [y_j^S \mid j \in J].$$

Injective. Assume $[x_i^S \mid i \in I] = [y_j^S \mid j \in J]$. Thus, there is a bijection $\pi : I \simeq J$ such that for all $i \in I$, $x_i^S = y_{\pi(i)}^S$. Hence, for all $i \in I$ there is $\theta_i : x_i^S \cong_S y_{\pi(i)}^S$, altogether forming a symmetry $\theta : [x_i^S \mid i \in I] \cong_S [y_j^S \mid j \in J]$.

Surjective. Consider a multiset $[x_1, \dots, x_n] \in \mathcal{M}_f((S))$. For each $1 \leq i \leq n$, consider a representative $x_i \in x_i$. Then, $[x_i \mid 1 \leq i \leq n]$ provides a pre-image. \square

Observe that there is no such bijection $\mathcal{M}_f((A)) \simeq (!A)$ for A non-strict, because

$$[\emptyset] \neq [\emptyset, \emptyset]$$

are distinct elements of $\mathcal{M}_f((A))$, but can only be represented in $(!A)$ with the same (empty) configuration since there are no events. The relational model may remember how many times a game is *not* visited, which makes no sense in terms of games. This is the reason behind the preservation of the *relative* Seely \sim -categorical structure only.

Among the several diagrams to check, the most subtle is preservation of promotion:

Lemma 10.4.13. *Consider S, T strict and $\sigma : !S \vdash T$.*

*Then, the following square commutes in **Rel**:*

$$\begin{array}{ccc} \mathcal{M}_f((S)) & \xrightarrow{((\sigma) \circ s_S^!)^!} & \mathcal{M}_f((T)) \\ s_S^! \downarrow & & \downarrow s_T^! \\ (!S) & \xrightarrow{(\sigma^!)} & (!T) \end{array}$$

Proof. We use a multiset notation for stopping positions of $!S$ and $!T$, inlining Proposition 10.4.12. The upper-right path computes to the pairs that can be written as

$$([x_{i,j} \mid i \in I, j \in J_i], [y_i \mid i \in I])$$

for $x_{i,j} \in (S)$ and $y_i \in (T)$ where we have that for all $i \in I$, $([x_{i,j} \mid j \in J_i], y_i) \in (\sigma)$. Thus for all $i \in I$, there is $x_i^\sigma \in \mathcal{C}^+(\sigma)$ witnessing that pair, *i.e.* we have

$$\partial_\sigma x_i^\sigma \in [x_{i,j} \mid j \in J_i] \vdash y_i.$$

Thus, relying on Proposition 8.3.3 we may form $[x_i^\sigma \mid i \in I] \in \mathcal{C}^+(\sigma^!)$ (by Lemma 8.1.8, $\sigma^!$ only differs from $!\sigma$ via its display map), and a direct verification ensures that

$$\partial_{\sigma^!} [x_i^\sigma \mid i \in I] \in ([x_{i,j} \mid i \in I, j \in J_i], [y_i \mid i \in I])$$

as required. The other direction follows the same decompositions. \square

Relative Seely \sim -functor. To wrap up, we introduce the missing components. First,

$$s^\top : \emptyset \cong (\top)$$

is simply the empty relation. For the product and the arrow, we prove:

Lemma 10.4.14. *Consider A, S, T $--$ boards with S, T strict. Then, we have bijections*

$$\begin{aligned} s_{S,T}^\& : (S) + (T) &\cong (S \& T) \\ s_{A,S}^\multimap & : (A) \times (S) &\cong (A \multimap S) \end{aligned}$$

Proof. For the product, recall first that by Lemma 8.2.9, we have a bijection

$$\mathcal{C}^{\neq \emptyset}(S) + \mathcal{C}^{\neq \emptyset}(T) \simeq \mathcal{C}^{\neq \emptyset}(S \& T)$$

which holds also for symmetries, and thus extends to symmetry classes. But now we note that S, T , and $S \& T$ are strict, therefore configurations of null payoff must be non-empty. It follows that the bijection above extends to stopping positions.

For the arrow, similarly recall that by Lemma 8.2.12, we have a bijection

$$\mathcal{C}^{\neq \emptyset}(A \multimap S) \simeq \mathcal{C}(A) \times \mathcal{C}^{\neq \emptyset}(S)$$

which also holds to symmetry and thus extends to symmetry classes. But again note that since S is strict, configurations with null payoff must be non-empty, and thus the above extends to a bijection on stopping positions. \square

The missing five coherence diagrams of relative Seely \sim -functors are direct, from an analysis of the symmetry classes reached by the component strategies. Projections, dereliction and monoidality are defined by lifting, hence the corresponding diagrams follow as for copycat in Proposition 10.4.5. For evaluation, this follows from the description of its $+$ -covered configurations in Proposition 8.2.27. Altogether:

Corollary 10.4.15. *We have a relative Seely \sim -functor $(-) : \mathbf{NTCG}\text{-}\mathbf{Vis} \rightarrow \mathbf{Rel}$.*

By Theorem A.2.3 we have a cartesian closed \sim -functor $(-) : \mathbf{NTCG}\text{-}\mathbf{Vis}_! \rightarrow \mathbf{Rel}_!$.

Once the deadlock-free lemma is established, this interpretation-preserving collapse from concurrent games to the relational model is strikingly simple, and emphasizes how the concurrent games model may be regarded as a causal and temporal extension of the relational model. This extension is of course far from negligible as it allows one to represent programming languages with non-commutative effects, such as references.

Positional equivalence. In this monograph, we shall exploit this concretely by reasoning on strategies up to the induced equivalence relation:

Definition 10.4.16. *Consider A, B $--$ boards, and $\sigma, \sigma' \in \mathbf{NTCG}\text{-}\mathbf{Vis}(A, B)$.*

*We say that σ and σ' are **positionally equivalent**, written $\sigma \equiv \sigma'$, iff $(\sigma) = (\sigma')$.*




Figure 10.16: Quotienting out evaluation order

$\frac{x_A \in \{A\} \quad x_B \in \{B\}}{x_A \otimes x_B \in \{A \otimes B\}}$		$\frac{x_A \in \{A\} \quad x_C \in \{C\}}{x_A \circ x_C \in \{A \circ C\}}$
$\frac{x \in \{A\}_i \quad (i \in I)}{(i, x) \in \{\&_{i \in I} A_i\}}$	$\frac{(x_C^i \in \{C\})_{i \in I}}{[x_C^i \mid i \in I] \in \{!C\}}$	$\frac{x_A \in \{A\} \quad x_B \in \{B\}}{x_A \vdash x_B \in \{A \vdash B\}}$

Figure 10.17: Syntax for positions for A, B, C --boards with C strict

Concretely, positional equivalence amounts to quotienting out information about the evaluation order, as illustrated in Figure 10.16. It is a very brutal quotient, which would clearly be inconsistent in the presence of any non-commutative effect. But since $(-)$ is a relative Seely \sim -functor, positional equivalence *is* a congruence within visible causal strategies: it is preserved by all operations used in computing the interpretation. In other words, there is another relative Seely \sim -category where *visible* strategies are considered up to positional equivalence – we shall not move to this \sim -category, but the finite definability result of Chapter 12 will hold up to positional equivalence only.

Back to \rightarrow -Strat. The developments presented in this chapter have no interaction with mixed boards and well-bracketing, hence transport transparently from \mathbf{NTCG}_1 to \rightarrow -Strat. Hence there are \rightarrow -Vis and \rightarrow -Inn, two cartesian closed sub- \sim -categories of \rightarrow -Strat with morphisms respectively restricted to *visible* and *parallel innocent* strategies. They both support the interpretation of $\mathbf{PCF}_{//}$, with cartesian closed \sim -functors:

$$\begin{aligned} (-) &: \rightarrow\text{-Vis} \rightarrow \mathbf{Rel}_1 \\ (-) &: \rightarrow\text{-Inn} \rightarrow \mathbf{Rel}_1; \end{aligned}$$

preserving the interpretation. Consequently, \rightarrow -Vis and \rightarrow -Inn both support positional equivalence which is preserved by all operations used in the interpretation of $\mathbf{PCF}_{//}$.




Figure 10.18: A well-bracketed innocent strategy with a non well-bracketed gcc

A syntax for positions. It will be convenient to have a syntactic description of the shape of positions on various \multimap -boards constructions. We give this description in Figure 10.17, directly applying the bijection $s_{A,B}^\otimes$ from Proposition 10.4.11, $s_{A,C}^\circ$ and $s_{C,D}^\&$ from Lemma 10.4.14, $s_S^!$ Proposition 10.4.12, and the obvious variant of $s_{A,B}^\otimes$ for \vdash .

10.5 Globularity

Though we regard parallel innocence (and *visibility*) as the key conditions to understand the causal shape of pure parallel programs, it is not yet precise enough to entail our definability result. For definability purposes, we need yet again more rigid constraints on the causal shape of strategies, taking the form of a final condition called *globularity*.

10.5.1 Motivation and Definition

Well-bracketed gccs. In Section 10.1.2, we introduced *pre-innocence* with the intuition that a gcc may be regarded as a standalone sequential program, and pre-innocence brings constraints as to how these “sequential threads” may be forked and merged. But if these “threads” do correspond to good old sequential computations, they should certainly satisfy a condition akin to sequential well-bracketing (see Definition 3.2.5).

Strategies in \multimap -**Strat** already satisfy a well-bracketing condition (Definition 9.3.2). However, it is based in the *logical* well-bracketing of non-alternating plays (Definition 5.1.2), which is insufficient to ensure that gccs are (sequentially) well-bracketed: this is illustrated in Figure 10.18, displaying a well-bracketed innocent causal strategy with a non well-bracketed gcc highlighted. This invites the condition, for σ in \multimap -**Inn**:

wb-threads: any $\rho \in \text{gcc}(\sigma)$ is well-bracketed in the sense of Definition 3.2.5.

However, this is not yet enough to fully capture higher-order pure parallel computation. The final condition required has to do with which gccs can be legally merged.





Figure 10.19: An innocent non-globular strategy

Intuitively, all causal merges in PCF_{\parallel} are generated by the parallel *let* binding

$$\text{let } \begin{pmatrix} x_1 & = & N_1 \\ x_2 & = & N_2 \end{pmatrix} \text{ in } M$$


where the two branches are causally merged when the computation of N_1 has *returned*, and the computation of N_2 has *returned* also – so all Questions involved in the computation of N_1 have an answer, and likewise for N_2 . But not all parallel innocent strategies satisfy this pattern: the strategy shown in Figure 10.19 is parallel innocent, satisfies condition *wb-threads*, but is not definable within PCF_{\parallel} as there is a causal merge of computation branches that have not *returned* yet. Banning this invites the condition:

globules: for any diagram with $X = \{m_1, \dots, m_n\}$ and $Y = \{n_1, \dots, n_p\}$ disjoint:



in σ , then every question in X (resp. Y) is answered in X (resp. Y).

We call such a diagram a *globule*. Note that 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. Additionally, one can also observe that globules must always have Question/Answer assignments as in the diagram:



Indeed, if m_1 was an answer, it would be maximal in σ (by Lemma 6.1.16, as by *answer-closing*, answers are maximal in A) and the merge would be impossible. So m_1 is a question, and by *globules* it has an answer in $\{m_1, \dots, m_{2n}\}$. By Lemma 6.1.16 this answer depends immediately on m_1 in σ and so must be m_2 . Repeating this we get the description above. Hence by *globules*, parallel threads that eventually might merge have to follow a strict call/return discipline. Altogether, we set:

Definition 10.5.1. Consider A, B mixed boards, and $\sigma \in \rightarrow\text{-}\mathbf{Inn}(A, B)$.
Then σ is **globular** if it satisfies conditions *wb-threads* and *globules*.

10.5.2 Composition of Globularity

In this section, we develop the compositional structure of globularity.

Basic categorical structure. As usual when introducing a new condition on strategies, we must show that it is compatible with all our constructions on strategies – and as is often the case, the most crucial point is compatibility with *copycat* and *composition*.

Recall that the definition of globularity concerns morphisms in $\rightarrow\text{-}\mathbf{Inn}$, defined as (a subcategory of) a Kleisli category, *i.e.* its morphisms are certain strategies on $!A \vdash B$ where A and B are mixed boards. In examining compatibility of globularity with the categorical structure, it is convenient to work with the basic constructions in **NTCG** (*copycat*, *composition*) rather than those in the Kleisli category (Kleisli composition). Here we shall do that, keeping in mind the slight abuse that technically the definition of globularity refers to the Question/Answer labelling, which was not required on boards, but arises only in mixed boards, via the link with the arena.

Proposition 10.5.2. The copycat strategy $\mathfrak{c}_A : A \vdash A$ is globular.


Proof. The condition *wb-threads* boils down to the proof that P-views of the alternating copycat are well-bracketed in the sense of Definition 3.2.5. *Globularity* holds vacuously, as the event structure for copycat has no causal merge. \square

As usual, this reasoning applies to all copycat strategies involved in the interpretation: projections, structural morphisms for the monoidal structure, dereliction, etc.

Now, we establish stability of globularity under composition:

Proposition 10.5.3. Consider $\sigma : A \vdash B$ and $\tau : B \vdash C$ globular.
Then, $\tau \odot \sigma : A \vdash C$ is globular.

Proof. For *wb-threads*, we first prove that all gccs in $\tau \odot \sigma$ are well-bracketed – this is a variation of the fact that alternating innocent strategies are well-bracketed iff their meager form is (Proposition 3.2.13). For *globules*, by the same reasoning as in Proposition 10.3.2, a globule in $\tau \odot \sigma$ originates in a diagram picturing causal chains



in $\tau \odot \sigma$, constructed as and satisfying the same properties as in Proposition 10.3.2 – in particular n has polarity ℓ without loss of generality. But then as in Proposition 10.3.2,

$$\text{gce}(\ulcorner \kappa_{\leq n}^1 \urcorner^\sigma, \ulcorner \kappa_{\leq p}^2 \urcorner^\sigma) = \text{gce}(\kappa_{\leq n}^1, \kappa_{\leq p}^2)_\sigma = m'_\sigma,$$

by the forking lemma, *i.e.* σ sees the same fork as in $\tau \otimes \sigma$. Since σ is *globular*, all questions between m'_σ and n_σ in both $\ulcorner \kappa^1_{\leq n} \urcorner^\sigma$ and $\ulcorner \kappa^2_{\leq n} \urcorner^\sigma$ have an answer. But as observed above, gccs in $\tau \otimes \sigma$ are well-bracketed, which entails that all questions between m' and n in κ^1 and κ^2 also have an answer – hence, *globules* is satisfied in $\tau \odot \sigma$. \square

Interpretation of PCF_\parallel . As for parallel innocence, preservation of globularity for all other operations is straightforward. We omit the details of that, and of the fact that the basic primitives of PCF_\parallel yields globular strategies. Altogether, we get:

Theorem 10.5.4. *there is $\rightarrow\text{-Glob}$, a cartesian closed sub- \sim -category of $\rightarrow\text{-Strat}$, with morphisms restricted to globular strategies.*

Moreover, $\rightarrow\text{-Glob}$ supports the interpretation of PCF_\parallel .

10.6 History and Related Work

Parallel innocence and its accompanying theory (the forking lemma and stability under composition, the deadlock-free lemma and the collapse to **Rel**, globularity) have been developed in close collaboration with Simon Castellan during his PhD thesis⁴ – they first appear in published form in [Castellan et al., 2015]; though the presentation given here differs significantly, having matured quite a bit since then.

Parallel innocence was first developed aiming for a finite definability result for (essentially) PCF_\parallel , which was achieved in [Castellan et al., 2015]. It became a cornerstone of subsequent developments in concurrent games. Though surprisingly, it is not *parallel innocence* per se that turned out so useful in our later work, but *visibility* and the deadlock-free lemma! Indeed, this is one of the key ingredients in our subsequent results relating concurrent games and relational models, in the probabilistic [Castellan et al., 2018b], the general weighted case [Clairambault and Paquet, 2021], the quantum case [Clairambault and de Visme, 2020] and more recently for generalized species of structure [Clairambault et al., 2023b]. In a sense, the deadlock-free lemma explains why the relational models work at all for pure parallel programs: no deadlocks can arise in composing them, so that causal information has no impact on possible synchronizations and can safely be forgotten. No such miracle occurs for effectful programming languages, for which the causal information is unavoidable.

Related work. In [Melliès and Mimram, 2007], Melliès and Mimram have sketched a notion of innocent strategies in the setting of non-alternating asynchronous games (which are essentially the linear, causally deterministic part of **TCG**), with the purpose of eliminating deadlocks from composition. This appears with more details in Mimram’s PhD thesis [Mimram, 2008]. According to their definition, innocence amounts to a correctness criterion inspired from linear logic proof nets. We are not aware of a direct link with the present developments, though this would be interesting to explore.

⁴More precisely, visibility and parallel innocence were invented on the 7th of february 2014, a (probably) rainy friday – as witnessed by Figure 8.5.

Chapter 11

Sequentiality

The conditions of *parallel innocence* and *globularity* developed in the previous chapter enforce *pure computation*, which may still be parallel. In contrast, the conditions developed in this chapter aim to reinstate *sequentiality*, which may still be impure.

Sequentiality is probably the simplest of our conditions on strategies, but it still comes with a few subtleties. Firstly, the concept of sequentiality is of course tied to *time*, which is notably absent from concurrent strategies! So it makes sense that the definition will have to refer to linearizations of strategies, in the style developed in the unfolding of causal to non-alternating strategies in Chapter 9. Secondly, a phenomenon occurs that is similar to what we encountered for parallel innocence: as *non-determinism* arises from the combination of parallelism and state, *sequentiality* must reinstate it – but not quite under the same form as in the previous chapter: it is hard to unify *sequential* determinism and *causal* determinism (see discussion in Section 14.2.2).

Outline. The developments of this chapter will proceed as follows. In Section 11.1, we develop sequentiality, aiming to show that sequential strategies form a sub-relative Seely \sim -category **NTCG-Seq**, that refines to a sub-cartesian closed \sim -category \rightarrow -**Seq** of \rightarrow -**Strat**, supporting the interpretation of IA. In Section 11.2, we show that the non-alternating unfolding \mathcal{U} -Unf of Chapter 9 refines into an interpretation-preserving

$$\Downarrow\text{-Unf} : \rightarrow\text{-Seq} \rightarrow \Downarrow\text{-Strat}$$

which entails that \rightarrow -**Seq** is intensionally fully abstract for IA. In Section 11.3, we show that in combination with *globularity*, \Downarrow -Unf yields an interpretation-preserving functor

$$\Downarrow\text{-Unf} : \rightarrow\text{-SeqGlob} \rightarrow \Downarrow\text{-InnWB}$$

proving \rightarrow -**SeqGlob** intensionally fully abstract for PCF.

11.1 Sequentiality

Sequentiality has no interaction with mixed boards or well-bracketing, hence we carry it out in the relative Seely \sim -category **NTCG**. Likewise, for succinctness, the examples shall be based on the *small* interpretation of ground types (see Section 9.1.1) – although the actual interpretation remains the same as for $\mathbf{IA}_{//}$, *i.e.* large.

11.1.1 Definition of Sequentiality


What is a sequential causal strategy? The right answer is more subtle than it looks.

Sequential event structures. As strategies are certain event structures, it seems reasonable to start by looking at what is a sequential event structure. Event structures are concurrent generalizations of trees or forests: there is a natural representation of forests as event structures, which entails a canonical notion of sequentiality:

Definition 11.1.1. An event structure E is **sequential** if it satisfies the conditions:

- forestial: for all $e_1, e_2 \leq_E e'$, then $e_1 \leq_E e_2$ or $e_2 \leq_E e_1$.
- #-branching: for all $e_1, e_2 \in E$, $e_1 \#_E e_2$ iff they are causally incomparable.

In other words, a sequential event structure has the shape of a forest as in



whose configurations, always totally ordered, correspond to branches of the forest.

However, it is quickly apparent that this is not useful as a definition of sequential causal strategies. Technically, *receptivity* entails that conflicts between negative moves should be exactly as in the game. But negative events may not conflict in the game, as in *e.g.* the initial moves of **!B**; so no strategy can actually be sequential on **!B**.

More fundamentally, it turns out that although **IA** is indeed a sequential language, the shape of causality in the corresponding strategies is certainly *not* sequential. We show two examples: first in Figure 11.1 we repeat the example of Figure 10.3. This is (an affine linear version of) the interpretation of the term in (10.1), a term of **IA** which should hence be deemed sequential – we see that even though the program itself does not spawn parallel threads, the strategy cannot be sequential as in Definition 11.1.1 as it must show how the program reacts when Opponent does. A similar phenomenon occurs for pure sequential innocent programs: Figure 11.2 shows a configuration of

$$\vdash \lambda f^{\mathbb{U} \rightarrow \mathbb{U}}. f \text{ skip}; f \text{ skip}; \mathbf{tt} : (\mathbb{U} \rightarrow \mathbb{U}) \rightarrow \mathbb{B}$$

which has a typical shape for pure sequential programs. We see that although the causal shape of pure sequential programs is indeed forestial, the branches do *not* conflict.




Figure 11.1: A sequential strategy Figure 11.2: A sequential innocent configuration

Alternating linearizations. So how shall we instead define sequentiality?

Let us set the stage for the definition. First, recall that the definition of *linearizations* (Definition 9.2.1) associates to any event structure the set $\mathcal{L}(E)$ of words of $|E|$, which are non-repetitive and whose prefixes form configurations. Now through its display map, σ has polarities, so we can look at those *alternating linearizations*:

Definition 11.1.2. Consider E an event structure with polarities.

The set $\Downarrow\text{-}\mathcal{L}(E)$ of **alternating linearizations** comprises those $s \in \mathcal{L}(E)$ such that polarities alternate and the first move (if s is non-empty) is negative.

We insist that $\Downarrow\text{-}\mathcal{L}(\sigma)$ comprises words of events of σ : we have not yet projected σ to the game. Its elements are internal witnesses of sequential executions, invisible to the environment. Despite this, we can perform on $\Downarrow\text{-}\mathcal{L}(\sigma)$ certain operations usually done on alternating legal plays on arenas. In particular, we can compute the *P-view*:

Definition 11.1.3. Consider $\sigma \in \text{NTCG}(A, B)$ pointed, and $s \in \Downarrow\text{-}\mathcal{L}(\sigma)$.

Its **P-view** is the subsequence $\ulcorner s \urcorner$ defined by:

$$\begin{aligned} \ulcorner \varepsilon \urcorner &= \varepsilon \\ \ulcorner sm^+ \urcorner &= \ulcorner s \urcorner m^+ \\ \ulcorner sn^+ s' m^- \urcorner &= \ulcorner s \urcorner n^+ m^- && \text{if } m \text{ points to } n, \\ \ulcorner sm^- \urcorner &= m^- && \text{if } m \text{ is minimal,} \end{aligned}$$

where m **points to** n if $\partial_\sigma(n) \rightarrow_{A \vdash B} \partial_\sigma(m)$; or if $\partial_\sigma(m)$ is minimal in A and $n = \text{init}(m)$.

We say $s \in \Downarrow\text{-}\mathcal{L}(\sigma)$ is **P-visible** if for all $t \sqsubseteq s$, $\ulcorner t \urcorner \in \Downarrow\text{-}\mathcal{L}(\sigma)$.

This mimics the traditional definition (Definition 3.2.7), except that sequences $s \in \Downarrow\text{-}\mathcal{L}(\sigma)$ do not have pointers; those are derived from the immediate causality in the game as in Lemma 9.2.8. Observe that this is where the hypothesis *pointed* is needed: as causal strategies play on $A \vdash B$, the causal dependency from initial moves in B to initial moves in A may only be recovered from the immediate causality in σ : pointedness ensures that moves minimal in A have a unique initial dependency which must be minimal in B , which ensures that that the justifier is always well-defined.

Sequentiality. Programs of IA will respect sequentiality only within a sequential execution: we wish to impose that in any state reachable in a sequential execution, Player will not throw parallel threads. But as explained in Section 10.1, as for parallel innocence, sequentiality must also reinstate a notion of determinism. Altogether we get:

Definition 11.1.4. A strategy $\sigma \in \text{NTCG}(A, B)$ is *sequential* if:

- pointed: if $s \in \sigma$, there is a unique $\text{init}(s) \leq_\sigma s$ minimal in σ ,
- sequential determinism: if $tn_1^+, tn_2^+ \in \Downarrow\text{-}\mathcal{L}(\sigma)$, then $n_1 = n_2$,
- sequential visibility: if $s \in \Downarrow\text{-}\mathcal{L}(\sigma)$, writing $\ulcorner s \urcorner = s_1 \dots s_n$ and $1 \leq j \leq n$,
if s_j is non-minimal then there is $1 \leq i < j$ s.t. $j(s_j) = s_i$,

where the justifier $j(m)$ of $m \in \sigma$ refers to Definition 10.2.1.

For *sequential determinism*, more than asking that the unfolding acts deterministically on alternating plays, this condition ensures that no *internal* non-deterministic choice is alternatingly reachable, even when this choice would yield no observable non-deterministic behaviour (surprisingly for us, this turned out necessary).

Sequential visibility is perhaps puzzling, as P-visibility of alternating strategies (Definition 3.2.8) is usually associated not to sequentiality, but to the absence of higher-order state (see Section 3.3.2). 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 **NTCG** the causality due to the syntax tree blurs with that passing via state or semaphores. So strategies arising from the interpretation of IA_{\parallel} are “morally” P-visible but formalizing this is nontrivial¹. For us it is not worth the trouble as P-visibility is not required for full abstraction for IA_{\parallel} ². Consequently, it suffices to reinstate it once we restrict to sequential strategies.

The strategy in Figure 10.6 is not sequential: it fails sequential determinism as two Player moves are available after the initial move. In contrast, the strategy in Figure 11.1 is sequential: although it may perform non-deterministic choices and play Player moves in parallel, none of these behaviours are reachable in an alternating execution.

We mention a direct consequence of sequentiality:

Lemma 11.1.5. Consider $\sigma \in \text{NTCG}(A, B)$ a sequential causal strategy.

Then, it automatically satisfies the following additional condition:

- reachable sequentiality: if $tn^+ \in \mathcal{L}(\sigma)$, if $t \in \Downarrow\text{-}\mathcal{L}(\sigma)$ then $tn \in \Downarrow\text{-}\mathcal{L}(\sigma)$.

Proof. Consider $tn^+ \in \mathcal{L}(\sigma)$ with $t \in \Downarrow\text{-}\mathcal{L}(\sigma)$. If $tn \notin \Downarrow\text{-}\mathcal{L}(\sigma)$, this necessarily means that tn^+ fails alternation, i.e. $t = sm^+$. But then, we argue that $sn^+ \in \mathcal{L}(\sigma)$ as well. Indeed, this may fail only if $m^+ \rightarrow_\sigma n^+$. But by *courteous*, this implies $\partial_\sigma(m) \rightarrow_{A+B} \partial_\sigma(n)$ as well. But as a board, $A \vdash B$ is *alternating*, contradiction. \square

¹This was done by Laird in the first interleaving games model [Laird, 2001b], via threading information.

²Proposition 5.1.11 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.

This means that restricting a sequential strategy to its alternating linearizations only cuts at Opponent moves: it never prevents Player from playing.

11.1.2 Categorical Structure

We show that sequential causal strategies form a sub- \sim -category of **NTCG**.

Sequentiality of copycat. First, we show that the copycat strategy is sequential:

Lemma 11.1.6. *Consider A a $--$ -board. Then, $\mathfrak{c}_A : A \vdash A$ is sequential.*

Proof. First, \mathfrak{c}_A is pointed, as its causal shape is forestial. For the other two conditions, we observe that the even-length alternating linearizations of \mathfrak{c}_A are exactly those words on $s \in |A_1 \vdash A_2|^*$ (with indices used for disambiguation only) such that for all even-length prefix $t \sqsubseteq s$, $t \upharpoonright A_1 = t \upharpoonright A_2$, with the restriction \upharpoonright defined in the obvious way. The two conditions *sequential determinism* and *sequential visibility* follow easily. \square

Composition of sequential determinism. As often when introducing a new condition, much of the work required is in checking stability under composition. Here in particular, something is striking: recall from Section 4.1.3 that proving associativity of composition for alternating strategies involves the study of a “state diagram” for interactions (in contrast, nothing of the sort shows up when constructing the associator for the composition of causal strategies) – this diagram intervenes again here.

Let us start by introducing some further terminology. Consider $\sigma : A \vdash B$ and $\tau : B \vdash C$ two causal strategies, and $p \in \tau \circledast \sigma$ an event of the interaction. We have analyzed above Lemma 6.2.14 the different possible polarities of p – at least one of $p_\sigma \in \sigma$ and $p_\tau \in \tau$ must be defined, possibly both if p occurs in B , in which case p does not appear in the composition. Otherwise p occurs in A or C and we have $p \in \tau \odot \sigma$; let us write $p_\odot = p \in \tau \odot \sigma$ if that is the case, and undefined otherwise.

Applying these projections move-by-move, from $s \in \mathcal{L}(\tau \circledast \sigma)$ we get:

Lemma 11.1.7. *Consider $s = s_1 \dots s_n \in \mathcal{L}(\tau \circledast \sigma)$. Then, forming*

$$\begin{aligned} s_\sigma &= (s_1)_\sigma \dots (s_n)_\sigma \\ s_\tau &= (s_1)_\tau \dots (s_n)_\tau \\ s_\odot &= (s_1)_\odot \dots (s_n)_\odot, \end{aligned}$$

removing moves where the projection is undefined, we have that

$$u_\sigma \in \mathcal{L}(\sigma), \quad u_\tau \in \mathcal{L}(\tau), \quad u_\odot \in \mathcal{L}(\tau \odot \sigma).$$

Proof. From the definition, it is immediate that for all $x^\sigma \in \mathcal{C}(\sigma)$, $x^\tau \in \mathcal{C}(\tau)$ causally compatible, we have $(x^\tau \circledast x^\sigma)_\sigma = x^\sigma \in \mathcal{C}(\sigma)$ (where $(x^\tau \circledast x^\sigma)_\sigma$ is obtained as for linearizations above, applying $(-)_\sigma$ event-by-event), $(x^\tau \circledast x^\sigma)_\tau = x^\tau \in \mathcal{C}(\tau)$ and $(x^\tau \circledast x^\sigma)_\odot \in \mathcal{C}(\tau \odot \sigma)$ (by definition of $\tau \odot \sigma$). The claims on linearizations follow. \square

A linearization $u \in \mathcal{L}(\sigma)$ may be – or not – alternating, *i.e.* $u \in \Downarrow\text{-}\mathcal{L}(\sigma)$. If it is, then as for legal plays on arenas (Definition 3.1.4) two cases are possible: if it has even length (so that the last move is by Player), we say that it is **in state O** . If it has odd length (so that the last move is by Opponent), we say that it is **in state P** .

Putting the projections of the lemma above to work, we have the state diagram:

Lemma 11.1.8. *Consider A, B, C --boards, and $\sigma : A \vdash B$, $\tau : B \vdash C$ sequential.*

For any $u \in \mathcal{L}(\tau \circledast \sigma)$ such that $u_\circ \in \Downarrow\text{-}\mathcal{L}(\tau \circ \sigma)$ we have $u_\sigma \in \Downarrow\text{-}\mathcal{L}(\sigma)$ and $u_\tau \in \Downarrow\text{-}\mathcal{L}(\tau)$, and we are in one of the following three cases:

- (1) $u_\sigma, u_\tau, u_\circ$ are respectively in state O, O, O ,
- (2) $u_\sigma, u_\tau, u_\circ$ are respectively in state O, P, P ,
- (3) $u_\sigma, u_\tau, u_\circ$ are respectively in state P, O, P .

Proof. By induction on u . If u is empty, this is clear. Consider $um \in \mathcal{L}(\tau \circledast \sigma)$. 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 IH, u_σ is alternating in state O , so ends with a Player move. But so, $u_\sigma m_\sigma^+ \in \mathcal{L}(\sigma)$ with $u_\sigma \in \Downarrow\text{-}\mathcal{L}(\sigma)$, so $u_\sigma m_\sigma \in \Downarrow\text{-}\mathcal{L}(\sigma)$ since σ satisfies reachable sequentiality, contradiction. So, m occurs in A or C – assume *w.l.o.g.* in A . Since u_\circ is in state O , m is negative – then, it is direct that um satisfies (3).

(2) First assume m occurs in A . Since $u_\circ \in \Downarrow\text{-}\mathcal{L}(\tau \circ \sigma)$ 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 \neq and we transition to (3), and if m occurs in C it has polarity \neq and we transition to (1).

(3) Symmetric to (2). □

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 4.2. None of the interacting agents can be the first to break alternation, so the interaction ends up alternating. It follows that $\tau \circ \sigma$ satisfies *sequential determinism*:

Proposition 11.1.9. *Consider $\sigma : A \vdash B$ and $\tau : B \vdash C$ sequential causal strategies.*

Then, $\tau \circ \sigma$ satisfies sequential determinism.

Proof. Consider $tn_1^+, tn_2^+ \in \Downarrow\text{-}\mathcal{L}(\tau \circ \sigma)$ completed to $u_1 n_1, u_2 n_2 \in \mathcal{L}(\tau \circledast \sigma)$, and u' the greatest common prefix of $u_1 n_1$ and $u_2 n_2$, with $u' m_1 \sqsubseteq u_1 n_1$ and $u' m_2 \sqsubseteq u_2 n_2$. Necessarily, the visible restriction of u' is a prefix $t' \sqsubseteq tn_i$, so in particular $t' \in \Downarrow\text{-}\mathcal{L}(\tau \circ \sigma)$. We distinguish cases on Lemma 11.1.7 applied to u' .

For (1) (state OOO), m_1 and m_2 must both be the next negative event appearing in t , so $m_1 = m_2$, contradiction. For (2) (state OPP), m_1 and m_2 both have polarity \neq and we have $u'_\tau \in \Downarrow\text{-}\mathcal{L}(\tau)$, $u'_\tau(m_1)_\tau^+, u'_\tau(m_2)_\tau^+ \in \Downarrow\text{-}\mathcal{L}(\tau)$. 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. For (3) (state POP), the reasoning is symmetric. □

By the state diagram, in an alternating interaction, only one agent has control at any point. So any alternatingly reachable non-deterministic choice in $\tau \circ \sigma$ can be attributed

to σ and τ . Since they are both deterministic, no such choice exists. We single out the linearizations of an interaction $u \in \mathcal{L}(\tau \otimes \sigma)$ that satisfy the state diagram of interactions as in Lemma 11.1.8: the *alternating interactions*.

Definition 11.1.10. Consider $u \in \mathcal{L}(\tau \otimes \sigma)$. It is an **alternating interaction**, written $u \in \Downarrow\text{-}\mathcal{L}(\tau \otimes \sigma)$, if $u_\circ \in \Downarrow\text{-}\mathcal{L}(\tau \odot \sigma)$, $u_\sigma \in \Downarrow\text{-}\mathcal{L}(\sigma)$ and $u_\tau \in \Downarrow\text{-}\mathcal{L}(\tau)$.

The notation $\Downarrow\text{-}\mathcal{L}(\tau \otimes \sigma)$ should hopefully not cause confusion, since $\tau \otimes \sigma$ does not have polarities: the polarity of synchronized events is indetermined.

A fundamental property is that an alternating linearization between sequential strategies is necessarily witnessed by an alternating interaction:

Lemma 11.1.11. Consider $\sigma : A \vdash B$ and $\tau : B \vdash C$ sequential, and $u \in \mathcal{L}(\tau \otimes \sigma)$. If $u_\circ \in \Downarrow\text{-}\mathcal{L}(\tau \odot \sigma)$, then $u \in \Downarrow\text{-}\mathcal{L}(\tau \otimes \sigma)$.

Proof. By induction on u . Consider $um \in \mathcal{L}(\tau \otimes \sigma)$ such that $(um)_\circ \in \Downarrow\text{-}\mathcal{L}(\tau \odot \sigma)$. By induction hypothesis, $u \in \Downarrow\text{-}\mathcal{L}(\tau \otimes \sigma)$, so it must be in one of the states of Lemma 11.1.7. For (1) (state OOO). Then necessarily m is negative, and it is immediate that the invariant is preserved. For (2) (state OPP), then we reason by cases depending on when m occurs. If it occurs in A , then it must be positive in $\tau \odot \sigma$ since $(um)_\circ \in \Downarrow\text{-}\mathcal{L}(\tau \odot \sigma)$. Then m_σ is defined and positive. So $u_\sigma \in \Downarrow\text{-}\mathcal{L}(\sigma)$ and $u_\sigma m_\sigma^+ \in \mathcal{L}(\sigma)$, thus by Lemma 11.1.5 we have $u_\sigma m_\sigma^+ \in \Downarrow\text{-}\mathcal{L}(\sigma)$ as well, contradicting that u_σ is in state O. Thus m occurs in B or C , so that m_τ is defined. But then $u_\tau \in \Downarrow\text{-}\mathcal{L}(\tau)$ and $u_\tau m_\tau \in \mathcal{L}(\tau)$, thus by Lemma 11.1.5 we have $u_\tau m_\tau \in \Downarrow\text{-}\mathcal{L}(\tau)$, thus m_τ is positive. If m occurs in B then the invariant is preserved and um is in state POP, if m occurs in C then the invariant is preserved and um is in state OOO. Finally, (3) (state POP) is symmetric. \square

Composition of sequential visibility. Next, we must prove that the composition of sequential strategies satisfies *sequential visibility*. This is a direct variation of the stability under composition of P-visible strategies in alternating game semantics.

As announced, we generalize P-views to alternating interactions:

Definition 11.1.12. For $u \in \Downarrow\text{-}\mathcal{L}(\tau \otimes \sigma)$, we define the subsequence $\ulcorner u \urcorner$ with

$$\begin{aligned} \ulcorner \varepsilon \urcorner &= \varepsilon \\ \ulcorner um \urcorner &= \ulcorner u \urcorner m && (\text{if } m \text{ is positive or synchronized}) \\ \ulcorner un^+ u' m^- \urcorner &= \ulcorner u \urcorner n^+ m' && (\text{if } m \text{ is justified by } n) \\ \ulcorner um^- \urcorner &= m^- && (m^- \text{ minimal}) \end{aligned}$$

Those are the same clauses as in Definition 11.1.3, applying to interactions rather than alternating linearizations of a causal strategy. The idea is the same: we browse the sequence, going back in time and following pointers of moves by the external Opponent.

The main argument behind the composition of sequential visibility is:

Lemma 11.1.13. Consider $u \in \Downarrow\text{-}\mathcal{L}(\tau \otimes \sigma)$.

Writing $\ulcorner u \urcorner = u_1 \dots u_n$ and for $1 \leq j \leq n$, if u_j is non-minimal in $\tau \otimes \sigma$ then there is $1 \leq i < j$ such that $j(u_j) = u_i$ – where $j(m)$ refers to Definition 10.2.3.

Proof. We show by mutual induction the following properties, for $u \in \downarrow\mathcal{L}(\tau \otimes \sigma)$:

- (1) writing $\ulcorner u \urcorner = u_1 \dots u_n$ and for $1 \leq j \leq n$, if u_j is non-minimal in $\tau \otimes \sigma$ then there is $1 \leq i < j$ such that $j(u_i) = u_i$,
- (2) if $u = vm$ with m_σ defined, $\ulcorner u_\sigma \urcorner$ is a subsequence of $\ulcorner u \urcorner_\sigma$,
- (3) if $u = vm$ with m_τ is defined, $\ulcorner u_\tau \urcorner$ is a subsequence of $\ulcorner u \urcorner_\tau$,

and we refer to (1) as $\ulcorner u \urcorner$ being *well-justified*.

For (1), write $u = vn$. If n is negative, then either it is minimal (then the property is trivial), or $u = v_1 l^+ v_2 n^-$ where $j(n^-) = l^+$ (say *w.l.o.g.* l is positive for σ). Then by induction hypothesis, $\ulcorner v_1 l^+ \urcorner$ is well-justified and as $\ulcorner v_1 l^+ v_2 n^- \urcorner = \ulcorner v_1 l^+ \urcorner n^-$ with $j(n) = l$, $\ulcorner u \urcorner$ is well-justified as well. Otherwise, $\ulcorner u \urcorner = \ulcorner v \urcorner n$ and by induction hypothesis, $\ulcorner v \urcorner$ is well-justified, but we must show that the justifier of n is in $\ulcorner v \urcorner$. Necessarily n is positive for σ or τ , say σ *w.l.o.g.*. As $(vn)_\sigma = v_\sigma n_\sigma^+ \in \downarrow\mathcal{L}(\sigma)$, as σ is sequentially visible, n_σ must point within $\ulcorner v_\sigma \urcorner$. Now by Lemma 11.1.8, since n is positive for σ , v is in state POP. If the last move of v occurs in C , since v_\odot has state P it must be negative. But since v_τ has state O it must be Player – contradiction, so the last move of v occurs in A or B and its projection to σ is defined. Thus by (2) of the induction hypothesis on v , $\ulcorner v_\sigma \urcorner$ is a subsequence of $\ulcorner v \urcorner_\sigma$. Now by sequential visibility of σ , the justifier of n_σ is in $\ulcorner v_\sigma \urcorner$, thus in $\ulcorner v \urcorner_\sigma$, hence it $\ulcorner v \urcorner$.

For (2), assume m_σ is defined. If it is positive, then by Lemma 11.1.8 v must be in state POP; write $v = v' n$. If n occurs in C , then it must be negative since v_\odot is in state P, but this contradicts that v_τ is in state O. Thus, n occurs in A or C , but then n_σ is defined (and negative). Hence, by induction hypothesis, $\ulcorner v_\sigma \urcorner$ is a subsequence of $\ulcorner v \urcorner_\sigma$. But $\ulcorner (um)_\sigma \urcorner = \ulcorner u_\sigma \urcorner m_\sigma$ and $\ulcorner u \urcorner_\sigma = \ulcorner v \urcorner_\sigma m_\sigma$, hence $\ulcorner u_\sigma \urcorner$ is a subsequence of $\ulcorner u \urcorner_\sigma$ as required. Next if m_σ is negative, write $u = v_1 n v_2 m$ with $j(m) = n$. Two subcases arise. If m occurs in A , then it is negative and $\ulcorner u \urcorner = \ulcorner v_1 \urcorner n m$. Likewise, $\ulcorner u_\sigma \urcorner = \ulcorner (v_1)_\sigma \urcorner n_\sigma^+ m_\sigma^-$, so the result follows by induction hypothesis. Finally if m occurs in B , then by (1), $\ulcorner u \urcorner$ is well-justified, hence n appears in $\ulcorner v_1 n v_2 \urcorner$, i.e. $\ulcorner v_1 n v_2 \urcorner = \ulcorner v_1 \urcorner n v'_2$ and $\ulcorner u \urcorner_\sigma = \ulcorner v_1 \urcorner_\sigma n_\sigma (v'_2)_\sigma m_\sigma$. Likewise by definition $\ulcorner u_\sigma \urcorner = \ulcorner (v_1)_\sigma \urcorner n_\sigma m_\sigma$. Finally by induction hypothesis $\ulcorner (v_1)_\sigma \urcorner$ is a subsequence of $\ulcorner v_1 \urcorner_\sigma$, hence the result follows.

For (3), it is the same reasoning as above. \square

From this, we may conclude:

Proposition 11.1.14. *Consider $\sigma : A \vdash B$ and $\tau : B \vdash C$ sequential causal strategies. Then, $\tau \otimes \sigma$ satisfies sequential visibility.*

Proof. Consider $s \in \downarrow\mathcal{L}(\tau \otimes \sigma)$, and write $\ulcorner s \urcorner = s_1 \dots s_n$, consider $1 \leq j \leq n$ with s_j non-minimal. Necessarily, there is $u \in \mathcal{L}(\tau \otimes \sigma)$ such that $u_\odot = s \in \downarrow\mathcal{L}(\tau \otimes \sigma)$. By Lemma 11.1.11, we have $u \in \downarrow\mathcal{L}(\tau \otimes \sigma)$. Write $\ulcorner u \urcorner = u_1 \dots u_m$ – from the definition it is immediate that $\ulcorner u \urcorner_\odot = \ulcorner s \urcorner$, so there is a unique j' such that $u_{j'} = s_j$ (necessarily non-minimal). By Lemma 11.1.13, there is some $1 \leq i' < j'$ such that $j(u_{i'}) = u_{i'}$. First, $u_{i'}$ may occur in B (if $u_{j'}$ is minimal in A and $u_{i'}$ minimal in B). Then by Lemma 11.1.13, there is $1 \leq k' < i'$ such that $j(u_{k'}) = u_{k'}$, necessarily minimal in C . But then $j(u_{j'}) = (u_{k'})$ for $\tau \otimes \sigma$ – but $u_{k'} = s_k$ for some $1 \leq k < j$, and the property is proved. If $u_{i'}$ does not occur in B , then it is visible and $j(u_{j'}) = u_{i'}$ in $\tau \otimes \sigma$ as well. As above $u_{i'} = s_i$ for some $1 \leq i < j$, which concludes the proof. \square

Putting everything together, we have:

Corollary 11.1.15. *The \sim -category **NTCG** admits a lluf sub- \sim -category **NTCG-Seq** with morphisms restricted to causal sequential strategies.*

Proof. Firstly, copycat is sequential by Lemma 11.1.6. Consider $\sigma \in \mathbf{NTCG}(A, B)$ and $\tau \in \mathbf{NTCG}(B, C)$ sequential. Firstly, the proof that $\tau \odot \sigma \in \mathbf{NTCG}(A, C)$ is *pointed* is as in Proposition 10.2.6 (which did not rely on visibility). Composition of *sequential determinism* and *sequential visibility* are Propositions 11.1.9 and 11.1.14. \square

Next, we move to the rest of the categorical structure.

11.1.3 A Relative Seely \sim -Category

All structural morphisms of **NTCG** are obtained by lifting, and it is immediate (as for Lemma 11.1.6) that they are sequential. The next step is to prove that other operations on strategies (tensor, pairing, currying and promotion) preserve sequentiality. We shall consider them all, focusing on the two non-trivial cases: tensor and promotion.

Tensor. Consider $\sigma : A \vdash B$ and $\tau : C \vdash D$. Similarly to interactions, if $m \in \sigma \otimes \tau$, write $m_\sigma \in \sigma$ as m' if $m = (1, m')$ and undefined otherwise; and $m_\tau \in \tau$ partially defined symmetrically. For $s \in \mathcal{L}(\sigma \otimes \tau)$, write $s_\sigma \in \mathcal{L}(\sigma)$ and $s_\tau \in \mathcal{L}(\tau)$ defined in the obvious way. Using these projections, the stability of sequentiality under tensor exploits a state case analysis, similar to Lemma 11.1.8:

Lemma 11.1.16. *Take A, B, C, D --boards, $\sigma : A \vdash B$, $\tau : C \vdash D$ sequential.*

If $s \in \Downarrow\text{-}\mathcal{L}(\sigma \otimes \tau)$, then $s_\sigma \in \Downarrow\text{-}\mathcal{L}(\sigma)$, $s_\tau \in \Downarrow\text{-}\mathcal{L}(\tau)$. Moreover, 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 ,

must hold.

Proof. Straightforward by induction on s , using reachable sequentiality of σ and τ . \square

Thus, in a tensor of sequential strategies $\sigma \otimes \tau$, at most one of σ and τ has control at a given point in time, and only Opponent is able to switch. From this, we may prove:

Proposition 11.1.17. *Take A, B, C, D --boards, $\sigma : A \vdash B$, $\tau : C \vdash D$ sequential.*

Then, $\sigma \otimes \tau$ is sequential.

Proof. Pointed: this is straightforward, as the event structure $\sigma \otimes \tau$ is simply $\sigma \parallel \tau$.

Sequential determinism: Consider $sm^+, sn^+ \in \Downarrow\text{-}\mathcal{L}(\sigma \otimes \tau)$. By Lemma 11.1.16, s_σ, s_τ, s must be in one of the states OOO , OPP , or POP . If it is OOO , then s is in state O , contradicting $sm^+ \in \Downarrow\text{-}\mathcal{L}(\sigma \otimes \tau)$. If it is OPP , then m_τ^+, n_τ^+ must be defined, or contradict that s_σ is in state O – but then, $(sm^+)_\tau = s_\tau m_\tau^+ \in \mathcal{L}(\tau)$, and $(sn)_\tau = s_\tau n_\tau^+ \in \mathcal{L}(\tau)$. As $s_\tau \in \Downarrow\text{-}\mathcal{L}(\tau)$ is in state P , thus $s_\tau m_\tau^+, s_\tau n_\tau^+ \in \Downarrow\text{-}\mathcal{L}(\tau)$, hence $m_\tau^+ = n_\tau^+$ by *sequential determinism* of τ . Hence, $m^+ = n^+$ as required. Finally, if s is in state POP ,

the reasoning is symmetric.

Sequential visibility: Consider $s \in \Downarrow\text{-}\mathcal{L}(\sigma \otimes \tau)$. By Lemma 11.1.16, we have $s_\sigma \in \Downarrow\text{-}\mathcal{L}(\sigma)$ and $s_\tau \in \Downarrow\text{-}\mathcal{L}(\tau)$. Moreover, it is immediate from Lemma 11.1.16 that only Opponent may switch between σ and τ , and consequently $\ulcorner s \urcorner$ is entirely in σ or in τ . Hence, sequential visibility of σ and τ allow us to conclude. \square

Promotion. The proof for promotion is analogous. Consider A and B strict $--$ -boards, and $\sigma : !A \vdash B$ sequential. Recall that every move $m \in \sigma^!$ has the form $m = (i, m')$ where $i \in \mathbb{N}$ is its *copy index*. For $j \in \mathbb{N}$ we then partially define $m_j = m'$ if $i = j$ and undefined otherwise. Again, if $s \in \mathcal{L}(\sigma \otimes \tau)$, then its projection s_i for $i \in \mathbb{N}$ is defined in the obvious way. Using these, we state the promotion state case analysis:

Lemma 11.1.18. *Consider A, B strict $--$ -boards, and $\sigma : !A \vdash B$ sequential.*

For any $s \in \Downarrow\text{-}\mathcal{L}(\sigma^!)$, for any $i \in \mathbb{N}$, $s_i \in \Downarrow\text{-}\mathcal{L}(\sigma)$ 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 . \square

As for tensor, we conclude:

Proposition 11.1.19. *Consider A, B strict $--$ -boards, and $\sigma : !A \vdash B$ sequential.*

Then, $\sigma^! : !A \vdash !B$ is a sequential strategy.

Proof. Similar to the proof of Proposition 11.1.17. \square

Remainer of the structure. That pairing preserves sequentiality is straightforward; so is currying as the internal event structure is unchanged (along with the justifier function $j(-)$). Altogether, we have the ingredients to conclude the relative Seely structure:

Theorem 11.1.20. *There is **NTCG-Seq**, a relative Seely sub- \sim -category of **NTCG**, with morphisms restricted to sequential causal strategies.*

Proof. Follows from Corollary 11.1.15, Propositions 11.1.17 and 11.1.19, and omitted direct verifications for pairing, currying and structural isomorphisms. \square

11.1.4 \rightarrow -Seq and Interpretation of IA

As a relative Seely \sim -category, **NTCG-Seq** induces a cartesian closed \sim -category:

Definition 11.1.21. *There is \rightarrow -Seq, a cartesian closed sub- \sim -category of \rightarrow -Strat.*

Its objects are mixed boards, its morphisms are those of \rightarrow -Strat restricted to sequential strategies, and the equivalence relation is as in \rightarrow -Strat.

For stateless primitives (i.e. PCF primitives, reference and semaphore queries, bad variables and semaphores), it is direct by inspection to show that the corresponding strategies are sequential – we omit those verifications. Finally, we also have:

Proposition 11.1.22. *Consider $\sigma \in \rightarrow\text{-Seq}(\mathbf{V}, A)$ sequential with A well-opened. Then, $\sigma \odot \text{cell} \in \rightarrow\text{-Strat}(\mathbf{T}, A)$ is also sequential.*

Proof. Although cell is not a strategy (it fails *courtesy*), it does satisfy conditions *sequential determinism* and *sequential visibility* (the latter involves the notion of *justifier* which is defined in Definition 10.2.1 only for *pointed* strategies, but this is needed only for strategies from A to B – it is not necessary for $\text{cell} : \mathbf{V}$). It follows from the same arguments as in Proposition 11.1.9 and 11.1.14 that $\sigma \odot \text{cell}$ satisfies these two conditions – it is also obvious from A well-opened that $\sigma \odot \text{cell}$ is pointed. \square

The exact same reasoning establishes that for $\sigma \in \rightarrow\text{-Seq}(\mathbf{S}, A)$ with A well-opened, $\sigma \odot \text{lock} : A$ is sequential as well. Altogether we have the ingredients for:

Corollary 11.1.23. *The interpretation of IA_{\parallel} in $\rightarrow\text{-Strat}$ of Chapter 9 restricts to an adequate interpretation of IA in $\rightarrow\text{-Seq}$.*

Proof. That the interpretation of IA terms yields sequential strategies follows from Theorem 11.1.20, verifications for basic primitives, Proposition 11.1.22 for new references, and the analogous for new semaphores. Adequacy is an immediate consequence of adequacy of the adequacy of the interpretation of IA_{\parallel} in $\text{NTCG}_!$ (Theorem 9.1.5). \square

11.2 The Alternating Unfolding

Now that we have constructed $\rightarrow\text{-Seq}$ and proved that the interpretation of IA indeed yields sequential causal strategies, we examine the unfolding of $\rightarrow\text{-Seq}$ onto $\Downarrow\text{-Strat}$.

11.2.1 Alternating Unfolding on a Mixed Board

The alternating unfolding of a sequential strategy is a simple refinement of the non-alternating unfolding of Section 9.2 – recall that for causal $\sigma : A$ on mixed board A , its *unfolding* was obtained by first considering all linearizations of σ , then displaying these linearizations to the game, and finally recovering plays with pointers by pointifixion.

We now adapt this to the alternating case.

Alternating linearizations. Recall from Definition 11.1.2 the set $\Downarrow\text{-}\mathcal{L}(E)$ of the alternating linearizations of an event structure with polarities E , which applies to a causal strategy $\sigma : A$ thanks to the polarities inherited from A . Our first observation is:

Lemma 11.2.1. *Consider $\sigma : A$ a causal strategy on a board A . The length-preserving monotone function of Lemma 9.2.4 restricts to*

$$\begin{aligned} \partial_{\sigma} : \Downarrow\text{-}\mathcal{L}(\sigma) &\rightarrow \Downarrow\text{-}\mathcal{L}(A) = \Downarrow\text{-}\text{Plays}(A) \\ \varepsilon &\mapsto \varepsilon \\ sm &\mapsto \partial_{\sigma}(s)\partial_{\sigma}(m). \end{aligned}$$

Proof. Straightforward. \square

We recall our convention in such a situation: alternating linearizations in σ are often denoted as $s^\sigma \in \Downarrow\text{-}\mathcal{L}(\sigma)$ with the causal strategy added as superscript, and their display is $s_A^\sigma = \partial_\sigma(s^\sigma) \in \Downarrow\text{-}\text{Plays}(A)$. As in the non-alternating case, this displayed linearization enjoys the main properties expected of what would be a play-based alternating strategy on A – to prove that, the following lemma shall be useful:

Lemma 11.2.2. *Consider A a board, and $\sigma : A$ a sequential causal strategy.*

For any $s \in \Downarrow\text{-}\text{Plays}(\sigma)$, there is a unique $t \in \Downarrow\text{-}\mathcal{L}(\sigma)$ such that $s = \partial_\sigma(t)$.

Proof. Immediate by induction, using *receptivity* and *sequential determinism*. \square

We obtain a strategy in the sense of *simple games* (Definition 3.1.2):

Proposition 11.2.3. *Consider $\sigma : A$ a sequential strategy on board A . Then, the set*

$$\Downarrow\text{-}\text{Plays}(\sigma) = \partial_\sigma(\Downarrow\text{-}\mathcal{L}(\sigma))$$

satisfies the following properties:

- non-empty: $\varepsilon \in \Downarrow\text{-}\text{Plays}(\sigma)$,
- prefix-closed: *for all $t \in \Downarrow\text{-}\text{Plays}(\sigma)$, is $s \sqsubseteq t$, then $s \in \Downarrow\text{-}\text{Plays}(\sigma)$,*
- receptive: *if $s \in \Downarrow\text{-}\text{Plays}(\sigma)$ and $sa^- \in \Downarrow\text{-}\text{Plays}(A)$, then $sa \in \Downarrow\text{-}\text{Plays}(\sigma)$,*
- deterministic: *if $sa^+, sb^+ \in \Downarrow\text{-}\text{Plays}(\sigma)$, then $a^+ = b^+$,*

Proof. *Non-empty, prefix-closed, receptive.* Same proof as for Proposition 9.2.5.

Deterministic. Immediate by Lemma 11.2.2 and sequential determinism. \square

With the appropriate phrasing it also follows that $\Downarrow\text{-}\text{Plays}(\sigma)$ is P -visible, but we have not defined the P -view of alternating plays on A and it does not seem worth the detour. Thus, following the non-alternating case, we shall now apply the above to a sequential strategy on a mixed board, and follow it up with pointifixion.

Alternating linearizations up to symmetry. The above proposition does not say anything about symmetry. Recall that we proved in Proposition 7.1.15 that for A a gws (hence, in particular, for A a board), the set of plays $\Downarrow\text{-}\text{Plays}(A)$ equipped with the equivalence relation \cong_A (from Definition 7.1.14) is automatically an AJM game. Accordingly, we expect that $\Downarrow\text{-}\text{Plays}(\sigma)$ is *uniform* in the sense of AJM games – this is natural as a consistency check, but it shall also be required later on.

To prove this, the first step is a symmetry-aware version of Lemma 11.2.2. Here for $s^\sigma, t^\sigma \in \Downarrow\text{-}\mathcal{L}(\sigma)$, we write $s^\sigma \cong_\sigma t^\sigma$ for the equivalence relation defined in exactly the same way as the AJM equivalence on plays on gws (Definition 7.1.14).

Lemma 11.2.4. *Consider A a --board, $\sigma : A$ a sequential causal strategy.*

For all $s^\sigma, t^\sigma \in \Downarrow\text{-}\mathcal{L}(\sigma)$, if $s_A^\sigma \cong_A t_A^\sigma$ then $s^\sigma \cong_\sigma t^\sigma$.

Proof. By induction on the length of linearizations – note from $s_A^\sigma \cong_A t_A^\sigma$ that s^σ and t^σ have the same length. If they are both empty then there is nothing to prove.

Consider $s^\sigma m^- \in \Downarrow\text{-}\mathcal{L}(\sigma)$ displaying to $s_A^\sigma a^- \in \Downarrow\text{-}\text{Plays}(\sigma)$, and $t^\sigma n^- \in \Downarrow\text{-}\mathcal{L}(\sigma)$ displaying to $t_A^\sigma b^- \in \Downarrow\text{-}\text{Plays}(\sigma)$, s.t. $s_A^\sigma a^- \cong_A t_A^\sigma b^-$. By IH, we have $s^\sigma \cong_\sigma t^\sigma$. But then it immediately follows by \sim -receptivity of σ that $s^\sigma m^- \cong_\sigma t^\sigma n^-$ as required.

Now, consider $s^\sigma m^+ \in \Downarrow\text{-}\mathcal{L}(\sigma)$ displaying to $s_A^\sigma a^+ \in \Downarrow\text{-}\text{Plays}(\sigma)$, and $t^\sigma n^+ \in \Downarrow\text{-}\mathcal{L}(\sigma)$ displaying to $t_A^\sigma b^+ \in \Downarrow\text{-}\text{Plays}(\sigma)$, such that $s_A^\sigma a^+ \cong_A t_A^\sigma b^+$. By induction hypothesis, we have $s^\sigma \cong_\sigma t^\sigma$. Now, by *extension* for \cong_σ , we must have $t^\sigma p^+ \in \Downarrow\text{-}\mathcal{L}(\sigma)$ such that $s^\sigma m^+ \cong_\sigma t^\sigma p^+$. But by *sequential determinism*, it follows that $n^+ = p^+$, hence we do have $s^\sigma m^+ \cong_\sigma t^\sigma n^+$ as required. \square

The following lemma has two purposes: it will ensure that strategies obtained as displayed linearizations are indeed uniform in the sense of AJM strategies, and that two positively isomorphic sequential strategies will yield equivalent AJM strategies.

Lemma 11.2.5. *Consider A a \leftarrow -board and $\sigma, \tau : A$ sequential causal strategies.*

If $\sigma \approx \tau$ (as in Section 7.2.1), then $\Downarrow\text{-}\text{Plays}(\sigma) \approx \Downarrow\text{-}\text{Plays}(\tau)$ (as in Definition 3.4.4).

Proof. If $\sigma \approx \tau$, by Definition 7.2.4 there is an iso $\varphi : \sigma \cong \tau$ of ess satisfying

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

which in particular entails that for all $s^\sigma \in \Downarrow\text{-}\mathcal{L}(\sigma)$, we have $s_A^\sigma \cong_A \varphi(s^\sigma)_A$. Combined with Lemma 11.2.4, this has an important direct consequence: if $s^\sigma \in \Downarrow\text{-}\mathcal{L}(\sigma)$ and $s^\tau \in \Downarrow\text{-}\mathcal{L}(\tau)$ are such that $s_A^\sigma \cong_A s_A^\tau$, then it also follows that $\varphi(s^\sigma)_A \cong_A s_A^\tau$, and consequently $\varphi(s^\sigma) \cong_\tau s^\tau$ by Lemma 11.2.4 – we shall use that in the proof below.

\rightarrow -simulation: consider $s_A^\sigma a^+ \in \Downarrow\text{-}\text{Plays}(\sigma)$, $s_A^\tau \in \Downarrow\text{-}\text{Plays}(\tau)$ such that $s_A^\sigma \cong_A s_A^\tau$. By our auxiliary statement, we necessarily have $\varphi(s^\sigma) \cong_\tau s^\tau$. Now writing $s^\sigma m^+ \in \Downarrow\text{-}\mathcal{L}(\sigma)$ the witness for $s_A^\sigma a^+ \in \Downarrow\text{-}\text{Plays}(\sigma)$, we have $\varphi(s^\sigma)\varphi(m^+) \in \Downarrow\text{-}\mathcal{L}(\tau)$ with $\varphi(s^\sigma) \cong_\tau s^\tau$, hence by *extension* for τ there is n^+ such that $\varphi(s^\sigma)\varphi(m^+) \cong_\tau s^\tau n^+$, with

$$s_A^\sigma a^+ \cong_A \partial_\sigma(s^\sigma m^+) \cong_A \partial_\tau(s^\tau n^+) = s_A^\tau b^+ \in \Downarrow\text{-}\text{Plays}(\tau)$$

as required. The case for \leftarrow -simulation is symmetric.

\rightarrow -receptive: consider $s_A^\sigma a^- \in \Downarrow\text{-}\text{Plays}(\sigma)$, $s_A^\tau \in \Downarrow\text{-}\text{Plays}(\tau)$, and $s_A^\sigma a^- \cong_A s_A^\tau b^-$. By the observation above, $\varphi(s^\sigma) \cong_\tau s^\tau$. Write $s^\sigma m^- \in \Downarrow\text{-}\mathcal{L}(\sigma)$ such that $\partial_\sigma(s^\sigma m^-) = s_A^\sigma a^-$, and write $\varphi(m^-) = n^-$. Then $\varphi(s^\sigma)n^- \in \Downarrow\text{-}\mathcal{L}(\tau)$ displayed to $\varphi(s^\sigma)_A c^-$, and

$$\varphi(s^\sigma)_A c^- \cong_A s_A^\sigma a^- \cong_A s_A^\tau b^-,$$

where $\varphi(s^\sigma) \cong_\tau s^\tau$, hence by \sim -receptivity of τ there must be a unique extension

$$\varphi(s^\sigma)n^- \cong_\tau s^\tau p^-$$

such that $\partial_\tau(s^\tau p^-) = s_A^\tau b^-$. Finally, \leftarrow -receptive is symmetric. \square

In particular, we obtain as a direct corollary:

Corollary 11.2.6. *If $\sigma : A$ is a sequential strategy, then $\Downarrow\text{-}\text{Plays}(\sigma)$ is an AJM strategy.*

Proof. From Proposition 11.2.3 we lack *uniformity*, obvious from Lemma 11.2.5. \square

However, our goal is to unfold to strategies in HO games rather than AJM games.

Pointifixion of Alternating Plays From a sequential causal strategy, we wish to obtain an alternating strategy in the sense of Chapter 3; but those play on *arenas* rather than *boards*. Hence the situation is exactly the same as in Section 9.2: it is unclear how to automatically obtain an arena in the sense of Chapter 3 from a board, hence we unfold sequential causal strategies playing on *mixed boards* in the sense of Definition 9.2.6 – comprising the data of the board, the arena, along with the link between the two.

As in Corollary 9.2.11 for the non-alternating case, we have:

Proposition 11.2.7. *Consider A a mixed board. Then, $(\hat{\cdot})$ restricts to bijections*

$$\begin{aligned} (\hat{\cdot}) &: \Downarrow\text{-Plays}(!A)/\cong_{!A} \simeq \Downarrow\text{-Plays}(\underline{A}), \\ (\hat{\cdot}) &: \Downarrow\text{-Plays}(A)/\cong_A \simeq \Downarrow\text{-Plays}_*(\underline{A}), \end{aligned}$$

preserving length and prefix.

Proof. Simply the restriction of Corollary 9.2.11 to alternating plays. \square

Corollary 9.2.11 stated a similar bijection between $\Downarrow\text{-Plays}(!A)$ up to symmetry and *pre-plays* $\Downarrow\text{-PrePlays}(\underline{A})$. The mismatch with the above is because while this bijection does not account for any well-bracketing constraint, non-alternating plays with pointers were defined as always logically well-bracketed (Definition 5.1.3) whereas alternating plays have no such hardwired condition³ (Definition 3.1.4).

11.2.2 Alternating Unfoldings of Sequential Strategies

Now, we are in position to define the *alternating unfolding* of a sequential causal strategy. As for non-alternating strategies, we shall actually define *three* unfoldings:

- (1) $\Downarrow\text{-Unf}_\bullet$ unfolds $\sigma : A$ to an alternating thread-strategy on \underline{A} ,
- (2) $\Downarrow\text{-Unf}_!$ unfolds $\sigma : !A$ to an alternating strategy on \underline{A} ,
- (3) $\Downarrow\text{-Unf}$ unfolds $\sigma : !A \vdash B$ to an alternating strategy on $\underline{A} \Rightarrow \underline{B}$,

for all mixed boards A and B .

Unfolding to a thread-strategy. First, we define the pointed alternating unfolding of a sequential causal strategy on a mixed board A . The definition is essentially the same as Definition 9.2.12 for the non-alternating unfolding of causal strategies:

Definition 11.2.8. *Consider A a mixed board, and $\sigma : A$ a sequential causal strategy. The *unfolding* of σ is the set $\Downarrow\text{-Unf}_\bullet(\sigma) = \{(\hat{\cdot})(s_A^\sigma) \in \Downarrow\text{-Plays}(\underline{A}) \mid s^\sigma \in \mathcal{L}(\sigma)\}$.*

³In general, we consider it better to impose conditions on strategies rather than on plays whenever possible, so as to not unnecessarily constrain the behaviour of the execution environment. But for non-alternating strategies we were brought to assume all plays to be logically well-bracketing, as the non-alternating strategies for basic $\text{IA}_{//}$ primitives are hard to define rigorously otherwise: they have too many plays!

We show that this is indeed a valid alternating thread-strategy. With respect to Proposition 9.2.13 in the non-alternating case, the difficulty is to prove determinism: this puts together (1) sequential determinism, (2) Proposition 11.2.7 ensuring that alternating plays with pointers on \underline{A} exactly represent alternating plays on A up to symmetry, and (3) Lemma 11.2.4 showing uniqueness of the witness in $\Downarrow\text{-}\mathcal{L}(\sigma)$ for a plays in $\Downarrow\text{-}\text{Plays}(\sigma)$, up to symmetry. Putting those together, we get the proposition:

Proposition 11.2.9. *Consider $\sigma : A$ a causal strategy on mixed board A .*

Then, $\Downarrow\text{-}\text{Unf}_!(\sigma) : \underline{A}$ is a P -visible thread-strategy in the sense of Definition 4.3.2.

Proof. It is a set of threads since A is strict. Non-empty and prefix-closed are direct.

Receptive: consider $t = (\hat{\cdot})(s_A^\sigma)$ for $s^\sigma \in \Downarrow\text{-}\mathcal{L}(\sigma)$, with $tb^- \in \Downarrow\text{-}\text{Plays}_!(\underline{A})$. By Lemma 9.2.10, there is $s_A^\sigma a^- \in \Downarrow\text{-}\text{Plays}(A)$ such that $(\hat{\cdot})(s_A^\sigma a^-) = tb^-$; but by construction $s_A^\sigma a^-$ is alternating, thus in $\Downarrow\text{-}\text{Plays}(A)$. Now, by Proposition 11.2.3, there is $s^\sigma m \in \Downarrow\text{-}\mathcal{L}(\sigma)$ displaying to $s_A^\sigma a^-$, and *receptive* follows.

Deterministic: Consider $sa^-b_1^+, sa^-b_2^+ \in \Downarrow\text{-}\text{Unf}_!(\sigma)$, so there are $t^\sigma m^\sigma n^\sigma \in \Downarrow\text{-}\mathcal{L}(\sigma)$ and $u^\sigma c^\sigma d^\sigma \in \Downarrow\text{-}\mathcal{L}(\sigma)$ s.t. $sa^-b_1^+ = (\hat{\cdot})(t_A^\sigma m_A^\sigma n_A^\sigma)$ and $sa^-b_2^+ = (\hat{\cdot})(u_A^\sigma c_A^\sigma d_A^\sigma)$. Thus

$$t^\sigma m_A^\sigma n_A^\sigma \in \Downarrow\text{-}\text{Plays}(\sigma), \quad u^\sigma c_A^\sigma d_A^\sigma \in \Downarrow\text{-}\text{Plays}(\sigma),$$

with $(\hat{\cdot})(t_A^\sigma m_A^\sigma) = (\hat{\cdot})(u_A^\sigma c_A^\sigma) = sa^-$, which entails $t_A^\sigma m_A^\sigma \cong_A u_A^\sigma c_A^\sigma$ by Proposition 11.2.7. But then, by Lemma 11.2.4, we must actually have

$$t^\sigma m^\sigma \cong_\sigma u^\sigma c^\sigma,$$

but hence by extension, there is $t^\sigma m^\sigma n^\sigma \cong_\sigma u^\sigma c^\sigma e^\sigma$. But by *sequential determinism* for σ , this means $e^\sigma = d^\sigma$, so that $t^\sigma m^\sigma n^\sigma \cong_\sigma u^\sigma c^\sigma d^\sigma$. Thus their display are also symmetric, so their pointifixion must be equal by Proposition 11.2.7.

P-visible: Straightforward by definition of sequential visibility. \square

Unfolding to an alternating strategy. Next, we show how to unfold a sequential causal strategy to a full alternating strategy – again the definition closely mimics the non-alternating case, more specifically Definition 9.2.14

Definition 11.2.10. *Consider A a mixed board, and $\sigma : !A$ a sequential causal strategy.*

*The **unfolding** of σ is the set $\Downarrow\text{-}\text{Unf}_!(\sigma) = \{(\hat{\cdot})(s_A^\sigma) \mid s^\sigma \in \Downarrow\text{-}\mathcal{L}(\sigma)\}$.*

This is a well-formed alternating strategy on arena \underline{A} :

Proposition 11.2.11. *Consider A a mixed board and $\sigma : !A$ a sequential strategy.*

Then, $\Downarrow\text{-}\text{Unf}_!(\sigma) : \underline{A}$ is a P -visible alternating strategy as in Definition 3.1.5.

Proof. Same proof as for Proposition 11.2.9. \square

As in the non-alternating case, if $\sigma : A$ is a sequential causal strategy, we may unfold it to a thread-strategy $\Downarrow\text{-}\text{Unf}_!(\sigma) : \underline{A}$ and then obtain the corresponding single-threaded alternating strategy $\sigma^!$, or we may directly unfold its promotion as an alternating strategy with $\Downarrow\text{-}\text{Unf}_!(\sigma^!)$ – the two options coincide, as we show now.

Lemma 11.2.12. *Consider sequential causal $\sigma : A$, and $\sigma^! : !A$ its promotion.*

Then, $(\Downarrow\text{-Unf}_(\sigma))^! = \Downarrow\text{-Unf}_!(\sigma^!)$.*

Proof. \subseteq . We actually show by induction on n that for all $s = s_1 \dots s_n \in (\Downarrow\text{-Unf}_*(\sigma))^!$, there is $t = t_1 \dots t_n \in \Downarrow\text{-}\mathcal{L}(\sigma^!)$ such that $s = (\hat{\cdot})(t_{!A})$ and for all prefix $s_1 \dots s_k$, writing

$$s_{\pi(1)} \dots s_{\pi(k)} = [s_1 \dots s_k]$$

the current thread of $s_1 \dots s_k$ for $\pi : \{1, \dots, p\} \rightarrow \{1, \dots, k\}$ some injection, we have $t_{\pi(1), \dots, \pi(p)} \in \Downarrow\text{-}\mathcal{L}(\sigma)$ and $s_{\pi(1)} \dots s_{\pi(k)} = (\hat{\cdot})(\partial_\sigma(t_{\pi(1)} \dots t_{\pi(p)}))$; the proof by induction is straightforward, relying on Lemma 9.2.10 and receptivity for negative extensions, on Lemma 11.2.4 and *extension* of isomorphism families for positive extensions.

\supseteq . Necessarily, any $sa^+ \in \Downarrow\text{-Unf}_!(\sigma^!)$ comes from some $tm^+ \in \Downarrow\text{-}\mathcal{L}(\sigma^!)$, where m is in the component of $\sigma^!$ with copy index i . It is then straightforward that the moves of tm^+ with the same initial hereditary justifier as m^+ are exactly those events in copy index i , written $t'm^+ \in \Downarrow\text{-}\mathcal{L}(\sigma)$; and then we must have that

$$[sa^+] = [(\hat{\cdot})(\partial_{\sigma^!}(tm^+))] = (\hat{\cdot})(\partial_\sigma(t'm^+)) \in \Downarrow\text{-Unf}_*(\sigma)$$

as required, concluding the proof of the inclusion. \square

Unfolding from mixed board A to mixed board B . As in the non-alternating case, we finish with the definition of the unfolding of a sequential causal strategy from one mixed board to another – again this must deal with the fact that morphisms from A to B in $\rightarrow\text{-Seq}$ are causal sequential strategies on $!A \vdash B$ (with at most one initial move) whereas morphisms in $\Downarrow\text{-Strat}$ are alternating P-visible strategies on $\underline{A} \Rightarrow \underline{B}$, whose plays have arbitrarily many initial moves. Thus to deal with these differences, we set:

Definition 11.2.13. *Consider $\sigma : !A \vdash B$ a sequential strategy. Its **unfolding** is:*

$$\Downarrow\text{-Unf}(\sigma) = \Downarrow\text{-Unf}_!(\Lambda(\sigma)^!) : \underline{A} \Rightarrow \underline{B}.$$

It is a consequence of Proposition 11.2.11 that this is well-defined:

Proposition 11.2.14. *For any $\sigma \in \rightarrow\text{-Seq}(A, B)$, we have $\Downarrow\text{-Unf}(\sigma) \in \Downarrow\text{-Strat}(\underline{A}, \underline{B})$.*

Proof. In complement to Proposition 11.2.11, it remains to prove that $\Downarrow\text{-Unf}(\sigma)$ is single-threaded. But by Lemma 11.2.12, we have $\Downarrow\text{-Unf}(\sigma) = \Downarrow\text{-Unf}_*(\Lambda(\sigma))^!$ which is single-threaded by Proposition 4.3.3; this concludes the proof. \square

11.2.3 Unfolding the Basic Categorical Structure

Next, we aim to show that the unfolding defined just above yields a \sim -functor

$$\Downarrow\text{-Unf}(-) : \rightarrow\text{-Seq} \rightarrow \Downarrow\text{-Strat}$$

which we will show preserves the interpretation of IA .

Identities. First, we show that unfolding preserves identities. Of course, one must keep in mind that $\rightarrow\text{-Seq}$ is built on top of the Kleisli category \mathbf{NTCG}_1 , so that the identity to be unfolded is actually dereliction $\mathbf{der}_A : !A \vdash A$.

Proposition 11.2.15. *Consider A a mixed board. Then,*

$$\Downarrow\text{-Unf}(\mathbf{der}_A) = \Downarrow\text{-}\mathfrak{C}_A \in \Downarrow\text{-Strat}(\underline{A}, \underline{A}).$$

Proof. By Lemma 11.2.12 and Proposition 4.3.4 it suffices to show the equality

$$\Downarrow\text{-Unf}_*(\Lambda(\mathbf{der}_A)) = \Downarrow\text{-}\mathfrak{C}_A^*,$$

where $\Downarrow\text{-}\mathfrak{C}_A^*$ denotes the set of *threads* of $\Downarrow\text{-}\mathfrak{C}_A$.

For \subseteq , consider $s \in \Downarrow\text{-Unf}_*(\Lambda(\mathbf{der}_A))$, meaning that $s = \hat{\tau}$ with $t = \partial_{\Lambda(\mathbf{der}_A)}(u)$ for some $u \in \Downarrow\text{-}\mathcal{L}(\mathfrak{C}_A)$. Recalling that $|\mathfrak{C}_A| = |A_1 \vdash A_2|$ with tags used for disambiguation, we observed in the proof of Lemma 11.1.6 that alternating linearizations on \mathfrak{C}_A are exactly those $v \in |A_1 \vdash A_2|^*$ such that for all event-length $w \sqsubseteq v$, we have $w \upharpoonright A_1 = w \upharpoonright A_2$. Now any even-length $s' \sqsubseteq s$ comes from $w \sqsubseteq u$ an even-length prefix of u on which this observation applies, and it is direct that $s' \upharpoonright \underline{A}_1 = s' \upharpoonright \underline{A}_2$.

For \supseteq , consider $s \in \Downarrow\text{-}\mathfrak{C}_A^*$, assuming first that it has even length – hence $s \upharpoonright \underline{A}_1 = s \upharpoonright \underline{A}_2$. By Lemma 9.2.10, there is $t \in \Downarrow\text{-Plays}(A)$ such that $\hat{\tau} = s \upharpoonright \underline{A}_1 = s \upharpoonright \underline{A}_2$. Interleaving two copies of t , it is straightforward to get $u \in \Downarrow\text{-Plays}(A_1 \vdash A_2)$ such that $u \upharpoonright A_1 = u \upharpoonright A_2 = t$, and that for all even-length prefix $v \sqsubseteq u$, we have $v \upharpoonright A_1 = v \upharpoonright A_2$ as well, hence $u \in \Downarrow\text{-}\mathcal{L}(\mathfrak{C}_A)$ – it is then a direct verification that writing $w = \partial_{\Lambda(\mathbf{der}_A)}(u)$, we have $\hat{w} = s$ as required. Finally, if s has odd length, we use the above without the last move, and conclude by *receptivity* of $\Downarrow\text{-Unf}_*(\mathbf{der}_A)$. \square

Composition. Here, the reasoning is pretty much the same as in the non-alternating case, except that plays and witnesses have to be ensured alternating – which follows easily from sequentiality. Thus we give fewer details than in the non-alternating case:

Proposition 11.2.16. *Consider $\sigma \in \rightarrow\text{-Seq}(A, B)$ and $\tau \in \rightarrow\text{-Seq}(B, C)$. Then:*

$$\Downarrow\text{-Unf}(\tau \odot_! \sigma) = \Downarrow\text{-Unf}(\tau) \odot \Downarrow\text{-Unf}(\sigma)$$

Proof. First of all, we prove the inclusion

$$\Downarrow\text{-Unf}_*(\Lambda(\tau \odot \sigma^!)) \subseteq \Downarrow\text{-Unf}(\tau) \odot \Downarrow\text{-Unf}(\sigma),$$

this follows by the same reasoning as in Lemma 9.3.4, additionally observing by induction that the interaction we obtain is in $\Downarrow\text{-Unf}(\sigma) \parallel \Downarrow\text{-Unf}(\tau)$ (in the sense of Definition 4.1.7) exploiting that σ and τ satisfy *reachable sequentiality* (Lemma 11.1.5).

Next, we show the following inclusion:

$$\llbracket \Downarrow\text{-Unf}(\tau) \odot \Downarrow\text{-Unf}(\sigma) \rrbracket \subseteq \Downarrow\text{-Unf}_*(\Lambda(\tau \odot \sigma^!)),$$

in this case this follows exactly as in 9.3.5, as the play for which we construct a witness in $\tau \odot \sigma^!$ is by hypothesis alternating.

From these two inclusions, the desired equality follows as in Corollary 9.3.6. \square

Altogether, we have proved:

Theorem 11.2.17. *Unfolding yields a \sim -functor $\Downarrow\text{-Unf} : \rightarrow\text{-Seq} \rightarrow \Downarrow\text{-Strat}$.*

Proof. The only thing remaining to check is that $\Downarrow\text{-Unf}$ preserves equivalence, which is proved exactly as in Theorem 9.3.7. \square

11.2.4 Unfolding the Interpretation of IA

As in the non-alternating case, it is an immediate verification that the alternating unfolding preserves the cartesian closed structure: preservation of projections is a variation of Proposition 11.2.15, it is a simple verification that unfolding preserves currying and its inverse, so that evaluation is preserved as well – altogether we get

$$\Downarrow\text{-Unf}(-) : \rightarrow\text{-Seq} \rightarrow \Downarrow\text{-Strat}$$

a cartesian closed \sim -functor. The unfolding is continuous with respect to the dcpo structure on hom-sets, so that the fixpoint operator is preserved. Finally, it follows by inspection that the unfolding preserves the interpretation of all IA primitives, so that:

Theorem 11.2.18. *We have a cartesian closed \sim -functor:*

$$\Downarrow\text{-Unf}(-) : \rightarrow\text{-Seq} \rightarrow \Downarrow\text{-Strat},$$

preserving the interpretation in the sense that for $\Gamma \vdash M : A$ any term of IA, we have

$$\Downarrow\text{-Unf}(\llbracket M \rrbracket_{\rightarrow\text{-Seq}}) = \llbracket M \rrbracket_{\Downarrow\text{-Strat}}.$$

Intensional full abstraction. We motivated the unfolding

$$\Downarrow\text{-Unf}(-) : \rightarrow\text{-Seq} \rightarrow \Downarrow\text{-Strat}$$

as a way to show that $\rightarrow\text{-Seq}$ is intensionally fully abstract for IA, *i.e.* to inherit in concurrent games the full abstraction result for IA in alternating game semantics (Theorem 3.3.5). However, full abstraction does not hold for $\Downarrow\text{-Strat}$ but only with respect to *well-bracketed* P-visible alternating strategies, *i.e.* for $\Downarrow\text{-WB}$.

Strategies in $\rightarrow\text{-Seq}$ already embark two notions of well-bracketing: the winning mechanism implemented in **NTCG** (see Section 8.2) already bans call/cc, and we also imported logical well-bracketing from non-alternating games in Definition 9.3.2 – so we may hope that the unfolding actually lifts to a cartesian closed \sim -functor

$$\Downarrow\text{-Unf}(-) : \rightarrow\text{-Seq} \rightarrow \Downarrow\text{-WB},$$

but that is not the case: we showed in Figure 10.18 a sequential causal strategy whose alternating unfolding is not well-bracketed. The strategy is winning: its $+$ -covered configurations have a positive payoff. Its non-alternating unfolding is logically well-bracketed (this matches the counter-example to the fact that logical well-bracketing implies well-bracketing, seen in Figure 5.3). Nevertheless, the play following the long causal chain is not well-bracketed as the final move does not answer the pending question. However this does not prevent the interpretation of IA in $\rightarrow\text{-Seq}$ to be intensionally fully abstract, because this phenomenon cannot occur in plays where the initial move has an answer.

Lemma 11.2.19. *Let $s \in \Downarrow\text{-Plays}(\underline{A})$ be P- and O-visible. Assume that s has the form*

$$s = \dots s_i \cdots s_j \dots$$

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 detailed proof appears in [Clairambault and Harmer, 2010, Lemma 5]. \square

From this, we may easily deduce the following:

Lemma 11.2.20. *Consider $s \in \Downarrow\text{-Plays}(\underline{A})$ O- and P-visible, and such that any Question has an answer. Then, s is well-bracketed in the sense of Definition 3.2.5.*

Proof. Assume s has a well-bracketing failure, i.e. as in:

$$s = \dots \mathbf{q}_1^Q \cdots \mathbf{q}_2^Q \cdots a^A \dots$$

with \mathbf{q}_2 unanswered when playing a . By *answer-closing*, no further move can point to a . Thus by Lemma 11.2.19, no further move can point to \mathbf{q}_2 which must therefore remain unanswered, contradicting the hypothesis that every question has an answer. \square

From all this, we may now conclude:

Theorem 11.2.21. *The interpretation of IA in $\rightarrow\text{-Seq}$ is intensionally fully abstract.*

Proof. Let $\vdash M, N : A$ be terms in IA , and assume that $\llbracket M \rrbracket \neq \llbracket N \rrbracket$, i.e. there is a test $\alpha \in \rightarrow\text{-Seq}(\llbracket A \rrbracket, \llbracket \mathbb{U} \rrbracket)$ 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 $\alpha' = \Downarrow\text{-Unf}(\alpha)$, it follows that

$$\alpha' \odot \llbracket M \rrbracket_{\Downarrow\text{-Strat}} \Downarrow \quad \alpha' \odot \llbracket N \rrbracket_{\Downarrow\text{-Strat}} \Uparrow,$$

since, by Theorem 11.2.18, the alternating unfolding preserves composition and the interpretation of IA . Here, α' is a P-visible alternating strategy, but because of the phenomena highlighted in Figure 10.18 it may not be well-bracketed as in Definition 3.2.6. 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 $\Downarrow\text{-Strat}$. The initial question of s has an answer. Moreover, $\llbracket M \rrbracket$ is the alternating restriction of $\Downarrow\text{-Unf}(\llbracket M \rrbracket_{\rightarrow\text{-Strat}})$ which is logically well-bracketed, and likewise α' is logically well-bracketed, thus s is logically well-bracketed and hence all its questions are answered. It is P-visible and O-visible since both $\llbracket M \rrbracket$ and α' are P-visible. Hence, by Lemma 11.2.20, it is well-bracketed as in Definition 3.2.6.

Consider α' restricted to prefixes of s . Now α' is well-bracketed as in Definition 3.2.5, and it distinguishes $\llbracket M \rrbracket$ and $\llbracket N \rrbracket$, hence $M \neq N$ by Theorem 4.3.11. \square

So far we have seen that:

- $\rightarrow\text{-Strat}$ is intensionally fully abstract for \mathbf{IA}_{\parallel} (Theorem 9.3.10),
- $\rightarrow\text{-Seq}$ is intensionally fully abstract for \mathbf{IA} (Theorem 11.2.21).

It remains to prove two matching theorems in the presence of parallel innocence.

11.3 Sequential Innocence

In this section, we study the combination of sequentiality and parallel innocence, and establish links with traditional innocence as introduced in Section 3.2.

Let us start with the definition:

Definition 11.3.1. *A causal strategy $\sigma : A$ is **sequential innocent** if it is both sequential (as in Definition 11.1.4) and parallel innocent (as in Definition 10.1.11).*

As both sequential and parallel innocent causal strategies form a cartesian closed sub- \sim -category of $\rightarrow\text{-Strat}$, there is a cartesian closed \sim -category $\rightarrow\text{-SeqInn}$ of sequential innocent strategies. As PCF may be interpreted in both $\rightarrow\text{-Seq}$ and $\rightarrow\text{-Inn}$, it follows that it has an adequate interpretation in $\rightarrow\text{-SeqInn}$ as well. In this section we shall prove that this interpretation is intensionally fully abstract, by proving that $\Downarrow\text{-Strat}(-)$ sends sequential innocent strategies to innocent alternating strategies as in Definition 3.2.9, and rely on Theorem 4.3.10 to deduce intensional full abstraction.

11.3.1 Causal Analysis of Sequential Innocence

Basic causal shape. So as to establish the link with traditional innocence, we first aim to understand better what constraints on the causal shape are brought by the combination of parallel innocence, and sequentiality. Roughly speaking, parallel innocence only allows Player to merge threads he spawned, whereas sequentiality prevents Player to spawn parallel threads: it follows that the causal shape of a sequential innocent strategy must be purely forestial, with branchings the entire responsibility of Opponent:

Lemma 11.3.2. *Consider A a board, and $\sigma : A$ a sequential, parallel innocent 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 6.1.16, 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' must be the only immediate dependency of m_1 , and the only immediate dependency of 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




Figure 11.3: A sequential innocent causal strategy

m' , contradicting its minimality. Hence, $[m'] \cup \{m_i\} \in \mathcal{C}(\sigma)$ for $i \in \{1, 2\}$.

Also writing $[m']$ for the linearization in $\Downarrow\text{-}\mathcal{L}(\sigma)$ with events in the same order,

$$[m'], [m']m_1, [m']m_2 \in \mathcal{L}(\sigma),$$

but by Lemma 6.1.16, $[m']m_1^+$ 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 \bar{O} -branching. Finally, as for any causal strategy $\sigma : A$ on A alternating, we have \rightarrow_σ is alternating as well. \square


Description. The causal shape of a sequential innocent strategy is a *forest*, but it might not be obvious to the reader what this forest actually represents. Thus we start this discussion by looking at the interpretation as a causal strategy of the simple term

$$\vdash \lambda f^{\mathbb{U} \rightarrow \mathbb{U}}. f \text{ skip}; f \text{ skip}; \mathbf{tt} : (\mathbb{U} \rightarrow \mathbb{U}) \rightarrow \mathbb{B}.$$


of pure PCF. The interpretation of this term yields the strategy in Figure 11.3.

This event structure is infinite, since Opponent may play any of his available moves an arbitrary number of times, using any natural number as copy index.

Intuitively, this diagram follows the structure of the term: once computation is initiated, Player asks for the output of f . Each time f calls its argument, Player answers \checkmark^+ . Note though that unlike in traditional game semantics, the strategy shows that these calls are causally independent from each other. Likewise, the game permits Opponent to have f return multiple times. Each time this happens, this triggers another call to f , which may call its argument any number of times, and return any number of times.



It seems clear from this diagram that the strategy contains a lot of redundant information: by symmetry, it intuitively suffices to know how the strategy behaves if Opponent performs each available action exactly once, say with copy index 0:



Recovering P-views. Next we show that, as expected from the discussion above, the branches of a sequential innocent strategy do correspond to P-views.

For all $\rho \in \text{gcc}(\sigma)$, $(\hat{-})(\partial_\sigma(\rho))$ is a P -view.

Proof. By Lemma 11.3.2, σ is forestal, so that prefixes of ρ are down-closed and inform configurations. Hence, we actually have $\rho \in \mathcal{L}(\sigma)$. Moreover, by Lemma

6.1.16, immediate causality alternates in σ and in particular, $\rho \in \Downarrow\text{-}\mathcal{L}(\sigma)$. In particular, $\partial_\sigma(\rho) \in \Downarrow\text{-}\text{Plays}(A)$ and hence $(\hat{\leftarrow})(\partial_\sigma(\rho)) \in \Downarrow\text{-}\text{Plays}(\underline{A})$.

To prove that it is a P-view, it remains to show that Opponent moves point to the previous move. But if we have $s = s_1 \dots s_i^+ s_{i+1}^- \dots = (\hat{\leftarrow})(\partial_\sigma(\rho))$, those must come from

$$\rho = \rho_1 \rightarrow_\sigma \dots \rightarrow_\sigma \rho_i^+ \rightarrow_\sigma \rho_{i+1}^- \rightarrow_\sigma \dots ,$$

then $\partial_\sigma(\rho_i^+) \rightarrow_A \partial_\sigma(\rho_{i+1}^-)$ by *courtesy*, which entails that s_{i+1}^- points to s_i^+ . \square

Thus indeed, we may regard a sequential innocent strategy as a presentation of an expanded version of the forest of P-views, with an explicit choice of copy indices.

11.3.2 The Unfolding Preserves Innocence

We now aim to show that the unfolding of a sequential innocent strategy is innocent in the sense of traditional Hyland-Ong game semantics, as introduced in Chapter 3.

Alternating linearizations. In order to prove that, our first step will be to study alternating linearizations of sequential innocent strategies, in particular understanding how the mechanism of P-views captures the causal structure.

Recall the *P-view* of an alternating linearization, introduced in Definition 11.1.3.

Lemma 11.3.4. *Consider $\sigma \in \text{NTCG}(A, B)$ sequential innocent.*

Then, for all $tm \in \Downarrow\text{-}\mathcal{L}(\sigma)$, we have $\ulcorner tm \urcorner = [m]_\sigma$.

Proof. Here we treat $[m]_\sigma$ as the sequence induced by its total ordering.

The crucial observation is that if $tm^- n^+ \in \Downarrow\text{-}\mathcal{L}(\sigma)$, then necessarily $m^- \rightarrow_\sigma n^+$. To prove that, we prove by induction that for any $t \in \Downarrow\text{-}\mathcal{L}(\sigma)$: (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 $tm^- \in \Downarrow\text{-}\mathcal{L}(\sigma)$, 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 $tm^+ \in \Downarrow\text{-}\mathcal{L}(\sigma)$, 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 11.3.2, and in particular 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 $t_1 m^+ t_2 n^- \in \Downarrow\text{-}\mathcal{L}(\sigma)$ such that $j(n) = m$ – so $\ulcorner t_1 m t_2 n \urcorner = \ulcorner t_1 \urcorner mn$ then we must have $\partial_\sigma(m) \rightarrow \partial_\sigma(n)$ hence $m \rightarrow_\sigma n$ by Lemma 6.1.16. From these two facts, the lemma is a direct verification by induction on t . \square

We regard this as the fundamental reason why the mechanism of P-views works: on alternating linearizations, it recaptures exactly the causal structure explicitly carried by concurrent strategies – a strength of concurrent games is that beyond the alternating




Figure 11.5: Witness for Figure 9.19

deterministic realm, P-views are no longer available while concurrent strategies remain.

As an illustration, Figure 11.5 shows the configuration explored in the play of

$$\llbracket \lambda f^{\mathbb{U} \rightarrow \mathbb{U}}. f \text{ skip}; f \text{ skip}; \mathbf{tt} \rrbracket : \llbracket (\mathbb{U} \rightarrow \mathbb{U}) \rightarrow \mathbb{B} \rrbracket$$

in Figure 9.19 (note that this is another presentation of the diagram on the left hand side of Figure 10.12) – the numbers in red correspond to the order in which the linearization proceeds. As proved above, at each stage the P-view is exactly 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 visiting 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 move⁴.

Innocence of the unfolding. Now, we show that as expected, the alternating unfolding of a sequential innocent strategy is innocent in the traditional sense.

Proposition 11.3.5. *Consider A a mixed board, and $\sigma : A$ sequential innocent.*

Then, $\Downarrow\text{-Unf}_1(\sigma^1) : \underline{A}$ is an innocent alternating strategy (as in Definition 3.2.9).

Proof. We already know from Proposition 11.2.11 that $\Downarrow\text{-Unf}_1(\sigma^1)$ is P-visible.

Consider $sa^-b^+, ta^- \in \Downarrow\text{-Unf}(\sigma^1)$ s.t. $\ulcorner sa^- \urcorner = \ulcorner ta^- \urcorner$. By definition, there are witnesses $um \in \Downarrow\text{-}\mathcal{L}(\sigma^1)$ for sa^-b^+ and $vo \in \Downarrow\text{-}\mathcal{L}(\sigma^1)$ for ta^- . By Lemma 11.3.4,

$$\ulcorner um \urcorner = [m]_{\sigma^1}, \quad \ulcorner vo \urcorner = [o]_{\sigma^1}$$

⁴Different explorations of the same configuration 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 Mellès' presentation of innocence [Mellès, 2004a].

where $[m]_{\sigma^!}, [o]_{\sigma^!}$ are seen as the sequences induced by their total ordering. Those are gccs, and as in Proposition 11.3.3 they must be alternating, thus $\ulcorner um \urcorner, \ulcorner vo \urcorner \in \Downarrow\text{-}\mathcal{L}(\sigma^!)$.

Now, from $\ulcorner sa^- \urcorner = \ulcorner ta^- \urcorner$, by Proposition 11.2.7 we have that $\partial_{\sigma^!} \ulcorner um \urcorner \cong_A \partial_{\sigma^!} \ulcorner vo \urcorner$, which entails that $\ulcorner um \urcorner \cong_{\sigma^!} \ulcorner vo \urcorner$ by Lemma 11.2.4. Now, as $\ulcorner um \urcorner$ extends with n , by *extension* there must be $\ulcorner vo \urcorner p = \ulcorner vop \urcorner \in \Downarrow\text{-}\mathcal{L}(\sigma^!)$ with $\ulcorner umn \urcorner \cong_{\sigma^!} \ulcorner vop \urcorner$. Now it follows that $|vo| \cup \{p\} \in \mathcal{C}(\sigma^!)$ by *determinism*, so that $vop \in \Downarrow\text{-}\mathcal{L}(\sigma^!)$. We may project it to $ta^-b^+ = (\hat{\cdot})(\partial_{\sigma^!}(vop))$. Finally, since $\ulcorner umn \urcorner \cong_{\sigma^!} \ulcorner vop \urcorner$ we must have

$$\partial_{\sigma^!}(\ulcorner umn \urcorner) \cong_{!A} \partial_{\sigma^!}(\ulcorner vop \urcorner)$$

hence $(\hat{\cdot})(\partial_{\sigma^!}(\ulcorner umn \urcorner)) = (\hat{\cdot})(\partial_{\sigma^!}(\ulcorner vop \urcorner))$, and since $\ulcorner sa^-b^+ \urcorner = \ulcorner (\hat{\cdot})(\partial_{\sigma^!}(umn)) \urcorner = (\hat{\cdot})(\partial_{\sigma^!}(\ulcorner umn \urcorner))$ and likewise for vop it follows that $\ulcorner sa^-b^+ \urcorner = \ulcorner ta^-b^+ \urcorner$. \square

Of course, from the above we also have a characterization of its P-views:

Proposition 11.3.6. *Consider A a mixed board, and $\sigma : A$ sequential innocent.*

The (non-empty) P-views of $\Downarrow\text{-Unf}_i(\sigma^!)$ are exactly those $(\hat{\cdot})(\partial_{\sigma^!}([m]_{\sigma^!}))$ for $m \in \sigma$.

Proof. If $m \in \sigma$, then, picking some arbitrary copy index $i \in \mathbb{N}$, we have $(i, m) \in \sigma^!$, and $[(i, m)]_{\sigma^!} \in \Downarrow\text{-}\mathcal{L}(\sigma^!)$ thus $(\hat{\cdot})(\partial_{\sigma^!}([(i, m)]_{\sigma^!})) \in \Downarrow\text{-Unf}_i(\sigma^!)$. Moreover, by Proposition 11.3.3, it is a P-view of $\Downarrow\text{-Unf}_i(\sigma^!)$. Finally, $(\hat{\cdot})(\partial_{\sigma^!}([(i, m)]_{\sigma^!})) = (\hat{\cdot})(\partial_{\sigma^!}([m]_{\sigma^!}))$.

Reciprocally, consider $s \in \Downarrow\text{-Unf}_i(\sigma^!)$ a non-empty P-view, *i.e.* we have $\ulcorner s \urcorner = s$. By definition, there is $tm \in \Downarrow\text{-}\mathcal{L}(\sigma^!)$ such that $s = (\hat{\cdot})(\partial_{\sigma^!}(tm))$, with $\ulcorner tm \urcorner = tm$. By Lemma 11.3.4, $\ulcorner tm \urcorner = [m]_{\sigma^!}$. But then, $m = (i, m')$ for some copy index $i \in \mathbb{N}$ and $(\hat{\cdot})(\partial_{\sigma^!}([m']_{\sigma^!})) = (\hat{\cdot})(\partial_{\sigma^!}([m]_{\sigma^!})) = s$ as required. \square

We may now wrap up the unfolding to innocent alternating strategies:

Theorem 11.3.7. *We have a cartesian closed \sim -functor:*

$$\Downarrow\text{-Unf}(-) : \multimap\text{-SeqInn} \rightarrow \Downarrow\text{-Inn},$$

preserving the interpretation in the sense that for $\Gamma \vdash M : A$ any term of PCF, we have

$$\Downarrow\text{-Unf}(\llbracket M \rrbracket_{\multimap\text{-SeqInn}}) = \llbracket M \rrbracket_{\Downarrow\text{-Inn}}.$$

Proof. Follows from Theorem 11.2.18 and Proposition 11.3.5. \square

11.3.3 Sequential Globularity

We refine this in the presence of the further constraint of *globularity* (Definition 10.5.1).

Let us start with:

Definition 11.3.8. *A strategy $\sigma \in \multimap\text{-Strat}(A, B)$ is **sequential globular** if it is both sequential (as in Definition 11.1.4) and globular (as in Definition 10.5.1).*

As both sequential and globular strategies form a cartesian closed sub- \sim -category of \rightarrow -**Strat**, there is a cartesian closed \sim -category \rightarrow -**SeqGlob** of sequential globular strategies. As PCF may be interpreted in both \rightarrow -**Seq** and \rightarrow -**Glob**, it follows that it has an adequate interpretation in \rightarrow -**SeqGlob** as well – here we shall prove that this interpretation is actually exactly the same as that in \Downarrow -**InnWB**.

We first show that the unfolding of a sequential globular strategy is well-bracketed:

Proposition 11.3.9. *Consider A and B mixed boards, and $\sigma \in \rightarrow$ -**SeqGlob**(A, B).*

*Then, \Downarrow -Unf(σ) $\in \Downarrow$ -**Inn**($\underline{A}, \underline{B}$) is well-bracketed (in the sense of Definition 3.2.6).*

Proof. Recall that \Downarrow -Unf(σ) is defined as \Downarrow -Unf_i($\Lambda(\sigma)^!$), which is innocent by Proposition 11.3.5. Moreover, by Proposition 11.3.6 its non-empty P-views are exactly the

$$s = (\hat{\cdot})(\partial_\sigma([m]_\sigma))$$

for $m \in \sigma$. But by *wb-threads*, $[m]_\sigma$ is a well-bracketed gcc; it entails directly that s is well-bracketed. Thus by Proposition 3.2.13, \Downarrow -Unf(σ) is well-bracketed. \square

Together with Theorem 11.2.18, we get:

Theorem 11.3.10. *The alternating unfolding yields*

$$\Downarrow$$
-Unf($-$) : \rightarrow -**SeqGlob** \rightarrow \Downarrow -**InnWB**

a cartesian closed \sim -functor preserving the interpretation of PCF.

In fact, we shall be able to prove later on:

Theorem 11.3.11. *The induced functor \Downarrow -Unf($-$) : \rightarrow -**SeqGlob**/ \approx \rightarrow \Downarrow -**InnWB** is full and faithful: for all A, B mixed boards, the alternating unfolding induces a bijection*

$$\Downarrow$$
-Unf($-$) : \rightarrow -**SeqGlob**(A, B)/ \approx \simeq \Downarrow -**InnWB**($\underline{A}, \underline{B}$).

Proof. This requires the notion of *meager form* of a parallel innocent causal strategy, introduced in Section 12.1.2 – hence, we postpone the proof until Section 12.1.3. \square

This should not surprise the reader: we have seen in Section 11.3.1 that the causal shape of a sequential innocent causal strategy is a forest, whose branches correspond to P-views enriched with explicit copy indices. Once we factor out the choice of copy indices via positive iso, we are left with exactly the same information as P-views.

11.3.4 Intensional Full Abstraction

At last, we are finally equipped to prove the final result of this chapter:

Theorem 11.3.12. *The interpretation of PCF in \rightarrow -**SeqGlob** is intens. fully abstract.*

Proof. Let $\vdash M, N : A$ be terms in PCF, and assume $\llbracket M \rrbracket \neq \llbracket N \rrbracket$, i.e. there is $\alpha \in \rightarrow\text{-SeqGlob}(\llbracket A \rrbracket, \llbracket \cup \rrbracket)$ s.t. $\alpha \odot_! \llbracket M \rrbracket \neq \alpha \odot_! \llbracket N \rrbracket$ – assume w.l.o.g. $\alpha \odot_! \llbracket M \rrbracket$ converges while $\alpha \odot_! \llbracket N \rrbracket$ diverges. Then it follows from Theorem 11.3.7 that

$$\Downarrow\uparrow\text{-Unf}(\alpha) \odot \llbracket M \rrbracket_{\Downarrow\uparrow\text{-Strat}} \Downarrow \qquad \Downarrow\uparrow\text{-Unf}(\alpha) \odot \llbracket N \rrbracket_{\Downarrow\uparrow\text{-Strat}} \Uparrow,$$

where $\Downarrow\uparrow\text{-Unf}(\alpha)$ is an innocent alternating strategy, but also well-bracketed by Theorem 11.3.10. Thus finally, it follows that $M \neq N$ by Theorem 4.3.10. \square

This leaves us with one, final theorem to prove in Part III: that the model $\rightarrow\text{-Inn}$ of parallel innocent strategies is intensionally fully abstract for PCF_{\parallel} . This will prove the most challenging of our full abstraction results, investigated in the next chapter.

11.4 History and Related Work

Though it seemed clear for a long time that concurrent strategies supported a fitting notion of sequentiality, the constructions presented in this chapter were not worked out until the long paper [Castellan and Clairambault, 2021], with a few significant surprises in store. The presentation given here is slightly simpler.

Related work. Though the results of this chapter are new, the developments resonate significantly with several earlier works. First and foremost, the structure of sequential innocent strategies is close to the designs of Girard’s *Ludics* [Girard, 2001], and part of our developments are reminiscent of Faggian and Hyland’s work on the connection between Ludics and Hyland-Ong games [Faggian and Hyland, 2002]. Again, the causal shape of sequential innocent strategies was also noticed by Melliès: diagrams such as in Figure 11.3 are very syntactic in nature, and may be written down via terms of the *non-uniform λ -calculus* introduced in [Melliès, 2004a].

Another notable aspect of this chapter is the analysis of polarities necessary to prove *sequentiality* stable under composition and tensoring (see e.g. Lemmas 11.1.8 and 11.1.16). Such lemmas are fundamental structures of the theory of sequentiality. In traditional presentations of alternating strategies, they show up early when proving determinism stable under composition, and associativity (Chapter 4); here instead they show up when proving sequentiality stable under composition. Usually they are kept implicit, relegated to technical appendices (if anything); but they are at the foreground of Melliès’ recent *template games* [Melliès, 2019a, Melliès, 2019b].

Chapter 12

Finite Definability for PCF_{\parallel}

We are approaching the end of the journey. Over the course of the previous three chapters, we have established the following intensional full abstraction results

$$\begin{array}{ll} \multimap\text{-}\mathbf{Strat} & \text{is fully abstract for } \mathbf{IA}_{\parallel}, \\ \multimap\text{-}\mathbf{Strat} + \textit{sequentiality} & \text{is fully abstract for } \mathbf{IA}, \\ \multimap\text{-}\mathbf{Strat} + \textit{globularity} + \textit{sequentiality} & \text{is fully abstract for } \text{PCF}, \end{array}$$


and we are left with the one outstanding objective:

$$\multimap\text{-}\mathbf{Strat} + \textit{globularity} \quad \text{is fully abstract for } \text{PCF}_{\parallel}.$$

This is also the most challenging of our full abstraction results: for the others we could leverage earlier work, but here we must prove finite definability from scratch.

Proving finite definability for globular strategies is a technical endeavour, involving a number of steps. First of all, what is “finite” in “finite definability”? Recall that this is already subtle for traditional innocent strategies, where finiteness is defined with respect to a *meager* representation: the forest of P-views (see Section 3.2.4). Accordingly, we shall start in Section 12.1 by providing a compact, *meager* form for parallel innocent strategies. Once this is done, we focus on the causal shape of globular strategies. There our observation, is that parallel behaviour is inherently *first-order*: two threads whose destiny is to be merged cannot delve into the higher-order structure. This allows us to prove a *factorization* result: any globular strategy can be decomposed into (the interpretation of) a pure λ -term, and a globular purely first-order strategy – this is detailed in Section 12.2. Finally, it remains to show finite definability for first-order globular strategies, which we do up to positional equivalence in Section 12.3.

Let us now delve in, starting with the meager form.

Figure 12.1: Meager form of a sequential innocent strategy on $(U \rightarrow U) \rightarrow B$

12.1 Meager Form

Analogously to traditional innocent strategies, globular causal strategies support a notion of *meager form*, where all Opponent replications are factored out.

Introducing the meager form We have already seen in Section 11.3.1 the computation of the meager form for a sequential innocent strategy, summed up in Figure 12.1: so as to factor out Opponent's duplications, it suffices to only keep those Opponent moves that have 0 as their copy index (it is actually not straightforward how to make this formal for arbitrary mixed boards, we shall see later on how to do that). In doing so, we get a useful notion of finite sequential innocent strategy: although the tree at the left hand side of Figure 12.1 is infinite, the strategy is finite because factoring out Opponent's duplications yields the finite tree on the right hand side of Figure 12.1.

Then, the full strategy can be fully recovered from it: all the branches of the diagram on the left hand side of Figure 12.1 are copies of some branch of the meager form, and thus with each new Opponent copy of a move, we can make an independent copy of the subsequent subtree, choosing copy indices so as to avoid collisions. We shall see that the only information lost in going from left to right is the specific choice of copy indices, which however does not matter up to positive isomorphism.

This motivates the meager form and its expansion for sequential innocent strategies; but in this section we develop this in general for *parallel* innocent strategies. We show an example of the computation of the meager form of a parallel innocent strategy in Figure 12.2 – again we shall see that no essential information is lost. The diagram on the right hand side informs the full causal strategy: there is one copy of the final answer for each pair of answers to the two arguments, so that Opponent providing a new copy of the value for either argument prompts a new independent copy of the final answer.




Figure 12.2: Meager form of parallel composition

12.1.1 Updating Mixed Boards

To make this formal, we need a way to specify the *copy index* of a move occurring deep in the board; forcing us to complete the definition of mixed boards (Definition 9.2.6).

Making copy indices explicit. If A is a board (and thus in particular, for a mixed board) its *causality* is forestial, thus any non-minimal event $a \in A$ has a unique **predecessor** for \rightarrow_A , written $\text{pred}(a) \rightarrow_A a$; by convention we also specify that $\text{pred}(a) = *$ if a is a minimal event. With these conventions in place, we complete Definition 9.2.6 (the intention is that the definition of mixed board was that below from the start, but with some of the components postponed, as they played no role until now).


Definition 12.1.1. A **mixed board** is (A, \underline{A}) with A a strict board, \underline{A} an arena, with

$$\text{lbl}_A : |A| \rightarrow |\underline{A}|, \quad \text{ind}_A : |A| \rightarrow \mathbb{N},$$

with lbl_A a **label** function preserving polarities, satisfying Definition 9.2.6, and ind_A an **indexing** function, with the following additional conditions:

- initialized: for all $\underline{a} \in \min(\underline{A})$ there is a unique $a \in \min(A)$ s.t. $\text{lbl}(a) = \underline{a}$, and it additionally satisfies $\text{ind}(a) = 0$.
- local conflict: if $a \rightsquigarrow a'$, we have $\text{pred}(a) = \text{pred}(a')$ and $\text{ind}(a) = \text{ind}(a')$,
- invariant conflict: if $a_1 \rightsquigarrow a_2$, $\text{lbl}(a_1) = \text{lbl}(b_1)$, $\text{lbl}(a_2) = \text{lbl}(b_2)$, $\text{pred}(b_1) = \text{pred}(b_2)$ and $\text{ind}(b_1) = \text{ind}(b_2)$, then $b_1 \rightsquigarrow b_2$.
- 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 $\text{pred}(a_1) = \text{pred}(a_2)$, then $a_1 = a_2$.
- wide: for any $a \in A$, $\underline{b} \in \underline{A}$ such that $\text{lbl}(a) \rightarrow_{\underline{A}} \underline{b}$, and $n \in \mathbb{N}$, there is $b \in A$ s.t. $\text{pred}(b) = a$, $\text{lbl}(b) = \underline{b}$ and $\text{ind}(b) = n$.
- +-transparent: for $\theta : x \cong_A y$, $\theta \in \mathcal{S}_+(A)$ iff for all $a^- \in x$, $\text{ind}(\theta(a)) = \text{ind}(a)$.
- transparent: for $\theta : x \cong_A y$, $\theta \in \mathcal{S}_-(A)$ iff for all $a^+ \in x$, $\text{ind}(\theta(a)) = \text{ind}(a)$.

In particular, the function ind_A explicitly associates to every move $a \in A$ of the board a **copy index**. While this definition looks intimidating, its intention is simply to bind

Figure 12.3: Copy indices for $\text{ground}(V)$

tightly the board and the arena so that events of the board are determined by their label in the arena, their predecessor in the board and their copy index. By *transparent*, the copy index is the only thing that symmetries can change. By *+transparent*, *positive symmetries* can only change indices of positive moves, while by *--transparent*, *negative symmetries* can only change indices of negative moves.

Accounting for the new component $\text{ind}_A : |A| \rightarrow \mathbb{N}$ of mixed boards, we must now define the copy indices for all our mixed board constructions. For basic mixed boards, recall that they were obtained by linking the arenas for ground types with the ground strict boards obtained as $\text{ground}(V)$ for some set V of values. Thus we must equip this with a copy index function: for a set of values $V = \{a, b, c, \dots\}$, we define the *copy indices* for the ground board $\text{ground}(V)$ (Definition 9.1.1) as specified in Figure 12.3. For all basic mixed boards, it is clear that Definition 12.1.1 is satisfied. We extend the copy index function to all constructions on mixed boards used in the interpretation, via:

$$\begin{aligned}
 \text{ind}_{A_1 \& A_2}(i, a) &= \text{ind}_{A_i}(a) \\
 \text{ind}_{A \Rightarrow B}(2, b) &= \text{ind}_B(b) \\
 \text{lbl}_{A \Rightarrow B}(1, (b, (i, a))) &= i && \text{if } a \in \min(A), \\
 \text{lbl}_{A \Rightarrow B}(1, (b, (i, a))) &= \text{ind}_A(a) && \text{otherwise,}
 \end{aligned}$$

for $B = \&_{i \in I} B_i$ with B_i well-opened. We omit the routine verification that these definitions indeed preserve the conditions of Definition 12.1.1.

Introducing explicit copy indices only makes explicit a notion that we have already referred to in the exposition; for instance in diagrams such as in Figure 11.5 where the greyed-out subscripts are nothing but the numbers given by the copy index function. In mixed boards, the board is essentially an expanded form of the arena, with all branches replaced with countably many copies: playing a new move consists in picking a justifier (if any), picking an enabled move in the arena, and picking a copy index¹

Properties. Equipped with this now completed notion of mixed board, we may now prove the property that was postponed after Definition 9.2.6.

¹In fact, we could have set $\rightarrow\text{-Strat}$ as having mere arenas as objects, and having strategies play on the generated board – this was the road suggested in [Castellan et al., 2019]. We prefer the present presentation, which allows for a smoother link with the relative Seely category structure of **NTCG** and the relational model.

Lemma 9.2.10. *Consider A a mixed board. We have the following properties:*

- (1) *For all $ta \in \mathcal{O}\text{-PrePlays}_{\mathcal{S}}(\underline{A})$, $s \in \mathcal{O}\text{-Plays}(A)$ such that $t = \mathcal{S}$, there is $sa' \in \mathcal{O}\text{-Plays}(A)$ such that $(\overset{\circ}{\leftarrow})(sa') = ta$;*
- (2) *For all $ta \in \mathcal{O}\text{-PrePlays}_{\mathcal{S}}(\underline{A})$, $s \in \mathcal{O}\text{-Plays}(!A)$ such that $t = \mathcal{S}$, there is $sa' \in \mathcal{O}\text{-Plays}(!A)$ such that $(\overset{\circ}{\leftarrow})(sa') = ta$.*

Proof. (1) Consider $ta \in \mathcal{O}\text{-PrePlays}_{\mathcal{S}}(\underline{A})$ and $s \in \mathcal{O}\text{-Plays}(A)$ such that $t = \mathcal{S}$. If t is empty, then a is initial in \underline{A} and the property follows by *initialized*. Otherwise, write

$$s = s_1 \dots s_n \quad t = \mathcal{S} = t_1 \dots t_n$$

where a points to t_i . Now, consider p a copy index not yet appearing in s . By *wide*, there is $a' \in A$ such that $\text{lbl}(a') = a$, $\text{pred}(a') = s_i$ and $\text{ind}(a') = p$. Then $sa' \in \mathcal{L}(A)$. Indeed, $|s| \cup \{a'\} \in \mathcal{C}(A)$: it is down-closed as $\text{pred}(a') \in |s|$ and conflict-free by *locally conflicting* since no other move in s has index p . Finally, $(\overset{\circ}{\leftarrow})(sa') = ta$.

(2) same reasoning as for (1). □

12.1.2 Meager Innocent Strategies

Now, we are equipped to define *meager parallel innocent strategies*, the notion analogous to the *meager innocent strategies* in the traditional setting (see Section 3.2.4).

Definition. In analogy to the traditional case, *meager innocent strategies* are supposed to be parallel innocent strategies, but playing on games not allowing any Opponent replications. We formalize this by restricting Opponent to copy index 0:

Definition 12.1.2. *Consider A a mixed board. Then, we set A_+ a strict board with*

$$|A_+| = \{a' \in |A| \mid \forall a^- \leq_A a', \text{ind}_A(a) = 0\},$$

*and all other components inherited. We call A_+ the **positive restriction** of A .*

This definition of the positive restriction has the immediate consequence that its only non-trivial symmetries are positive, since Opponent moves must have index 0:

Lemma 12.1.3. *For A a mixed board, if $\theta \in \mathcal{S}(A_+)$ then $\theta \in \mathcal{S}_+(A_+) \subseteq \mathcal{S}_+(A)$.*

This is obvious by *+transparent*. We now define meager parallel innocent strategies:

Definition 12.1.4. *Consider A a mixed board.*

*A **meager parallel innocent strategy** on A is a parallel innocent strategy on A_+ .*

For instance, the diagrams on the right hand side of Figures 12.1 and 12.2 are meager parallel innocent strategies. Note that on the left hand side of these diagrams, the infinite nature of the strategies forces us to resort to a symbolic notation so as to describe those finitely; also the notation conveys no information on the symmetries, although those are in principle part of the data. In contrast, for meager parallel innocent strategies:

Proposition 12.1.5. *Consider A a mixed board, and $\sigma : A_+$ meager parallel innocent. Then its symmetries are trivial, i.e. $\mathcal{S}(\sigma) = \{\text{id}_x : x \cong_\sigma x \mid x \in \mathcal{C}(\sigma)\}$.*

Proof. Consider $\theta : x \cong_\sigma y$ any symmetry. By Lemma 12.1.3, $\partial_\sigma \theta \in \mathcal{S}_+(A_+)$. Hence by *thin*, we have that $x = y$ and $\theta = \text{id}_x$; thus $\mathcal{S}(\sigma)$ is restricted to identities. \square

Thus when specifying a meager parallel innocent strategy σ , it is unnecessary to specify the isomorphism family $\mathcal{S}(\sigma)$, which is trivial. Thus it is entirely sufficient to describe σ with a diagram (as on the right hand side of Figures 12.1 and 12.2) describing the event structure concretely, with no ellipse or symbolic representation. Likewise, a *positive iso* between meager parallel innocent strategies boils down to an isomorphism

$$\varphi : \sigma \cong \tau$$

of plain event structures such that $\partial_\tau \circ \varphi \sim \partial_\sigma$, which amounts to renaming the internal events of σ and τ and changing the copy indices of positive moves.

Meager form. The above is interesting, *provided* it does indeed give an alternative but equivalent representation for usual parallel innocent strategies. Accordingly, we now define the *meager form* of parallel innocent strategies – in the obvious way:

Proposition 12.1.6. *Consider A a mixed board, and $\sigma : A$ a parallel innocent strategy. Then, setting the set of events*

$$|\text{mf}(\sigma)| = \{m \in |\sigma| \mid \partial_\sigma[m]_\sigma \in \mathcal{C}(A_+)\}$$

with all components inherited yields a meager parallel innocent strategy $\text{mf}(\sigma)$ on A . Moreover, $\text{mf}(-)$ preserves \approx .

Proof. There are a few conditions to check, which we list below.

Event structure with symmetry. It is easy that $\text{mf}(\sigma)$ is an event structure. For symmetry only the *extension* axiom is non-trivial: if $\theta : x \cong_\sigma y$ with $x, y \in \mathcal{C}(\text{mf}(\sigma))$, then by the same reasoning as in Proposition 12.1.5, $\theta = \text{id}_x$; *extension* follows.

Strategy, winning, visible, parallel innocent. Immediate verifications.

Preservation of \approx . Consider $\varphi : \sigma \approx \tau$ a positive isomorphism. Then, the important observation is that φ sends $\text{mf}(\sigma)$ to $\text{mf}(\tau)$. Indeed if $m \in \text{mf}(\sigma)$, then

$$\langle \partial_\sigma, \partial_\tau \circ \varphi \rangle_{[m]_\sigma} : \partial_\sigma[m]_\sigma \cong_A^+ \partial_\tau[\varphi(m)]_\tau$$

which by *+transparent*, being positive, must preserve the copy indices of negative events. Hence $\partial_\tau[\varphi(m)]_\tau \in \mathcal{C}(A_+)$, and φ restricts to $\text{mf}(\varphi) : \text{mf}(\sigma) \approx \text{mf}(\tau)$. \square

Note that in the construction of $\text{mf}(\sigma)$ we have never actually used the hypothesis that σ is parallel innocent, and the construction works in more generality. However, the operation is not very interesting for non-innocent strategies as it is lossy: if σ is not parallel innocent, one cannot in general reconstruct σ from $\text{mf}(\sigma)$.

Next, we show how to recover a parallel innocent strategy from a meager form.




Figure 12.4: Expansion of meager parallel composition

Expansion of a meager parallel innocent strategy. Fix A a mixed board, and $\sigma : A_+$ meager parallel innocent on A . We must infer from σ its behaviour if Opponent performs arbitrary duplications, by performing adequately many copies of its events. We shall achieve that by pairing each event of a meager parallel innocent strategy with the choice of a copy index for each of its duplicable Opponent dependencies, *i.e.* its negative dependencies save for the initial move – we call such choice an *exponential slice*. We must then assign a new copy index for positive moves, hashing together the positive move and the exponential slice. The process is displayed in Figure 12.4.

Consider $m \in \sigma$. How many copies of m should we include in the expanded form? The answer is that there should be one copy for each simultaneous choice of a copy index for each duplicable negative dependency of m . Write

$$[m]_{\sigma}^{-} = \{n^{-} \in |\sigma| \mid n^{-} \leq_{\sigma} m\}$$

for the **negative dependencies** of m . Then, we define:

Definition 12.1.7. An *exponential slice* for σ is a function, for some $m \in \sigma$,

$$\alpha : [m]_{\sigma}^{-} \rightarrow \mathbb{N}.$$

such that $\alpha(n) = 0$ for n minimal.

An *expanded event* is $m = (m, \alpha_m : [m]_{\sigma}^{-} \rightarrow \mathbb{N})$, we write $m = (m, \alpha) \in \exp(\sigma)$.

We write $\text{lbl}((m, \alpha)) = m$. If $\text{lbl}(m)$ is negative we say that m is **negative**, and write $\text{ind}((m, \alpha)) = \alpha(m)$ for its **copy index**. For $m = (m, \alpha), n = (n, \beta) \in \exp(\sigma)$, we write $m \sqsubseteq n$ iff $m \leq_{\sigma} n$, and for all $e^{-} \leq_{\sigma} m$, we have $\alpha(e^{-}) = \beta(e^{-})$. Note that if m is negative, as $\text{lbl}(m)$ has a unique predecessor in σ , m must have a unique predecessor in $\exp(\sigma)$, written $\text{pred}(m)$. We set the *minimal conflict* $m \sim_{\exp(\sigma)} n$ iff m and n are negative, $\text{pred}(m) = \text{pred}(n)$, $\text{lbl}(m) \sim_{\sigma} \text{lbl}(n)$ and $\text{ind}(m) = \text{ind}(n)$.

Towards defining a strategy, we first extend this to an event structure with symmetry:

Proposition 12.1.8. Consider $\exp(\sigma)$ with events $|\exp(\sigma)|$, components

$$\begin{aligned} m \leq_{\exp(\sigma)} n &\Leftrightarrow m \sqsubseteq n \\ m \#_{\exp(\sigma)} n &\Leftrightarrow \exists m' \sqsubseteq m, n' \sqsubseteq n, m' \sim_{\exp(\sigma)} n' \end{aligned}$$

and $\mathcal{S}(\exp(\sigma))$ comprising all order-isos preserving $\text{lbl}(-)$ between configurations.

Then, $\exp(\sigma)$ is an event structure with symmetry.

Proof. The less trivial point is that $\#_{\exp(\sigma)}$ is irreflexive so assume $m \#_{\exp(\sigma)} m$, meaning there are $n, o \sqsubseteq m$ such that $n \sim_{\exp(\sigma)} o$. So $\text{pred}(n) = \text{pred}(o) = p$, $\text{lbl}(n) \sim_{\sigma} \text{lbl}(o)$, and $\text{ind}(n) = \text{ind}(o)$. But this means we have $\text{lbl}(n) \leq_{\sigma} \text{lbl}(m)$ and $\text{lbl}(o) \leq_{\sigma} \text{lbl}(m)$ with $\text{lbl}(n) \sim_{\sigma} \text{lbl}(o)$, thus $\text{lbl}(m) \#_{\sigma} \text{lbl}(m)$, contradiction since σ is an event structure.

Event structure with symmetry. The non-trivial point is the *extension* axiom. Consider $\theta : x \cong_{\exp(\sigma)} y$, and $x \vdash_{\exp(\sigma)} (m, \alpha)$. If m is positive, then for each $n^- \leq_{\sigma} m$, we have $(n^-, \alpha \upharpoonright [n^-]_{\sigma}^-) \in x$, and we write $\theta((n^-, \alpha \upharpoonright [n^-]_{\sigma}^-)) = (n^-, \beta_n)$. Then we set:

$$\gamma(n^-) = \beta_n(n^-)$$

for all $n^- \leq_{\sigma} m$, yielding $(m, \gamma) \in \exp(\sigma)$. Then, we observe that $y \vdash_{\exp(\sigma)} (m, \gamma)$: indeed if $(p, \gamma') \leq_{\exp(\sigma)} (m, \gamma)$ then one can prove that $\theta((p, \alpha')) = (p, \gamma')$ for $\alpha' = \alpha \upharpoonright [p]_{\sigma}^-$. Moreover, $y \cup \{(m, \gamma)\}$ is consistent as immediate conflicts are negative. Likewise $\theta \vdash ((m, \alpha), (m, \gamma))$ by definition of $\mathcal{S}(\exp(\sigma))$. Finally if m is negative, then write $(n, \beta) = \text{pred}(m, \alpha)$ and $(n, \beta') = \theta(n, \beta)$. Set $\alpha' : [m]_{\sigma} \rightarrow \mathbb{N}$ as $\alpha'(o) = \beta(o)$ for $o <_{\sigma} m$, and $\alpha'(m) = i$ fresh; then it is straightforward that $\theta \vdash ((m, \alpha), (m, \alpha'))$. \square

Now, in order to obtain a strategy, we must specify its display map. In particular, we must set a copy index for all moves. For negative moves, it should clearly match that given by the exponential slice. For positive moves, we choose arbitrarily

$$\sharp : \Sigma_{m^+ \in |\sigma|} \mathbb{N}^{[m^+]_{\sigma}^-} \rightarrow \mathbb{N}$$

any injective function, which we shall use to fix a copy index for positive moves – note that the strategy obtained will not depend (up to positive iso) on this choice.

In order to better formulate the definition, we shall use the lemma:

Lemma 12.1.9. Consider $\underline{a} \in \underline{A}$ non-minimal, and write $\text{pred}(\underline{a}) = \underline{b}$ its predecessor.

Then for all $b \in A$ such that $\text{lbl}(b) = \underline{b}$, and $n \in \mathbb{N}$, there is a unique $a \in A$ such that $\text{lbl}(a) = \underline{a}$, $\text{pred}(a) = b$ and $\text{ind}(a) = n$.

Proof. Direct consequence of wide and jointly injective of mixed boards. \square

Notice also that each non-minimal $m \in \exp(\sigma)$ has a natural notion of **justifier**: for $m = (m, \alpha)$, we set $j(m) = (j(m), \alpha \upharpoonright [j(m)]_{\sigma}^-)$ – we make use of this below:

Lemma 12.1.10. *The function defined by induction on \leq_σ via the clauses*

$$\begin{aligned} \partial_{\exp(\sigma)}(m^-, \alpha) = a^- \quad \text{such that} & \begin{cases} \text{pred}(a^-) = \partial_{\exp(\sigma)}(j(m^-, \alpha)) \\ \text{lbl}(a^-) = \text{lbl}(\partial_\sigma(m^-)) \\ \text{ind}(a^-) = \alpha(m^-) \end{cases} \\ \partial_{\exp(\sigma)}(m^+, \alpha) = a^+ \quad \text{such that} & \begin{cases} \text{pred}(a^+) = \partial_{\exp(\sigma)}(j(m^+, \alpha)) \\ \text{lbl}(a^+) = \text{lbl}(\partial_\sigma(m^+)) \\ \text{ind}(a^+) = \sharp(m^+, \alpha) \end{cases} \end{aligned}$$

is a map of event structures with symmetry $\partial_{\exp(\sigma)} : \exp(\sigma) \rightarrow A$.

Proof. It is clear by construction that the image of a configuration is down-closed, and it is easily verified that $\partial_{\exp(\sigma)}$ is locally injective. Now considering $x \in \mathcal{C}(\exp(\sigma))$, we must show that $\partial_{\exp(\sigma)}(x)$ is conflict-free. Seeking a contradiction, assume there are $m = (m^-, \alpha)$ and $n = (n^-, \beta)$ s.t. $\partial_{\exp(\sigma)}(m) \sim_{\exp(\sigma)} \partial_{\exp(\sigma)}(n)$. By definition of $\partial_{\exp(\sigma)}$, $\text{lbl}(\partial_\sigma(m^-)) = \text{lbl}(\partial_{\exp(\sigma)}(m))$ and $\text{lbl}(\partial_\sigma(n^-)) = \text{lbl}(\partial_{\exp(\sigma)}(n))$. By *local conflict*, $\partial_{\exp(\sigma)}(m)$ and $\partial_{\exp(\sigma)}(n)$ must have the same predecessor, so $j(m^-, \alpha)$ and $j(n^-, \beta)$ have the same display. But they are compatible, thus by local injectivity they must be equal – in turn this entails that $\text{pred}(m^-) = \text{pred}(n^-)$, thus $\text{pred}(\partial_\sigma(m^-)) = \text{pred}(\partial_\sigma(n^-))$ by *courtesy*. Thus by *invariant conflict*, we have $\partial_\sigma(m^-) \sim_A \partial_\sigma(n^-)$ which entails $m^- \sim_\sigma n^-$. Summing up, $\text{lbl}(m) \sim_\sigma \text{lbl}(n)$, $\text{ind}(m) = \text{ind}(n)$ since $\text{ind}(\partial_{\exp(\sigma)}(m)) = \text{ind}(\partial_{\exp(\sigma)}(n))$ by *local conflict*, and $\text{pred}(m) = \text{pred}(n)$, thus $m \sim_{\exp(\sigma)} n$.

Finally if $\theta \in \mathcal{S}(\exp(\sigma))$, it is by definition an order-isomorphism preserving labels. From that it is obvious that it must preserve justifiers. It follows from the definition that $\partial_{\exp(\sigma)} \theta$ is an order-iso preserving labels, hence a symmetry by *transparent*. \square

There is still a lot of work to do in order to obtain a causal strategy; first we must construct a strategy in the sense of Definition 6.1.6. But in order to prove properties based on configurations (starting with *receptive*), we must understand better how configurations of $\exp(\sigma)$ relate to configurations of σ . We shall establish such a connection not in general, but in the particular case of so-called *linear* configurations:

Definition 12.1.11. *Consider A a mixed board, $\sigma : A$ a strategy.*

*Then, $x \in \mathcal{C}(\sigma)$ is **linear** if for all $n_1^-, n_2^- \in x$, if $\text{pred}(n_1^-) = \text{pred}(n_2^-)$ then $n_1^- = n_2^-$.*

Linear configurations are analogous to P-views: Opponent may point on any given Player move at most once. Their interest comes from the following observation:

Lemma 12.1.12. *If $\sigma : A$ is parallel innocent and $m \in \sigma$, $[m]_\sigma \in \mathcal{C}(\sigma)$ is linear.*

If $\sigma : A$ is meager parallel innocent and $m \in \exp(\sigma)$, then $[m]_{\exp(\sigma)}$ is linear.

Proof. The first is a direct consequence of pre-innocence (and Lemma 6.1.16); the second is straightforward from the definition, boiling down to the pre-innocence of σ . \square

Next, we observe that if $\sigma : A$ is a meager parallel innocent strategy, any linear configuration $x \in \mathcal{C}(\sigma)$ may be canonically regarded as a (linear) configuration of $\exp(\sigma)$ by adjoining the constant 0 copy index to all Opponent moves:

Lemma 12.1.13. *Consider A a mixed board, and $\sigma : A$ a meager parallel innocent. For each linear $x \in \mathcal{C}(\sigma)$, with $0_m : [m]_\sigma^- \rightarrow \mathbb{N}$ the constant 0 function, setting*

$$\exp(x) = \{(m, 0_m) \mid m \in x\},$$

we get $\exp(x) \in \mathcal{C}(\exp(\sigma))$ a linear configuration; and $\kappa_x : m \mapsto (m, 0_m)$ is an order-isomorphism $\kappa_x : x \cong \exp(x)$ compatible with display maps, i.e. $\partial_{\exp(\sigma)} \circ \kappa_x \sim^+ \partial_\sigma$.

Proof. Straightforward, since meager parallel innocent strategies play on A_+ . \square

It is also straightforward that κ_x preserves and reflects symmetry. Reciprocally, all linear configurations of $\exp(\sigma)$ are canonically symmetric to a configuration of σ :

Lemma 12.1.14. *Consider A a mixed board, meager $\sigma : A$, and $x \in \mathcal{C}(\exp(\sigma))$ linear. There are unique $\text{mf}(x) \in \mathcal{C}(\sigma)$ and $\theta_x : x \cong_{\exp(\sigma)} \exp(\text{mf}(x))$, given by*

$$\text{mf}(x) = \{m \mid (m, \alpha) \in x\}, \quad \theta_x = \{((m, \alpha), (m, 0_m)) \mid (m, \alpha) \in x\}.$$

Proof. Existence. We prove that this data satisfies the required conditions by induction on x . For x empty this is clear. If $x \vdash (m^-, \alpha)$, then by induction hypothesis, $\text{mf}(x) \in \mathcal{C}(\sigma)$. It is immediate that $\text{mf}(x) \cup \{m^-\}$ is down-closed. If it was conflicting, then there would be $n^- \in \text{mf}(x)$ such that $m^- \sim_\sigma n^-$. But then, by Lemma 6.1.17, we have $\partial_\sigma(m^-) \sim_A \partial_\sigma(n^-)$, which by *local conflict* entails $\text{pred}(\partial_\sigma(m^-)) = \text{pred}(\partial_\sigma(n^-))$, hence $\text{pred}(m^-) = \text{pred}(n^-)$ by Lemma 6.1.16. Writing $p^+ \in \sigma$ for this predecessor,

$$(p^+, \alpha') \rightarrow_{\exp(\sigma)} (m^-, \alpha), \quad (p^+, \beta') \rightarrow_{\exp(\sigma)} (n^-, \beta)$$

with all four events in $x \cup \{(m^-, \alpha)\}$ linear. If $\alpha' = \beta'$ we get $(m^-, \alpha) = (n^-, \beta)$ by linearity hence $m^- = n^-$, contradiction. Otherwise $(p^+, \alpha'), (p^+, \beta') \in x$ differ, but this directly entails a failure of linearity as well – contradiction. Now if $x \vdash (m^+, \alpha)$, as above $\text{mf}(x) \cup \{m^+\}$ is down-closed. It must also be conflict-free, since by *causal determinism*, minimal conflicts in σ are negative. It is direct to check that $\xi_x : x \cong \text{mf}(x)$ sending (m, α) to m is an order-isomorphism and that $\theta_x = \kappa_{\text{mf}(x)} \circ \xi_x$ is an order-iso preserving labels, hence a symmetry in $\exp(\sigma)$ as required.

Uniqueness. Assume we have $y \in \mathcal{C}(\sigma)$ and $\theta'_x : x \cong_{\exp(\sigma)} \exp(y)$. Then,

$$\theta'_x \circ \theta_x^{-1} : \exp(\text{mf}(x)) \cong_{\exp(\sigma)} \exp(y)$$

which entails $\exp(\text{mf}(x)) = \exp(y)$ (so that $\text{mf}(x) = y$) and $\theta'_x \circ \theta_x^{-1}$ is an identity by definition of symmetries in $\exp(\sigma)$; this concludes the proof. \square

Now, we may finally prove:

Proposition 12.1.15. *Consider A a mixed board, and $\sigma : A$ a meager parallel innocent. Then, $\exp(\sigma)$ with the components above is a parallel innocent strategy.*

Proof. Receptive. Consider $x \in \mathcal{C}(\exp(\sigma))$ and $\partial_{\exp(\sigma)} x \vdash_A a^-$. Call $c^+ \in \partial_{\exp(\sigma)} x$ the predecessor of a^- in x – there is a unique $m = (m, \alpha) \in x$ such that $\partial_{\exp(\sigma)} m = c^+$. By Lemma 12.1.12, $[m]_{\exp(\sigma)}$ is linear, and thus by Lemma 12.1.14 there is

$$\theta_x : [(m, \alpha)]_{\exp(\sigma)} \cong_{\exp(\sigma)} [(m, 0_m)]_{\exp(\sigma)},$$

mapping via $\partial_{\exp(\sigma)}$ to a symmetry in A which must extend by some (a^-, b^-) . By *wide*, we may assume that $\text{ind}(b^-) = 0$ so that the extension is in A_+ . By Lemma 12.1.13 and receptivity of σ , there is a unique $[m]_\sigma \vdash_\sigma n^-$ such that $\partial_\sigma(n^-) = b^-$. Extending α to $\alpha' : [n^-]_\sigma^- \rightarrow \mathbb{N}$ with $\alpha'(n^-) = \text{ind}(a^-)$, by *jointly injective* we get the required extension – it is compatible with x by definition of minimal conflict.

~receptive is similar, and all further conditions are straightforward. \square

This at last concludes the construction of $\exp(\sigma)$ as a parallel innocent strategy.

Equivalence of meager and expanded forms. Finally, so as to show that meager parallel innocent strategies provide an equivalent representation of parallel innocent strategies, we must prove that the two constructions $\text{mf}(-)$ and $\exp(-)$ preserve positive isomorphisms, and are inverses up to positive isomorphism.

Proposition 12.1.16. *The constructions $\text{mf}(-)$ and $\exp(-)$ preserve positive isos.*

Proof. It is obvious that $\text{mf}(-)$ preserves positive isomorphisms. Consider $\varphi : \sigma \approx \tau$ a positive iso between meager parallel innocent strategies on A ; we must construct

$$\exp(\varphi) : \exp(\sigma) \approx \exp(\tau)$$

a positive isomorphism between the expansions. We simply set:

$$\begin{aligned} \exp(\varphi) : \exp(\sigma) &\approx \exp(\tau) \\ (m, \alpha) &\mapsto (\varphi(m), \alpha') \end{aligned}$$

where for $\varphi(n)^- \leq_\tau \varphi(m)$, $\alpha'(\varphi(n)) = \alpha(n)$. It is direct that this is an order-iso; it preserves and reflects immediate conflict by *causal determinism* and *invariant conflict*. For positivity, for all $(m, \alpha) \in \exp(\sigma)$, $\exp(\varphi)$ on $[(m, \alpha)]_{\exp(\sigma)}$ factors as

$$\begin{array}{ccccccc} & & \kappa^{-1} \circ \theta & & \varphi & & \theta^{-1} \circ \kappa \\ & & \cong & & \cong & & \cong \\ [(m, \alpha)]_{\exp(\sigma)} & & [m]_\sigma & & [\varphi(m)]_\tau & & [(\varphi(m), \alpha')]_{\exp(\tau)} \\ & \searrow \partial_{\exp(\sigma)} & \searrow \partial_\sigma & \searrow \partial_\tau & \searrow \partial_{\exp(\tau)} & & \\ & & & A & & & \end{array}$$

omitting subscripts for readability – where every triangle commutes up to symmetry via Lemmas 12.1.13 and 12.1.14, and positivity follows by *+transparent*. \square

It remains to show that $\text{mf}(-)$ and $\exp(-)$ are inverses up to positive iso. We start with the easy direction, that a meager strategy can be recovered from its expansion:

Proposition 12.1.17. *Consider $\sigma : A$ a meager parallel innocent. Then, the function*

$$\begin{aligned} \Phi &: \sigma \rightarrow \text{mf}(\text{exp}(\sigma)) \\ m &\mapsto (m, 0_m), \end{aligned}$$

yields a positive isomorphism $\Phi : \sigma \approx \text{mf}(\text{exp}(\sigma))$.

Proof. We omit the direct verification that this is an isomorphism of ess. Positivity follows from Lemma 12.1.13 on prime configurations, extended by taking unions. \square

Finally we must show that analogously, a full parallel innocent strategy can always be recovered from its meager form. This means that given a parallel innocent strategy $\sigma : A$, we must provide a bijection between events of σ and of $\text{exp}(\text{mf}(\sigma))$, hence extract from any $m \in \sigma$ a *label* in $\text{mf}(\sigma)$ and an *exponential slice*.

We start by the following generalization of Lemma 12.1.14:

Lemma 12.1.18. *Consider A a mixed board, and $\sigma : A$ a parallel innocent strategy.*

For $x \in \mathcal{C}(\sigma)$ linear, there are unique $\text{mf}(x) \in \mathcal{C}(\text{mf}(\sigma))$ and $\theta_x : x \cong_\sigma \text{mf}(x)$.

Proof. Existence. By induction on x . If x is empty there is nothing to do. If $x \vdash_\sigma a^-$, then θ_x extends with some (a^-, b^-) . By *wide*, we may assume that $\text{ind}(b^-) = 0$, so that it fits in A_+ (such b^- cannot be in $\text{mf}(x)$ by linearity of x). If $x \vdash_\sigma a^+$, then the extension of θ_x automatically does the trick.

Uniqueness. As in Lemma 12.1.14. \square

In particular, given parallel innocent $\sigma : A$ and $m \in \sigma$ we know from Lemma 12.1.12 that $[m]_\sigma$ is linear, thus Lemma 12.1.18 gives us a unique $\text{mf}([m]_\sigma) = [\text{lbl}(m)]_\sigma \in \mathcal{C}(\text{mf}(\sigma))$ with $\theta_m : [m]_\sigma \cong_\sigma [\text{lbl}(m)]_\sigma$. As the notation suggests, $\text{lbl}(m)$ is considered as the *label* of m , i.e. the move of $\text{mf}(\sigma)$ that m is considered a duplicate of.

This duplicate is also characterized by an *exponential slice*, defined as:

$$\begin{aligned} \alpha &: [\text{lbl}(m)]_\sigma^- \rightarrow \mathbb{N} \\ \theta_m(n) &\mapsto \text{ind}(\partial_\sigma(n)), \end{aligned}$$

and writing $\alpha = \text{sl}(m)$ we get $(\text{lbl}(m), \text{sl}(m)) \in \text{exp}(\text{mf}(\sigma))$. Besides:

Proposition 12.1.19. *Consider $\sigma : A$ a parallel innocent. Then, the function*

$$\begin{aligned} \Psi &: \sigma \rightarrow \text{exp}(\text{mf}(\sigma)) \\ m &\mapsto (\text{lbl}(m), \text{sl}(m)) \end{aligned}$$

yields a positive isomorphism $\Psi : \sigma \approx \text{exp}(\text{mf}(\sigma))$.

Proof. Injective. We prove that if $\text{lbl}(m) = \text{lbl}(m')$ and $\text{sl}(m) = \text{sl}(m')$, then $m = m'$. Indeed if $\text{lbl}(m) = \text{lbl}(m')$, by definition this means there are two symmetries

$$\theta_m : [m]_\sigma \cong_\sigma [\text{lbl}(m)]_\sigma, \quad \theta_{m'} : [m']_\sigma \cong_\sigma [\text{lbl}(m)]_\sigma$$

which we may compose to get $\theta = \theta_{m'}^{-1} \circ \theta_m : [m]_\sigma \cong_\sigma [m']_\sigma$. Now, since $\text{sl}(m) = \text{sl}(m')$, it follows that $\partial_\sigma \theta$ preserves indices of negative events, so that it is positive by *+-transparent*. Hence, θ is an identity and $[m]_\sigma = [m']_\sigma$ since σ is *thin*.

Surjective. Consider $(n, \alpha) \in \exp(\text{mf}(\sigma))$, then we construct by induction on m some $m \in \sigma$ s.t. $\text{lbl}(m) = n$ and $\text{sl}(m) = \alpha$, along with a symmetry $\theta : [m]_\sigma \cong_\sigma [n]_\sigma$. If $(n, \alpha) \rightarrow_{\exp(\text{mf}(\sigma))} (n', \alpha')$ positive, by necessity $\alpha' = \alpha$. Then θ extends to $\theta' : [m']_\sigma \cong_\sigma [n']_\sigma$ so that $\text{lbl}(m') = n'$, and it is direct that $\text{sl}(m') = \alpha$ as required. If (n', α') is negative, then by *~-receptive* and *wide*, there is a unique $m \rightarrow_\sigma m'$ such θ extends with (m, m') , and $\text{ind}(\partial_\sigma(m)) = \alpha'(m')$ – so that $\text{lbl}(m') = n'$ and $\text{sl}(m') = \alpha'$.

Remaining verifications are direct, positivity following from *+-transparent*. \square

Altogether, we have proved:

Theorem 12.1.20. *Consider A a mixed board. The operations $\text{mf}(-)$ and $\exp(-)$ inform a one-to-one correspondence between parallel innocent strategies on A (up to positive iso), and meager parallel innocent strategies on A (up to positive iso).*

Proof. The statement packages Propositions 12.1.16, 12.1.17 and 12.1.19. \square

Thus, just as for traditional alternating innocent strategies, parallel innocent strategies have a compact representation via the *meager form*. This is important conceptually: in traditional alternating innocent strategies, the meager form may be regarded as *syntax* as meager sequential innocent strategies are alternative representations of a syntactic notion of normal form for PCF, the *PCF Böhm trees* – we reviewed this correspondence in Section 3.2.4. Here, by analogy, meager parallel innocent strategies may also be regarded as syntax, but *DAG-shaped* rather than tree-shaped.

We included Theorem 12.1.20 because of its conceptual importance, however we shall not rely on it in our definability theorem: we shall only use the notion of *meager form* insofar as it gives an adequate notion of *finiteness* for parallel innocent strategies, that extends the *finiteness* of traditional innocent strategies reviewed in Section 3.2.4.

12.1.3 The Meager Form of Sequential Globularity

Here we take a small detour, and prove Theorem 11.3.11 which we postponed until we had access to the meager form of parallel innocent strategies. We recall the statement:

Theorem 11.3.11. *The induced functor $\Downarrow\text{-Unf}(-) : \rightarrow\text{-SeqGlob}/\approx \rightarrow \Downarrow\text{-InnWB}$ is full and faithful: for all A, B mixed boards, the alternating unfolding induces a bijection*

$$\Downarrow\text{-Unf}(-) : \rightarrow\text{-SeqGlob}(A, B)/\approx \simeq \Downarrow\text{-InnWB}(\underline{A}, \underline{B}).$$

Proof. Faithful. Consider $\sigma, \tau : A$ sequential globular on mixed board A , and assume $\Downarrow\text{-Unf}_1(\sigma^!) = \Downarrow\text{-Unf}_1(\tau^!)$. In particular, they have the same P-views – recall from Proposition 11.3.6 that P-views of $\Downarrow\text{-Unf}_1(\sigma^!)$ are exactly those of the form

$$(\hat{\curvearrowright})(\partial_\sigma([m]_\sigma))$$

for all events $m \in \sigma$. But by Lemma 12.1.18, there is a unique

$$\theta_m : [m]_\sigma \cong_\sigma [\text{lbl}(m)]_\sigma$$

with $\text{lbl}(m) \in \text{mf}(\sigma)$. From that, it is clear that m and $\text{lbl}(m)$ generate the same P-view, i.e. $(\hat{\cdot})(\partial_\sigma([m]_\sigma)) = (\hat{\cdot})(\partial_\sigma([\text{lbl}(m)]_\sigma))$. Therefore, the P-views of $\Downarrow\text{-Unf}_i(\sigma^!)$ can be obtained simply from the events of the meager form of σ . Moreover, for any P-view s on \underline{A} , there is at most one $m \in \text{mf}(\sigma)$ generating s – if they were two, they would be symmetric by Lemma 11.2.4, and thus equal by Proposition 12.1.5. Thus, any P-view of $\Downarrow\text{-Unf}_i(\sigma^!) = \Downarrow\text{-Unf}_i(\tau^!)$ is witnessed by *exactly one* event of $\text{mf}(\sigma)$, and *exactly one* event of $\text{mf}(\tau)$. This entails a bijection $\varphi : |\text{mf}(\sigma)| \simeq |\text{mf}(\tau)|$, and it is a direct verification that this entails a positive isomorphism $\varphi : \text{mf}(\sigma) \approx \text{mf}(\tau)$. By Theorem 12.1.20, this entails $\sigma \approx \tau$ as well – from this follows that $\Downarrow\text{-Unf}(-)$ is faithful.

Full. Now, consider $\sigma : \underline{A}$ an innocent well-bracketed alternating strategy on \underline{A} . For every non-empty even-length P-view s of σ , choose a copy index $\#s \in \mathbb{N}$ so that this is injective. We then specify a meager causal innocent strategy with events

$$|\sigma| = \{s \in \sigma \mid s \text{ non-empty P-view}\},$$

with causality the prefix ordering, no conflict, trivial symmetries. We set the display map by induction: for odd-length sa^+b^- , we set it to the unique event b' with $\text{lbl}(b') = b$, $\text{pred}(b') = \partial_\sigma(sa^+)$, and $\text{ind}(b') = 0$ – this event exists by *wide*. For even-length $s = s_1a^-s_2b^+$ where b^+ points to a^- , we set the display to the unique event b' with $\text{lbl}(b') = b$, $\text{pred}(b') = \partial_\sigma(s_1a^-)$ and $\text{ind}(b') = \#s$. Altogether, it is easily verified that this defines a meager sequential innocent strategy on A and that $(\hat{\cdot})(\partial_\sigma([s]_\sigma)) = s$. Hence the P-views of $\Downarrow\text{-Unf}_i(\sigma^!)$ are exactly those of σ by Proposition 11.3.6, thus $\Downarrow\text{-Unf}_i(\sigma^!) = \sigma$ by Proposition 3.2.13 – from this follows that $\Downarrow\text{-Unf}(-)$ is full. \square

So that indeed, globular causal strategies are conservative extensions of traditional well-bracketed innocent strategies. Moreover, we established that by showing that the meager form of globular strategies specializes to the meager form of Section 3.2.4.

12.1.4 Finite Tests Suffice

Back on track, we now introduce a notion of *finiteness* for parallel innocent strategies:

Definition 12.1.21. Consider A a mixed board, and $\sigma : A$ a parallel innocent strategy.

Then σ is *finite* iff $|\text{mf}(\sigma)|_+ = \{s \in |\text{mf}(\sigma)| \mid \text{pol}(s) = +\}$ is finite.

In that case, the *size* of σ is the cardinal of $|\text{mf}(\sigma)|_+$.

As in the case of traditional innocent strategies, our forthcoming definability result will concern only *finite* parallel innocent strategies. So that finite definability does entail intensional full abstraction, we must show that if two strategies can be distinguished by a parallel innocent strategy, then they can be distinguished by a *finite* one.

Any parallel innocent strategy can be approximated by finite ones:

Proposition 12.1.22. Consider A a mixed board, and $\sigma : A$ a parallel innocent. Writing

$$D = \{\tau : A \mid \tau \text{ finite} \ \& \ \tau \triangleleft \sigma\}$$

yields a directed set of finite parallel innocent strategies, such that $\sigma = \vee D$.

Proof. By definition, $\vee D \triangleleft \sigma$. For the other direction, pick $m \in \sigma$. Considering

$$|\tau| = \{n \in |\sigma| \mid \forall p^+ \leq_\sigma n, \text{mf}(p^+) \leq_\sigma \text{mf}(m)\},$$

it is easily verified that inheriting from σ all additional components makes $\tau : A$ a parallel innocent strategy s.t. $\tau \triangleleft \sigma$. Moreover $m \in \tau$, and τ is finite since all positive events of $\text{mf}(\tau)$ are below $\text{mf}(m)$. Altogether $\tau \in D$ and $m \in \tau$, so $m \in \vee D$. \square

The development above concerns strategies on mixed boards, but recall that morphisms in $\rightarrow\text{-Strat}$ are strategies *between* mixed boards: a $\sigma \in \rightarrow\text{-Strat}(A, B)$ is a well-bracketed causal strategy on $!A \vdash B$. Nevertheless the above directly extends: $\sigma \in \rightarrow\text{-Inn}(A, B)$ is **finite** if $\Lambda(\sigma) : A \Rightarrow B$ is. From the above, we get:

Corollary 12.1.23. *Consider A a mixed board, and $\sigma_1, \sigma_2 : A$ parallel innocent.*

If there is $\alpha : !A \vdash \mathbf{U}$ parallel innocent such that

$$\alpha \odot_! \sigma_1 \Downarrow, \quad \alpha \odot_! \sigma_2 \Uparrow,$$

then there is $\alpha' \triangleleft \alpha$ finite and parallel innocent, such that

$$\alpha' \odot_! \sigma_1 \Downarrow, \quad \alpha' \odot_! \sigma_2 \Uparrow.$$

Finally if α is globular (resp. well-bracketed), then so is α' .

Proof. By Proposition 12.1.22, $\Lambda(\alpha) = \vee D$ for D some directed set of finite parallel innocent strategies – automatically innocent, as this is inherited from α . The result then directly follows by continuity of composition of causal strategies; preservation of well-bracketing and globularity are immediate. \square

12.2 Factorization

Our next goal is to prove a definability result for finite globular strategies, but it shall be more complex than in the traditional case. Indeed, unlike in Section 3.2.4, globular strategies have no “first Player move” to reproduce first syntactically. Hence we organize our definability process differently. Its core is a *factorization result* (Corollary 12.2.18): namely, that every finite globular $\alpha : !(\&A_i) \vdash \mathbb{X}$ may be obtained as

$$\alpha \equiv \text{fo}(\alpha) \odot_! \langle x_i \alpha_{k,1} \dots \alpha_{k,p_i} \mid i \in I, k \in K_i \rangle, \quad (12.1)$$

with $\text{fo}(\alpha)$ 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 12.3.1.

We first extract the components mentioned in (12.1): the *first-order substrategy* $\text{fo}(\alpha)$, and the *argument substrategies* $\alpha_{k,j}$. We use as illustration the meager strategy with typical linear configurations in Figure 12.5. 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




Figure 12.5: A finite meager globular strategy

$$f : \mathbb{B} \rightarrow \mathbb{U}, g : (\mathbb{U} \rightarrow \mathbb{B}) \rightarrow \mathbb{B} \vdash$$

$$\text{let} \left(\begin{array}{l} x = f \text{ tt} \\ y = g (\lambda z^{\mathbb{U}}. \text{let} \left(\begin{array}{l} u = z \\ v = g \perp \end{array} \right) \text{ in } v) \end{array} \right) \text{ in (if } y \text{ then (f ff; tt) else ff)} : \mathbb{B}$$

Figure 12.6: Factorization and definability for Figure 12.5

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 12.5, the primary questions are q_0^+ , q_1^+ and 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 12.5 there are three, respectively prompted by (i.e. causally depending on) q_0^- , q_2^- and q_1^- and colored accordingly.

The strategy of Figure 12.5 is exactly definable; as illustration we show the term in Figure 12.6, with subterms colored so as to match the four components of the strategy².

The proof of factorization is organized as follows: in Section 12.2.1 we extract the first-order part, in Section 12.2.2, we extract the argument substrategies, and in Section 12.2.3, we recombine them to obtain the original strategy back.

12.2.1 The Flow Substrategy

Shallow substrategy. Consider A a type, and $\alpha : \llbracket A \rrbracket$ a finite globular strategy.

Necessarily A has form $A_1 \rightarrow \dots \rightarrow A_n \rightarrow \mathbb{X}$ for $A_i = A_{i,1} \rightarrow \dots \rightarrow A_{i,p_i} \rightarrow \mathbb{X}_i$.

²In the end, our definability process will not quite give the term of Figure 12.6 but a sequential version as we only know how to define first-order strategies in general up to *positional equivalence* – see Section 12.3.1.

Recall \mathbb{X}, \mathbb{X}_i range over ground types, *i.e.* \mathbb{U}, \mathbb{B} and \mathbb{N} . Up to currying, we write

$$\alpha : !(\&_{1 \leq i \leq n} A_i) \vdash \mathbf{X},$$

omitting semantic brackets. We often shorten the left hand side part to $!(\&A_i)$, and reuse A for the board $!(\&A_i) \vdash \mathbf{X}$. Now, we start with the *shallow substrategy*. First:

Definition 12.2.1. Consider an event $m \in \alpha$.

Then m is *shallow* iff $[m]_\alpha$ contains exactly one Opponent question.

In Figure 12.5, the shallow events are exactly those in red. Note that if m is shallow, then the Opponent question in $[m]_\alpha$ is necessarily the initial Opponent question. As one can dive within the A_i s only via an Opponent question, this means that m and its causal history must remain in the *first-order* part of the arena.

The shallow substrategy is obtained by restricting α to shallow events:

Proposition 12.2.2. Set $\text{sh}(\alpha)$ with the shallow events of α , other components inherited.

Then, $\text{sh}(\alpha) : !(\&\mathbf{X}_i) \vdash \mathbf{X}$ is finite globular.

Proof. Write $\text{sh}(A)$ for the board $!(\&\mathbf{X}_i) \vdash \mathbf{X}$. 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 finite globular strategy are direct from α . \square

This captures the strategy induced by the red part of Figure 12.5.

The flow substrategy Eventually we wish to reconstruct α from its first-order and argument substrategies using the composition mechanisms available on causal strategies. But we cannot hope to do that merely with $\text{sh}(\alpha)$. Indeed, two Player questions in $\text{sh}(\alpha)$ playing in the same component \mathbf{X}_i must receive the same argument substrategy if α is to be reassembled compositionally – but there is no reason why two Player questions in the same component should always have the same argument substrategy!

Thus rather than working with $\text{sh}(\alpha)$, we need to relabel it to send distinct (non-symmetric) Player questions to distinct components. First we define:

Definition 12.2.3. A *primary question* of α is any $q^{\mathcal{Q},+} \in \text{mf}(\text{sh}(\alpha))$. We write \mathcal{Q} for the set of primary questions, and \mathcal{Q}_i for the primary questions displaying to \mathbf{X}_i .

By construction, $\mathcal{Q} = \uplus_{1 \leq i \leq n} \mathcal{Q}_i$ – as observed above, \mathcal{Q} is finite. A $m^{\mathcal{Q},+} \in \text{sh}(\alpha)$ might not be a primary question, but it must necessarily be symmetric to $\text{lbl}(m) \in \mathcal{Q}$ a primary question (as its *label*, *i.e.* its representative in the meager form).

Definition 12.2.4. The *flow substrategy* of α , written

$$\text{flow}(\alpha) : \bigotimes_{1 \leq i \leq n} \bigotimes_{q \in \mathcal{Q}_i} !\mathbf{X}_i \vdash \mathbf{X}$$




Figure 12.7: Configurations of the flow substrategy for Figure 12.5

is obtained as $\text{sh}(\alpha)$ with the display map $\partial_{\text{sh}(\alpha)}$ replaced with:

$$\begin{array}{ll}
 \partial_{\text{flow}(\alpha)}(m^{\mathcal{Q},-}) = (2, a) & \text{if } \partial_{\text{sh}(\alpha)}(m^{\mathcal{Q},-}) = (2, a) \\
 \partial_{\text{flow}(\alpha)}(m^{\mathcal{A},+}) = (2, a) & \text{if } \partial_{\text{sh}(\alpha)}(m^{\mathcal{A},+}) = (2, a) \\
 \partial_{\text{flow}(\alpha)}(m^{\mathcal{Q},+}) = (1, (i, (\text{bl}(m), (j, a)))) & \text{if } \partial_{\text{sh}(\alpha)}(m^{\mathcal{Q},+}) = (1, (j, (i, a))) \\
 \partial_{\text{flow}(\alpha)}(m^{\mathcal{A},-}) = (1, (i, (j(\text{bl}(m)), (j, a)))) & \text{if } \partial_{\text{sh}(\alpha)}(m^{\mathcal{A},-}) = (1, (j, (i, a))),
 \end{array}$$

i.e. sending each $q \in \mathcal{Q}$ and its answers to the copy of \mathbf{X}_i specified by indices i, q .

It is a finite globular strategy (though not a morphism in $\rightarrow\text{-Strat}$, but the same definition applies nonetheless). Figure 12.7 shows typical configurations of the flow substrategy for Figure 12.5, tagging each component by a primary question.

12.2.2 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. A slight complication comes from the fact that because answers in the shallow substrategy can be replicated, α actually comprises many copies of each argument substrategy – we solve that by temporarily only considering the part of α where shallow answers have copy index 0:

Definition 12.2.5. An event $m \in \alpha$ is **canonical** iff for all shallow $n^{\mathcal{A},-} \leq_\alpha m$, $\text{ind}(\partial_\alpha(n)) = 0$. We write $\|\alpha\|$ the set of canonical events of α .

A primary question is always canonical; so are all events in $\text{mf}(\alpha)$. Now, fix a primary question $q \in \mathcal{Q}_i$. It displays to an initial event in A_i , which is:

$$!A_{i,1} \multimap \dots \multimap !A_{i,p_i} \multimap \mathbf{X}_i.$$

Argument substrategies are accessed by Opponent questions pointing to (moves symmetric to) primary questions. Up to symmetry, there are p_i distinct 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 a finite globular strategy:

$$\alpha_{q,j} : !(\&A_i) \vdash !A_{i,j}.$$

Extracting events. To do this, we assign to all events of α tags, as follows:

Definition 12.2.6. Consider $m \in \|\alpha\|$. We write:

$$\begin{aligned} m \in \text{sh}(\alpha) &\Leftrightarrow m \text{ is shallow,} \\ m \in \alpha_{q,j} &\Leftrightarrow \text{there is } q \rightarrow_\alpha n^{\mathcal{Q},-} \text{ in component } j, \text{ such that } n \leq_\alpha m, \end{aligned}$$

where in the second clause, $q \in \mathcal{Q}_i$ is a primary question and $1 \leq j \leq p_i$.


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 subtle and involves all our structural constraints on strategies. For instance, without *globular*, in Figure 10.19 the move q_2^+ would be tagged for *two* distinct argument sub-strategies.

Lemma 12.2.7. Every $m \in \|\alpha\|$ receives exactly one tag following Definition 12.2.6.

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, then m is shallow by definition. Assume there are at least two. Take $n^{\mathcal{Q},-} \leq_\alpha m$ minimal such that it is a non-initial Opponent question. Then its immediate predecessor must be some primary question $q \in \mathcal{Q}_i$; 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 $m \in \text{sh}(\alpha)$. But m could be in two argument sub-strategies: 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' . We draw:



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 *pre-innocence*, m_1 and n_1 are positive. Moreover, since the board is forestal and alternating, it is immediate from Lemma 6.1.16 that m_{k+1} is positive, and m_k, n_p positive. Altogether we have a *globule*, so X and Y are complete by globularity – and moreover the diagram has the shape as in (10.2), and in particular q

$$!((B \rightarrow U) \& ((U \rightarrow B) \rightarrow B)) \vdash B$$

$$!((B \rightarrow U) \& ((U \rightarrow B) \rightarrow B)) \vdash ! (U \rightarrow B)$$

Figure 12.8: Ev. $|\alpha_{q^+,1}|$ of Figure 12.5Figure 12.9: Its reassignment in $\alpha_{q^+,1}$

is answered in ρ . Additionally as observed in (10.2), writing $q = m_i, m_{i+1}$ must answer q . But since $m \in \alpha_{q,j}$, $q \rightarrow_\alpha n^{Q,-} \leq_\alpha m$ 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, if j and j' were distinct, there would be necessarily distinct $q \rightarrow_\alpha n_1^{Q,-}$ in component j and $n_2^{Q,-}$ in component j' such that $n_1, n_2 \leq_\alpha m$, but this is in direct contradiction with *pre-innocence*. Thus $j = j'$, concluding the proof. \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 have 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.

Assigning display map. The main challenge in extracting the argument sub-strategies is to define what shall be their new display maps. Indeed, ∂_α still displays events in $\alpha_{q,j}$ to moves in $!(\&A_i) \vdash X$, whereas $\partial_{\alpha_{q,j}}$ must display then to $!(\&A_i) \vdash !A_{i,j}$.

For this, we must give special care to those events that correspond to the new $!A_{i,j}$ component on the right hand side, reindexing them as illustrated in Figures 12.8 and 12.9. For this we split $|\alpha_{q,j}|$ in two subsets: on the one hand, some events depend *statically*, i.e. with respect to \leq_A (through ∂_α) on the primary question q – in Figure 12.8, those are $q_0^-, q_0^+, \checkmark_0^-$ and b_0^+ . On the other hand, the remaining events must follow from new calls to variables in the context – in Figure 12.8, those are q_3^+ and b_0^- . 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 12.9.

It will be helpful to have notations for the canonical embeddings of the set of moves $|A_i|$ and $|A_{i,j}|$ into $|A|$. More precisely, for each primary question $q \in \mathcal{Q}_i$, we write

$$\iota_q(-) : |A_i| \rightarrow |A|$$

the injection adding the sequence of tags addressing A_i within A , originating from the tagged disjoint unions in the board constructions – in particular, ι_q maps the initial move of A_i to $\partial_\alpha(q)$. Likewise, $\iota_{q,j} : |A_{i,j}| \rightarrow |A|$ addresses the j -th argument of q . Then:

Definition 12.2.8. We define a display map for $\alpha_{q,j}$ by setting, for $m \in \alpha_{q,j}$:

$$\begin{aligned} \partial_{\alpha_{q,j}}(m) &= \iota_{\rho}(a) & \text{if } \partial_\alpha(m) = \iota_{q,j}(a), \\ \partial_{\alpha_{q,j}}(m) &= \partial_\alpha(m) & \text{otherwise,} \end{aligned}$$

where $\iota_\ell(a) = (1, a)$ and $\iota_{\rho}(a) = (2, a)$.

Altogether, this lets us extract $\alpha_{q,j}$ as intended:

Proposition 12.2.9. 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 $\alpha, \mathcal{S}(\alpha_{q,j}) = \{\theta \cap |\alpha_{q,j}|^2 \mid \theta \in \mathcal{S}(\alpha)\}$, and $\partial_{\alpha_{q,j}}$ in Definition 12.2.8.

Then, $\alpha_{q,j} : !(\&A_i) \vdash !A_{i,j}$ is a globular finite strategy.

Proof. A lengthy but routine verification. \square

Removing the bang. But we must obtain a morphism in $\rightarrow\text{-Glob}$, and thus eliminate the bang on the right hand side. This is done by composition with *dereliction*, i.e.

$$\alpha_{q,j}^\bullet = \text{der}_{A_{i,j}} \odot \alpha_{q,j} : !(\&A_i) \vdash A_{i,j},$$

it only remains to show that this satisfies the conditions which will later on allow us to proceed from $\alpha_{q,i}^\bullet$ with the inductive definability argument:

Proposition 12.2.10. Consider $q \in \mathcal{Q}_i$ a primary question, and $1 \leq j \leq p_i$.

Then, $\alpha_{q,j}^\bullet$ is finite globular with size strictly lesser than α .

Proof. Using Proposition 7.3.1, it is easy that $\alpha_{q,j}^\bullet$ is positively isomorphic to the strategy obtained from $\alpha_{q,j}$ by restricting the initial Opponent question to copy index 0. This also informs an injection of $\text{mf}(\alpha_{q,j}^\bullet)$ into $\text{mf}(\alpha)$ not reaching the primary question q , from which follows the announced size constraint. \square

Note that $\alpha_{q,j}$ is easily recovered from $\alpha_{q,j}^\bullet$, as we have

$$\alpha_{q,j} \approx (\text{der}_{A_{i,j}} \odot \alpha_{q,j})^\dagger = (\alpha_{q,j}^\bullet)^\dagger \quad (12.2)$$

via the “bang lemma” (Lemma 10.3.7) and by definition of $\alpha_{q,j}^\dagger$.

12.2.3 Recomposition

Having extracted from α its flow substrategy and its argument substrategies, our next aim is to prove that α can be reconstructed from those.

Reconstructing α . From the original strategy $\alpha : !(\&A_i) \vdash \mathbf{X}$, we have now extracted the *flow substrategy* and a family of *argument substrategies*:

$$\begin{aligned} \text{flow}(\alpha) &: \bigotimes_{1 \leq i \leq n} \bigotimes_{q \in \mathcal{Q}_i} !\mathbf{X}_i \vdash \mathbf{X} \\ \alpha_{q,j}^\bullet &: !(\&A_i) \vdash A_{i,j} \end{aligned} \quad \text{for each } q \in \mathcal{Q}_i \text{ and } 1 \leq j \leq p_i.$$

For each $q \in \mathcal{Q}_i$, we first form a strategy $\alpha_q : !(\&A_i) \vdash \mathbf{X}_i$ gathering all the arguments for the primary question q . This is constructed from the argument substrategies, using the internal language corresponding to the cartesian closed structure of \rightarrow -**Inn**:

$$\alpha_q = x_1 : A_1, \dots, x_n : A_n \vdash x_i \alpha_{q,1}^\bullet \dots \alpha_{q,p_i}^\bullet : \mathbf{X}_i.$$

Then, we plug each α_q onto the corresponding primary question in the flow substrategy. This is done relying on the relative Seely category structure of **NTCG-Inn**:

$$\text{recomp}(\alpha) = \text{flow}(\alpha) \odot \left(\bigotimes_{1 \leq i \leq n} \bigotimes_{q \in \mathcal{Q}_i} \alpha_q^\dagger \right) \odot \delta_{\&A_i} : !(\&A_i) \vdash \mathbf{X},$$

where for B a strict board 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 .

Positions of $\text{recomp}(\alpha)$. Recall that the *positional equivalence* \equiv , studied in Section 10.4, consists in listing the zero-payoff symmetry classes of configurations reached by (+-covered configurations of) a visible strategy. We expect that $\text{recomp}(\alpha) \approx \alpha$, but we shall only prove $\text{recomp}(\alpha) \equiv \alpha$ – this is simpler as positions compose relationally. We shall make use of the syntax for positions on $--$ -boards introduced in Figure 10.17.

We now analyse the positions of the recomposition (12.2.3). We start with:

Lemma 12.2.11. *Consider B a strict $--$ -board.*

Then, the positions (δ_B) of $\delta_B : !B \vdash (!B)^{\otimes n}$ are exactly those

$$\left(\sum_{1 \leq i \leq n} x_i \vdash \bigotimes_{1 \leq j \leq n} x_j \right) \in (!B \vdash (!B)^{\otimes n})$$

where $x_i \in (!B)$ for all $1 \leq i \leq n$.

Proof. By a direct variation of Lemma 6.4.4. □

Next we analyse the positions of α_q for $q \in \mathcal{Q}_i$:

Lemma 12.2.12. *For any $q \in \mathcal{Q}_i$, the positions (α_q) are exactly those of the form*

$$\left([(i, y_{q,1} \multimap \dots \multimap y_{q,p_i} \multimap v_q)] + \sum_{1 \leq j \leq p_i} x_{q,j} \right) \vdash v_q \in (!(\&A_i) \vdash \mathbb{X}_i)$$

where for each $1 \leq j \leq p_i$, $(x_{q,j} \vdash y_{q,j}) \in (\alpha_{q,j})$, and for $v_q \in (\mathbb{X}_i)$.

Proof. From the laws of relative Seely categories, α_q is positively isomorphic to

$$\begin{aligned} !(\&A_i) &\xrightarrow{\delta} \otimes_{p_i+1} !(\&A_i) \\ &\quad \downarrow (\otimes_{1 \leq j \leq p_i} (\alpha_{q,j}^*)^\dagger) \otimes \llbracket x_i \rrbracket \\ (\otimes_{1 \leq j \leq p_i} !A_{i,j}) \otimes A_i &\xrightarrow{\text{ev}} \mathbb{X}_i \end{aligned}$$

in **NTCG**. The lemma follows from Corollary 10.4.15, Lemma 12.2.11, (12.2) and characterisations analogous to Lemma 12.2.11 for other copycat-like strategies. \square

From all those, we may characterise the positions of $\text{recomp}(\alpha)$ as

Corollary 12.2.13. *The positions of $\text{recomp}(\alpha)$ are exactly those of the form*

$$\sum_{1 \leq i \leq n} \sum_{q \in \mathcal{Q}'_i} \left([(i, y_{q,1} \multimap \dots \multimap y_{q,p_i} \multimap v_q)] + \sum_{1 \leq j \leq p_i} x_{q,j} \right) \vdash v \in (!(\&A_i) \vdash \mathbb{X})$$

where for all $1 \leq i \leq n$, \mathcal{Q}'_i is a subset of \mathcal{Q}_i , and:

$$\left(\otimes_{1 \leq i \leq n} \otimes_{q \in \mathcal{Q}'_i} v_q \vdash v \right) \in (\text{flow}(\alpha)), \quad ((x_{q,j} \vdash y_{q,j}) \in (\alpha_{q,j}))_{q \in \mathcal{Q}'_i, 1 \leq j \leq p_i}.$$

Proof. Direct from Lemmas 12.2.11, 12.2.12 and Corollary 10.4.15. \square

Positions of α . Next, we must similarly characterise the positions of α .

We call **stopping configurations** of α those $x \in \mathcal{C}^+(\alpha)$ such that $\kappa(\partial_\alpha(x)) = 0$ (i.e. those configurations that contribute a point to the relational collapse), and such that all Opponent answers in x have copy index 0. Write $\text{stop}(\alpha)$ for this set; our first observation is that the relational collapse of α is entirely spanned by stopping configurations – the restriction to the copy index of Opponent answers has no consequence.

Lemma 12.2.14. *Consider $x_\Gamma \vdash x_\mathbb{X} \in (\alpha)$.*

Then, there is a stopping configuration $x \in \text{stop}(\alpha)$ such that $\partial_\alpha x \in x_\Gamma \vdash x_\mathbb{X}$.

Proof. By definition, there is $y \in \mathcal{C}^+(\alpha)$ such that $\partial_\alpha y \in x_\Gamma \vdash x_\mathbb{X}$. In particular, $\kappa(\partial_\alpha y) = 0$. By definition of payoff on ground types, this entails that all questions in y have *exactly one answer*. We construct $z \in \mathcal{C}(A)$ from $\partial_\alpha(y)$ using *wide*, by changing the copy index of every negative answer to 0 – by *--transparent*, we obtain a symmetry

$$\theta : \partial_\alpha(x) \cong_A^- z,$$

and applying Lemma 7.2.6 we get $y \cong_{\alpha} x$ and $\varphi : \partial_{\alpha} y \cong_A^+ z$. By $+$ -transparent, φ preserves the copy indices of negative events, thus the indices of negative answers in $\partial_{\alpha} y$ is 0. Finally, as $x \cong_{\alpha} y$ we have $\partial_{\alpha} y \in x_{\Gamma} \vdash x_{\Sigma}$, and y is a stopping configuration. \square

Restricting to stopping configurations is important, because events of stopping configurations are *canonical* in the sense of Definition 12.2.5 (i.e. the shallow Opponent answers on which they depend have copy index 0) – and argument sub-strategies were extracted from the canonical events of α , following Definition 12.2.6. Thanks to this, we shall be able to partition any $x \in \text{stop}(\alpha)$ into stopping configurations from the flow substrategy, and the argument sub-strategies. This partitionning is the key argument to link the positions of α with the positions of its sub-strategies.

Let us perform this partitioning. We define components of the partition as

$$\begin{aligned} \text{sh}(x) &= \{m \in x \mid m \in \text{sh}(\alpha)\} \\ x_{q,j} &= \{m \in x \mid m \in \alpha_{q,j}\}, \quad (\text{for } q \in \mathcal{Q}_i \text{ and } 1 \leq j \leq p_i) \end{aligned}$$

additionally we write $\mathcal{Q}^x = \mathcal{Q} \cap x$, and \mathcal{Q}_i^x likewise. We have:

Lemma 12.2.15. *Consider $x \in \text{stop}(\alpha)$.*

Then $\text{sh}(x) \in \text{stop}(\text{sh}(\alpha))$, $x_{q,j} \in \text{stop}(\alpha_{q,j})$ and x partitions as

$$x = \text{sh}(x) \uplus \left(\bigsqcup_{1 \leq i \leq n} \bigsqcup_{q \in \mathcal{Q}_i^x} \bigsqcup_{1 \leq j \leq p_i} x_{q,j} \right),$$

moreover this decomposition is compatible with display maps, in the sense that

$$\partial_{\alpha}(x) = \left(\bigsqcup_{1 \leq i \leq n} \bigsqcup_{q \in \mathcal{Q}_i^x} \left(\left[\iota_q(z_{q,1} \multimap \dots \multimap z_{q,p_i} \multimap v_q) \right] \uplus \left[\bigsqcup_{1 \leq j \leq p_i} \iota_{\ell}(y_{q,j}) \right] \right) \right) \uplus \iota_r(v),$$

where we have, for all $1 \leq i \leq n$, $q \in \mathcal{Q}_i^x$ and $1 \leq j \leq p_i$:

$$\partial_{\text{flow}(\alpha)}(x) = \iota_{\ell} \left(\bigotimes_{1 \leq i \leq n} \bigotimes_{q \in \mathcal{Q}_i^x} v_q \right) \uplus \iota_r(v) \quad \partial_{\alpha_{q,j}}(x_{q,j}) = \iota_{\ell}(y_{q,j}) \uplus \iota_r(z_{q,j}).$$

Proof. From Lemma 12.2.7, lengthy but straightforward verifications. \square

Reciprocally, any compatible collection of stopping configurations from the various substrategies may be assembled into a stopping configuration of α :

Lemma 12.2.16. *For $x \in \text{stop}(\text{sh}(\alpha))$ and $(x_{q,j} \in \text{stop}(\alpha_{q,j}))_{1 \leq i \leq n, q \in \mathcal{Q}_i^x, 1 \leq j \leq p_i}$,*

$$x \uplus \left(\bigsqcup_{1 \leq i \leq n} \bigsqcup_{q \in \mathcal{Q}_i^x} \bigsqcup_{1 \leq j \leq p_i} x_{q,j} \right) \in \text{stop}(\alpha).$$

Proof. Lengthy but straightforward verifications. \square

From this we may finally conclude the proof of factorization:

Corollary 12.2.17. *The strategies α and $\text{recomp}(\alpha)$ are positionally equivalent.*

Proof. Altogether, Lemmas 12.2.15 and 12.2.16 ensure that stopping configurations of α are in one-to-one correspondence with the data of stopping configurations $x \in \text{stop}(\text{sh}(\alpha))$ and families $(x_{q,j})_{1 \leq i \leq n, q \in \mathcal{Q}_i^x, 1 \leq j \leq p_i}$, in a way compatible with the display maps as stated in Lemma 12.2.15. Taking symmetry classes, we derive the same characterisation of complete positions of α as that of $\text{recomp}(\alpha)$ in Corollary 12.2.13. \square

Syntactic factorization. Finally, we reformulate this relying on the cartesian closed structure only. The **first-order substrategy** $\text{fo}(\alpha) \in \text{NTCG}_!(\&_{1 \leq i \leq n} \&_{q \in \mathcal{Q}_i} \mathbf{X}_i, \mathbf{X})$ is obtained in the obvious way from $\text{flow}(\alpha)$ using the relative Seely category structure. Using Corollary 12.2.17, Proposition 12.2.10, and laws of a relative Seely category:

Corollary 12.2.18. *Any globular strategy $\alpha : !(\&A_i) \vdash \mathbf{X}$ factors as*

$$\alpha \equiv \text{fo}(\alpha) \odot_! \langle x_i \alpha_{q,1}^* \dots \alpha_{q,p_i}^* \mid 1 \leq i \leq n, q \in \mathcal{Q}_i \rangle,$$

where $\text{fo}(\alpha) : !(\&\mathbf{X}_i) \vdash \mathbf{X}$ and $\alpha_{q,j}^* : !(\&A_i) \vdash A_{i,j}$ are globular.

If α is finite, so are $\text{fo}(\alpha)$ and $\alpha_{q,j}^$, the latter with size strictly lesser than that of α .*

12.3 Finite Definability and Full Abstraction

12.3.1 First-Order Definability

Assuming we are in the process of an inductive definability procedure, we may assume that we have terms for the argument sub-strategies $\alpha_{q,j}^*$. Thus, it remains to define

$$\text{fo}(\alpha) : !(\&\mathbf{X}_i) \vdash \mathbf{X}$$

the first-order substrategy.

Strict definability. This naturally asks the question: does finite definability hold for first-order globular strategies? It is fairly easy to see that the answer is negative:

Proposition 12.3.1. *The strategy with meager form pictured in Figure 12.10 is finite globular, but not definable in PCF_{//}.*

Sketch. By rewriting means, one may establish that if it was definable in PCF_{//}, then it would be definable just via parallel compositions of variables. However, such terms composed only of parallel compositions and variables are easily seen to yield series-parallel causal shapes; thus the example in Figure 12.10 is undefinable. \square

For strategies on the game $!(\&\mathbf{U}) \vdash \mathbf{U}$, it seems that we may get finite definability up to positive isomorphism by restricting globular strategies further to those that have a series-parallel causal shape. But for more complex datatypes, requiring a series-parallel causal shape is not enough to get back definability: an example phenomenon appears in


$!(U \& U \& U \& U \& U) \vdash U$


Figure 12.10: An undefinable strategy


 $!(U \& U \& B \& U \& U) \vdash U$


Figure 12.11: Dynamic causal shape

Figure 12.11, where the structure of forking/merging threads is impacted by the values. Though we do not have a proof, it seems clear that this is undefinable: in a term as in

$$\mathbf{let} \begin{pmatrix} x_1 = M_1 \\ x_2 = M_2 \\ x_3 = M_3 \end{pmatrix} \mathbf{in} N$$

then the first move in N will depend on the values obtained for x_1, x_2 and x_3 : there is no way to indicate that *e.g.* if x_2 is true, then N should not wait for x_1 .

Altogether, we have not been able to characterise the globular strategies definable in PCF_{||}. A natural option would be to reify the first-order globular strategies, *i.e.* to consider them as syntax and add a new syntactic construct for every first-order globular strategy. This seems interesting, because first-order globular strategies are indeed relevant from a computational point of view; however it seems hard to give then an appealing syntactic theory. Instead, we shall weaken our objective and only perform definability *up to positional equivalence* – which will suffice for our purposes³.

12.3.2 Positional First-Order Definability

Moving from positive isomorphism to positional equivalence is convenient, because positional equivalence may not respect the evaluation order. This lets us construct a purely sequential version of any finite globular first-order strategy.

From now on, fix a first-order finite globular strategy $\alpha : !(\&\mathbb{X}_i) \vdash \mathbb{X}$.

Necessary primary question. Once prompted, α may launch many computations in parallel, via *minimal primary questions* (*i.e.* that only depend on the Opponent initial

³Of course, the factorization result of Corollary 12.2.18 is already up to positional equivalence, but this is mostly for convenience: factorization should hold up to positive isomorphism.

question). The idea is that our sequential definability procedure will choose any of those (the choice does not matter up to positional equivalence), and evaluate it first. However, we must show that all these primary questions are actually *necessary*, in the sense that any stopping configuration necessarily includes this primary question, answered.

Lemma 12.3.2. *Consider $\alpha \in \rightarrow\text{-Glob}(\&\mathbb{X}_i, \mathbb{X})$ finite globular.*

Then for each minimal primary question q and $x \in \text{stop}(\alpha)$, we have $q \in x$.

Proof. Consider $x \in \text{stop}(\alpha)$ and seeking a contradiction, assume $q \notin x$. By *causal determinism*, q is compatible with x . Consider $t \in \mathcal{L}(x)$ any linearization of x , yielding

$$s = (\hat{\cdot})(\partial_{\Lambda(\alpha)}(t)) \in \mathfrak{O}\text{-Unf}(\alpha),$$

it is straightforward that it is (logically) well-bracketed. Now, α is a morphism in $\rightarrow\text{-Strat}$, and must thus be well-bracketed in the sense of Definition 9.3.2. Therefore,

$$sa^+ = (\hat{\cdot})(\partial_{\Lambda(\alpha)}(tq)) \in \mathfrak{O}\text{-Unf}(\alpha)$$

should be (logically) well-bracketed. But it is not, since q points to the initial Opponent question which must already be answered in x : a contradiction, thus $q \in x$. \square

Thus we may safely begin the sequentialization with any minimal primary question: by this lemma, we know that it will appear in all stopping configurations.

Residual. Choose $q \in \mathcal{Q}_{i_0}$ minimal. As q appears in every stopping configuration, it is safe to first make a call to x_{i_0} , then branch 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 12.3.3. *There is $\alpha_{(q)} \in \rightarrow\text{-Glob}(\&\mathbf{X}_i, \mathbf{X}_{i_0}, \mathbf{X})$ with the same ess as α and differing only via its display map, such that:*

- (1) *for $x \vdash w \in \langle \alpha_{(q)} \rangle$, then $x = x' + [(n+1, v)]$ such that $x' + [(i_0, v)] \vdash w \in \langle \alpha \rangle$,*
- (2) *for $x \vdash w \in \langle \alpha \rangle$, then $x = x' + [(i_0, v)]$ such that $x' + [(n+1, v)] \vdash w \in \langle \alpha_{(q)} \rangle$.*

Proof. The strategy $\alpha_{(q)}$ has the same components as α , with display map sending q and its answers to the new component. The characterisation of positions is direct. \square

Note that in the second clause, the decomposition $x = x' + [(i_0, v)]$ is not unique: there may be several values in component i , but one of them yields a position of $\alpha_{(q)}$.

We set the residual $\alpha_{(qv)}$ as $\alpha_{(qv)} = \alpha_{(q)} \odot_! \langle \text{der}_{\&\mathbf{X}_i}, v \rangle : !(\&\mathbf{X}_i) \vdash \mathbf{X}$ writing $v : !(\&\mathbf{X}_i) \vdash \mathbf{X}_{i_0}$ the constant strategy. In order to characterize its positions, we note:

Lemma 12.3.4. *For any $v \in V$, the positions of $\langle \text{der}_{\&\mathbf{X}_i}, v \rangle^!$ are exactly of the form*

$$x \vdash x + p \cdot [(n+1, v)] \quad \in \quad \langle !(\&\mathbf{X}_i) \vdash !((\&\mathbf{X}_i) \& \mathbf{X}_{i_0}) \rangle,$$

where $p \cdot [(n+1, v)]$ denotes the p -fold sum, and for any $x \in \langle !(\&\mathbf{X}_i) \rangle$ and $p \in \mathbb{N}$.

Proof. Straightforward verification. \square

Using this lemma, we link the positions of α and $\alpha_{(qv)}$.

Lemma 12.3.5. *Consider positions $x \in \{!(\&X_i)\}$ and $w \in (X)$.*

Then, $x \vdash w \in (\alpha)$ iff $x = x' + [(i_0, v)]$ with $x' \vdash w \in (\alpha_{(qv)})$.

Proof. By definition, we have $\alpha_{(qv)} = \alpha_{(q)} \odot \langle \text{der}_{\&X_i}, v \rangle^!$, so by Corollary 10.4.15,

$$(\alpha_{(qv)}) = ((\alpha_{(q)})) \odot (\langle \text{der}_{\&X_i}, v \rangle^!).$$

The lemma directly follows by Lemmas 12.3.3 and 12.3.4. \square

Reconstruction. We now rebuild a term. We have $\alpha_{(qv)}$ a globular finite strategy with size strictly lesser than that of α . By induction hypothesis, there is a term

$$x_1 : \mathbb{X}_1, \dots, x_n : \mathbb{X}_n \vdash N_{(qv)} : \mathbb{X}_{i_0},$$

for each $v \in V$, such that $\llbracket N_{(qv)} \rrbracket \equiv \alpha_{(qv)}$. We set $x_1 : \mathbb{X}_1, \dots, x_n : \mathbb{X}_n \vdash M : \mathbb{X}$ as

$$\begin{array}{ll} \text{case } x_{i_0} \text{ of} & \text{let } x = x_{i_0} \text{ in} \\ v_1 \mapsto N_{(qv_1)} & \text{if } x =_{\mathbb{X}} v_1 \text{ then } N_{(qv_1)} \\ v_2 \mapsto N_{(qv_2)} & \text{else if } x =_{\mathbb{X}} v_2 \text{ then } N_{(qv_2)} \\ \dots & \dots \\ v_p \mapsto N_{(qv_p)} & \text{else if } x =_{\mathbb{X}} v_p \text{ then } N_{(qv_p)} \\ & \text{else } \perp \end{array} \quad \stackrel{\text{def}}{=}$$

where $V = \{v_1, \dots, v_p\}$, using the syntactic sugar introduced in Section 2.2.3. 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 $\text{let } x = x_{i_0} \text{ in } M'$.

Analysis. 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 by an analysis of the positions of $\llbracket M' \rrbracket$.

Lemma 12.3.6. *Consider positions $x \in \{!(\&X_i)\}$ and $w \in (X)$.*

Then, $x \vdash w \in (\llbracket N_{(qv)} \rrbracket)$ iff there is $p \in \mathbb{N}$ such that $x + p \cdot [(n+1, v)] \vdash w \in (\llbracket M' \rrbracket)$.

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 $\llbracket M' \rrbracket \odot_! \langle \text{der}, v \rangle = \llbracket N_{(qv)} \rrbracket$. The claim then follows by Corollary 10.4.15 and Lemma 12.3.4. \square

It remains to take the interpretation of the **let** construction into account. Recall that

$$\llbracket M \rrbracket = \text{let}_{\mathbb{X}_{i_0}, \mathbb{X}} \odot_! \langle \pi_{i_0}, \Lambda^!(\llbracket M' \rrbracket) \rangle \quad : \quad !(\&X_i) \vdash X$$

where $\text{let}_{\mathbb{X}_{i_0}, \mathbb{X}} : !(\mathbb{X}_{i_0} \& (!\mathbb{X}_i \multimap \mathbb{X})) \vdash \mathbb{X}$. The positions of $\text{let}_{\mathbb{X}_{i_0}, \mathbb{X}}$ are as follows:

Lemma 12.3.7. *The positions of $\text{let}_{\mathbf{X}_{i_0}. \mathbf{X}}$ are exactly those of the form*

$$[(1, v)] + [(2, ((p \cdot [v]) \multimap w))] \vdash w \quad \in \quad ((!(\mathbf{X}_{i_0} \& (!\mathbf{X}_i \multimap \mathbf{X})) \vdash \mathbf{X}))$$

for $v \in (\mathbf{X}_{i_0})$, $w \in (\mathbf{X})$, and $p \in \mathbb{N}$.

Proof. A direct analysis of positions reached by stopping configurations of $\text{let}_{\mathbf{X}_{i_0}. \mathbf{X}}$. \square

We can now wrap up, showing that $\llbracket M \rrbracket$ has the same non-empty positions as α .

Lemma 12.3.8. *Consider $x \in (\&\mathbf{X}_i)$ and $w \in (\mathbf{X})$.*

Then, $x \vdash w \in (\llbracket M \rrbracket)$ iff $x = x' + [(i_0, v)]$ with $x' \vdash w \in (\llbracket N_{(qv)} \rrbracket)$.

Proof. By Lemmas 12.3.6 and 12.3.7. \square

Note that this exactly matches the characterisation of positions of α in Lemma 12.3.5. As $\alpha_{(qv)} \equiv \llbracket N_{(qv)} \rrbracket$ by induction hypothesis, it follows that $\llbracket M \rrbracket \equiv \alpha$. Summing up, we have proved finite first-order definability up to positional equivalence:

Theorem 12.3.9. *For $\alpha \in \multimap\text{-Glob}((\&\mathbf{X}_i), \mathbf{X})$ finite, 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 $\llbracket M \rrbracket \equiv \alpha$.

12.3.3 Finite Globular Definability

We may now conclude the proof of finite definability.

Theorem 12.3.10. *Let $\Gamma \vdash A$ be a PCF typing judgment, $\alpha \in \multimap\text{-Glob}(\llbracket \Gamma \rrbracket, \llbracket A \rrbracket)$ finite.*

Then, there is $\Gamma \vdash M : A$ such that $\llbracket M \rrbracket \equiv \alpha$.

Proof. Up to currying, $\alpha : !(\&_{1 \leq i \leq n} A_i) \vdash \mathbf{X}$, 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 12.2.18,

$$\alpha \equiv \text{fo}(\alpha) \odot_! \langle x_i \alpha_{q,1}^\bullet \dots \alpha_{q,p_i}^\bullet \mid 1 \leq i \leq n, q \in \mathcal{Q}_i \rangle,$$

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 globular finite strategy of size strictly smaller than α . By induction hypothesis, there is

$$x_1 : A_1, \dots, x_n : A_n \vdash N_{q,j} : A_{i,j}$$

s.t. $\llbracket N_{q,j} \rrbracket \equiv \alpha_{q,j}$. Let us write $\mathcal{Q}_i = \{q_{i,1}, \dots, q_{i,k_i}\}$. By Theorem 12.3.9 there is

$$x_{q_{1,1}} : \mathbb{X}_1, \dots, x_{q_{1,k_1}} : \mathbb{X}_1, \dots, x_{q_{n,1}} : \mathbb{X}_n, \dots, x_{q_{n,k_n}} : \mathbb{X}_n \vdash M_{\text{fo}} : \mathbb{X}$$

such that $\llbracket M_{\text{fo}} \rrbracket \equiv \text{fo}(\alpha)$. Then, we define the term $x_1 : A_1, \dots, x_n : A_n \vdash M : \mathbb{X}$ as

$$x_1 : A_1, \dots, x_n : A_n \vdash M_{\text{fo}}[x_i N_{q_{i,1},1} \dots N_{q_{i,l},p_i} / x_{q_{i,l}}] : \mathbb{X}.$$

Then we may finally compute

$$\begin{aligned}
\llbracket M \rrbracket &= \llbracket M_{\text{fo}} \rrbracket \odot_! \langle \llbracket x_i N_{q_{i,l},1} \dots N_{q_{i,l},p_i} \rrbracket \mid 1 \leq i \leq n, q_{i,l} \in \mathcal{Q}_i \rangle \\
&\equiv \llbracket M_{\text{fo}} \rrbracket \odot_! \langle x_i \llbracket N_{q,1} \rrbracket \dots \llbracket N_{q,n} \rrbracket \mid 1 \leq i \leq n, q \in \mathcal{Q}_i \rangle \\
&\equiv \text{fo}(\alpha) \odot_! \langle x_i \alpha_{q,1}^\bullet \dots \alpha_{q,p_i}^\bullet \mid 1 \leq i \leq n, q \in \mathcal{Q}_i \rangle
\end{aligned}$$

using the substitution lemma for cartesian closed categories, properties of the internal language, the properties of M_{fo} and $N_{q,j}$ and that \equiv is a congruence. \square

12.3.4 Intensional Full Abstraction

We may now prove our final full abstraction result:

Theorem 12.3.11. *The model $\rightarrow\text{-Glob}$ is intensionally fully abstract for PCF_{//}.*

Proof. Consider $\vdash M, N : A$ such that $M \simeq N$, and assume that $\llbracket M \rrbracket \neq \llbracket N \rrbracket$, i.e. there is a test $\alpha \in \rightarrow\text{-Glob}(\llbracket A \rrbracket, \llbracket \mathbb{U} \rrbracket)$ such that, w.l.o.g., $\alpha \odot_! \llbracket M \rrbracket \Downarrow$ and $\alpha \odot_! \llbracket N \rrbracket \Uparrow$.

By Corollary 12.1.23, we assume α is additionally finite. By Theorem 12.3.9, there is a term $x : A \vdash T : \mathbb{U}$ such that $\llbracket T \rrbracket \equiv \alpha$. Defining the context $C[-] = (\lambda x^A. T)[-]$, it follows from the laws of cartesian closed categories that

$$\llbracket C[M] \rrbracket = \llbracket (\lambda x^A. T) M \rrbracket \approx \llbracket T[M/x] \rrbracket \approx \llbracket T \rrbracket \odot_! \llbracket M \rrbracket \equiv \alpha \odot_! \llbracket M \rrbracket \Downarrow,$$

and $\llbracket C[N] \rrbracket \equiv \text{comp}(\alpha) \odot_! \llbracket N \rrbracket$. By Theorem 9.1.5, $C[M] \Downarrow$. By hypothesis $M \simeq N$, so $C[N] \Downarrow$. By Theorem 9.1.5 again, $\llbracket C[N] \rrbracket \Downarrow$, hence $\alpha \odot_! \llbracket N \rrbracket \Downarrow$, contradiction. \square

Summing up, in Part III we have constructed a cartesian closed \sim -category $\rightarrow\text{-Strat}$. We have defined two conditions on morphisms of $\rightarrow\text{-Strat}$, *sequentiality* and *globularity*, forming three sub-cartesian closed \sim -categories of $\rightarrow\text{-Strat}$: $\rightarrow\text{-Seq}$, $\rightarrow\text{-Glob}$ and $\rightarrow\text{-SeqGlob}$. We established four intensional full abstraction results

$\rightarrow\text{-Strat}$	is fully abstract for	$\text{IA}_{//}$	(Th. 9.3.10)
$\rightarrow\text{-Strat} + \text{sequentiality}$	is fully abstract for	IA	(Th. 11.2.21)
$\rightarrow\text{-Strat} + \text{globularity} + \text{sequentiality}$	is fully abstract for	PCF	(Th. 11.3.12)
$\rightarrow\text{-Strat} + \text{globularity}$	is fully abstract for	$\text{PCF}_{//}$	(Th. 12.3.11)

therefore achieving our objective to *disentangle parallelism and state*. A significant part of this consisted in establishing interpretation-preserving cartesian closed \sim -functors

$\mathcal{O}\text{-Unf}$:	$\rightarrow\text{-Strat}$	\rightarrow	$\mathcal{O}\text{-Strat}$	(Theorem 9.3.8)
$\Downarrow\text{-Unf}$:	$\rightarrow\text{-Seq}$	\rightarrow	$\Downarrow\text{-Strat}$	(Theorem 11.2.18)
$\Downarrow\text{-Unf}$:	$\rightarrow\text{-SeqGlob}$	\rightarrow	$\Downarrow\text{-Inn}$	(Theorem 11.3.7)
$(-)$:	$\rightarrow\text{-Glob}$	\rightarrow	$\text{Rel}_!$	(Corollary 10.4.15)

to previously established models of different nature, which altogether, or so we hope, substantiates our claim that *thin concurrent games* can play a valuable role at the interface of many denotational semantics of programming languages.

Part IV

Other Developments and Openings

Introduction to Part IV

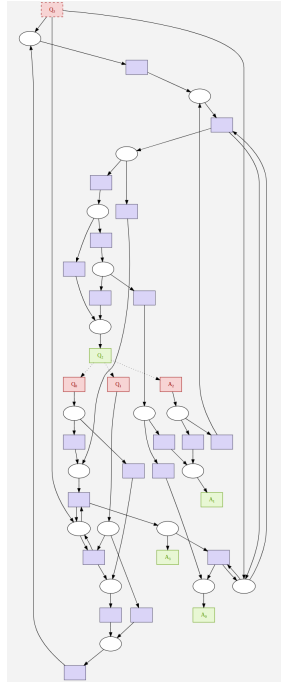


Figure 12.12: Concurrent strategies want cycles too!

In this final part, we conclude the monograph by providing some larger perspective on the landscape of concurrent games. First, in Chapter 13, we shall give a synthesis of a lot of the work done in concurrent games that could not be presented in this document. Then, in Chapter 14 we give some conclusions: we ask to what extent the work done so far substantiates the view evoked in the introduction, of concurrent games as a bridge between interactive semantics. We also then present a few open problems raised in the present development, and conclude with some perspectives.

Chapter 13

Further Work

In this monograph so far, we have given a detailed, technical account of what we consider to be the core results and developments in the theory of (thin) concurrent games. But this is not the full story; and though this monograph is already intimidatingly lengthy, we feel like a survey of the further work accomplished is called for. We show in Figure 13.1 all the works accounted for in this monograph, either in the main technical content or in this survey, grouped by themes and along with the (technical, rather than in terms of mere inspiration) dependencies between them.

This is not intended to be an exhaustive survey of concurrent games, but only of the lines of work that I have been leading. There are other interesting developments in concurrent games in which I did not participate; or for which my participation was minor; or that feel too loosely connected from the main thrust of the present document.

13.1 Effects and Concurrency

This section summarises work that was mostly developed in close collaboration with Simon Castellan, during his PhD thesis (2014-2017).

In 2015, our paper introducing **TCG**, our concept of parallel innocence, and our full abstraction result for parallel evaluation of PCF, was published [Castellan et al., 2015]. At that time, it felt ironic and a little embarrassing that we had something called *concurrent games* and yet had not developed semantics for actual concurrent effects.

It *was* very clear for us at the time how to interpret \mathbf{IA}_{\parallel} in **TCG**, but for a while, what was not clear was which theorem to prove! In the denotational semantics literature, the gold standard is full abstraction. We were confident that we could inherit intensional full abstraction by linearizing our model to [Ghica and Murawski, 2008] (though the details of this were not actually written up before this monograph – and they did require some care), but it felt unsatisfactory to only reprove an old result, and at that, one that

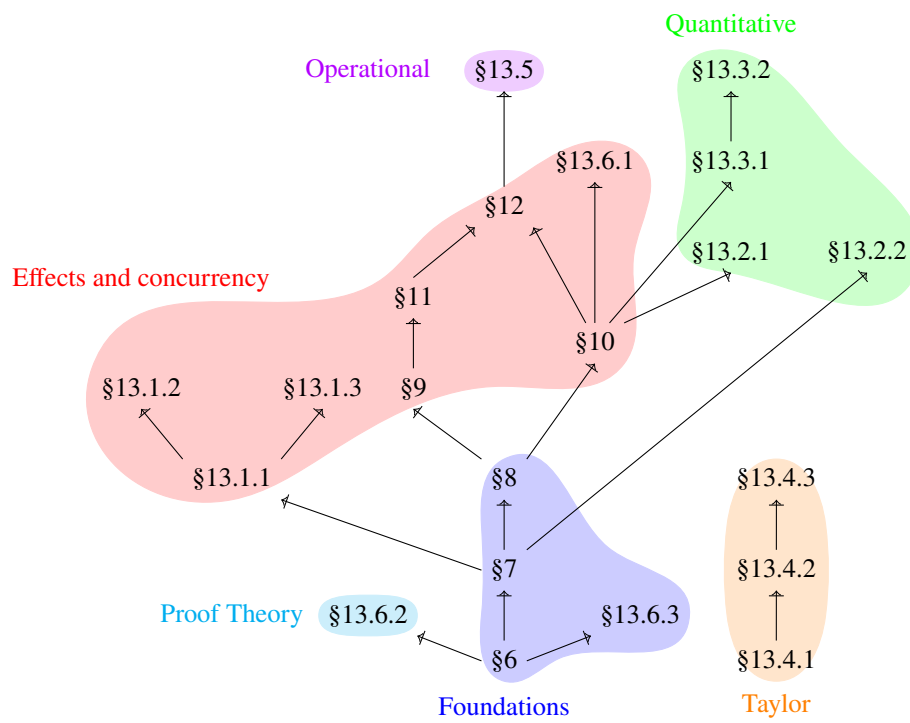


Figure 13.1: The event structure of contributions in concurrent games

does not actually need the causal information. So what is the right nail, for a hammer whose specialty is to uncover structure that by nature is *not* observable?

13.1.1 Non-canonical causal presentation

The mainstream approach, in formalizing the execution of concurrent systems, is by *interleaving* the executions of its sequential components – this is the approach followed by Ghica and Murawski’s fully abstract model for \mathbf{IA}_{\parallel} [Ghica and Murawski, 2008].

But considering a concurrent system formed of n sequential processes

$$P_1 \parallel \dots \parallel P_n,$$

even if each P_i performs only *one* atomic action, this induces $n!$ different executions: this is the *state explosion problem*, one of the main challenges to the algorithmic verification of concurrent systems. In contrast, the *truly concurrent* approach followed by concurrent games avoids interleavings and proposes a more compact representation of executions, that may in principle be helpful for algorithmic verification.

This makes concurrent games seem promising as a basis for an observational equivalence checker for \mathbf{IA}_{\parallel} , especially in the light of Ghica and Murawski’s full abstraction result [Ghica and Murawski, 2008] and the topic of *algorithmic game semantics* [Ghica and McCusker, 2003]. More precisely, the hope was to formalize the intuition that non-alternating strategies in the sense of \mathcal{O} -**Strat** already had an implicit causal structure, *i.e.* that a non-alternating strategy interpreting a term M was simply

$$\llbracket M \rrbracket = \mathcal{O}\text{-Unf}(\sigma_M)$$

for σ_M a **canonical causal explanation** for $\llbracket M \rrbracket$, a compact representation of the interleavings that may be used instead of the interleavings for algorithmic purposes.

In [Castellan and Clairambault, 2016] we showed that this was *not* the case, already in the affine case. We first built a concurrent games model of an affine version of \mathbf{IA}_{\parallel} , a corresponding non-alternating model, and an interpretation-preserving functor from the causal to the interleaving model. Then we showed that a non-alternating strategy could have *distinct* and *incomparable minimal* causal explanations – additionally, we proved this already happens within the interpretation of affine \mathbf{IA}_{\parallel} : there are two affine programs with incomparable *causal* behaviour that have the same non-alternating unfolding. This phenomenon is illustrated in Figure 13.2.

This does not show that concurrent games *cannot* help for checking the equivalence of higher-order concurrent programs, it only points out a fundamental obstacle to the route we had in mind. It is also a hint that the value of concurrent games is in its presentation of the *intensional behaviour*, not as an efficient way to present the interleavings. Thus, we then refocused our research on the intensional behaviour.

13.1.2 Non-angelic concurrent games

Another aspect of concurrent games that was clear from the start is that they record explicitly the point of non-deterministic branching. For instance, the program $x : \mathcal{U} \vdash$




Figure 13.2: As the goal is only to reconstruct interleavings, the branching structure does not matter, so rather than an event structure we can work only with a set of *augmentations*, *i.e.* the partial orders induced by configurations. In this example, both the second and the third augmentation can be removed without changing the interleavings – but no other; yielding two incompatible minimal causal explanations.




Figure 13.3: Recording non-angelic behaviour




Figure 13.4: After hiding




Figure 13.5: Composition with essential events

if choice then x ; \perp **else** x ; $\checkmark : \mathbb{U}$ yields the strategy of Figure 13.3, tracking explicitly the two incompatible calls to x and the fact that one yields to a divergence, while the other yields termination. In most earlier games models (with the notable exception of [Harmer and McCusker, 1999] for Erratic IA), the diverging branch would be ignored: this is a direct consequence of the traditional methodology of representing a program as a *set of plays*, *i.e.* a set of converging behaviour. We only record the converging behaviour, forgetting the non-deterministic branching information, leading to an *angelic* view of termination, *i.e.* may-convergence.

This discussion might lead the reader to expect that concurrent games are sound with respect to must-convergence, but that is not the case: *hiding* removes branches that do not prompt any visible behaviour, even if they witness possible non-deterministic branches. For instance, composing the strategy of Figure 13.3 with **skip** yields the interaction of Figure 13.4 which, after hiding, yields the constant strategy $\llbracket \text{skip} \rrbracket$ – because of hiding, we forget that something may have gone wrong in the computation.

Nevertheless, in [Castellan et al., 2018a], we proved that concurrent games are easily amended so as to recall more than angelic information on program behaviour. It is simply a matter of changing *hiding*, so that we retain events that are part of a minimal conflict: those are kept through hiding as events of neutral polarity, dubbed *essential*

events. With this modification, the composition above yields the strategy of Figure 13.5 where the events labelled \star are those essential events. With this interpretation, we keep more than angelic information: in fact, we show that (the transition system generated by) the strategy is *weakly bisimilar* to (the transition system generated by) the operational semantics, for an affine version of $\mathbf{IA}_{//}$. We are not aware of any other work capturing by denotational means such a precise information on program behaviour.

This felt a rather good illustration of the expressivity of concurrent games, and in particular of its feature that it retains the non-deterministic branching information. But, weak bisimulation does not really take causality into account.

13.1.3 Resource-tracking concurrent games

This is a collaboration with Aurore Alcolei during her PhD thesis, and Olivier Laurent.

To explore the expressiveness of causality in concurrent games, we needed to record information about programs that really depends on the causal independence between events, and is incompatible with an interleaving interpretation. As a natural objective, we thought of the following problem. Consider $\mathbf{IA}_{//}$ extended with a primitive **wait**(n) that takes a natural number n , and waits for n seconds. The program

$$\vdash \text{newref } r \text{ in let } \begin{pmatrix} x = \text{wait}(1); r := 1; \text{wait}(2) \\ y = \text{wait}(2); r := 2; \text{wait}(1) \end{pmatrix} \text{ in } !r : \mathbb{N}$$

always evaluates in 6 seconds according to an interleaving operational semantics, as all the **wait** instructions must be handled before termination. But running this program on a multicore architecture, one will observe faster termination times: the program should evaluate to 2 in about 3 seconds, but – if the thread for x gets stuck – we can in principle observe a termination to 1 in at least 4 seconds. This can be captured operationally via a semantics where threads can wait *in parallel*, but it seems hard to capture denotationally as one must be able to *see* when parts of the computation are causally independent.

Investigating this question, we had two main inspirations: firstly [Ghica, 2005], where Ghica builds a variant of $\mathcal{G}\text{-Strat}$ to capture semantically the (interleaving) execution time for an $\mathbf{IA}_{//}$ -like language; and secondly [Laird et al., 2013], where the authors use the relational model weighted over the tropical or arctic semirings to capture resource usage for a PCF-like language extended with non-deterministic choice.

In [Alcolei et al., 2019] we introduce a version of concurrent games parametrized by an algebraic structure we call a *resource bimonoid*, that tracks resource consumption *sequentially* and *in parallel* – our main example is \mathbb{R}_+ for time, equipped with $+$ (for sequential resource consumption) and \max (for parallel resource consumption). In this model, strategies track resource consumption: each positive move comes equipped with an expression describing the resource usage necessary to play this move, as a function of the resources used to reach its causal dependencies. Building on this, we give (among other results) an interpretation of an affine $\mathbf{IA}_{//}$ -like language extended with **wait**, correctly predicting execution time in a maximally parallel evaluation model. This is illustrated in Figure 13.6, that shows the interaction yielding the interpretation of the program above, where each event is annotated by its minimal time of appearance.




Figure 13.6: Predicting minimal parallel evaluation time

This is neat as it is an application that really exploits the causal information available in concurrent games. It is also the first occurrence of a trend that we followed next: the enriching of concurrent games with quantitative information.

13.2 Quantitative Concurrent Games

This section summarizes work developed in collaboration with Hugo Paquet and Glynn Winskel, as part of Hugo Paquet's PhD thesis¹; and – regarding the work on semantics of quantum programming – as part of Marc de Visme's PhD thesis (2017-2020).

The original motivation for this line of work was two-fold. The first was an intriguing gap in the literature, originally investigated by Hugo during a research internship at IRIF with Christine Tasson. There were two fully abstract models for fairly close probabilistic languages: probabilistic game semantics [Danos and Harmer, 2000] for probabilistic IA, and the weighted relational model [Laird et al., 2013] for probabilistic PCF. While these models were fairly close in spirit, it was very unclear how to actually relate them formally. Could we have a time-forgetting map from games to weighted relations? From reading Melliès' papers on asynchronous games [Melliès, 2005], it was clear to me that the source of such a time-forgetting map should not be Danos and

¹Hugo did his PhD in Cambridge with Glynn Winskel, defended in 2019, but I spent most of the year 2016 in Cambridge, and a lot of that work was carried out during that time.

Harmer’s model, but rather a concurrent games version, for reasons similar to those that make the relational collapse of Section 10.4 tick.

The second motivation was to give a notion of *probabilistic innocence*, i.e. to capture the interactive behaviour of pure parallel programs. It had been known for long that the naive theory of non-deterministic innocence does not work in sequential game semantics [Harmer, 1999]; and it seemed that such a theory could be obtained in probabilistic concurrent games thanks to their explicit handling of the branching information².

In both cases, this seemed to be problems where concurrent games were naturally helpful, so we embarked in our journey in probabilistic concurrent games.


13.2.1 Probabilistic PCF

When we started working on this, there was already a notion of probabilistic concurrent strategy, due to Glynn [Winskel, 2013a], which we first review.

The initial idea is simple: a probabilistic concurrent strategy is a strategy σ with

$$v_\sigma : \mathcal{C}(\sigma) \rightarrow [0, 1]$$

a *probability valuation* that expresses *the probabilistic weight of the Player choices made so far*. So for instance, a strategy for a fair random boolean would have



with $v(\emptyset) = 1$, $v(\{q^-\}) = 1$, $v(\{q^-, tt^+\}) = v(\{q^-, ff^+\}) = \frac{1}{2}$ – playing Opponent moves has no impact on the probability valuation as only the weight of the Player choices is tracked (this is similar to the approach in [Danos and Harmer, 2000]).

Glynn’s probabilistic strategies must satisfy additional conditions. The easy ones are

$$\begin{aligned} \text{normalized: } & v(\emptyset) = 1, \\ \text{receptive: } & \text{if } x^\sigma \subseteq^- y^\sigma, \text{ then } v(x^\sigma) = v(y^\sigma), \end{aligned}$$

but there is more. If from a configuration $x^\sigma \in \mathcal{C}(\sigma)$, σ can do two events s_1^+ and s_2^+ which are additionally conflicting, then we expect to have the inequality

$$\frac{v(x^\sigma \cup \{s_1^+\})}{v(x^\sigma)} + \frac{v(x^\sigma \cup \{s_2^+\})}{v(x^\sigma)} \leq 1 :$$

the first term of the sum is the probability of playing s_1^+ in configuration x^σ , the second is the probability of playing s_2^+ in configuration x^σ ; the condition amounts to asking that those form a sub-probability distribution – note the same condition can be written

$$v(x^\sigma) \geq v(x^\sigma \cup \{s_1^+\}) + v(x^\sigma \cup \{s_2^+\}).$$

²Other works had already arrived at notions of non-deterministic innocence via adopting strategies with explicit branching structure [Castellan et al., 2014, Tsukada and Ong, 2015].

In sequential probabilistic game semantics [Danos and Harmer, 2000], a similar condition suffices. But the soundness of this relies on the conflict between s_1^+ and s_2^+ , so that $x^\sigma \cup \{s_1^+\}$ and $x^\sigma \cup \{s_2^+\}$ delineate two disjoint parts of the probability space. In concurrent games it is possible that $x^\sigma \cup \{s_1^+, s_2^+\}$ is also a configuration, which is accounted for twice in the sum above. Correcting for this yields the inequality

$$v(x^\sigma) \geq v(x^\sigma \cup \{s_1^+\}) + v(x^\sigma \cup \{s_2^+\}) - v(x^\sigma \cup \{s_1^+, s_2^+\}),$$

but a similar phenomenon occurs for three concurrent events, in which case

$$\begin{aligned} v(x^\sigma) \geq & v(x^\sigma \cup \{s_1^+\}) + v(x^\sigma \cup \{s_2^+\}) + v(x^\sigma \cup \{s_3^+\}) \\ & - v(x^\sigma \cup \{s_1^+, s_2^+\}) - v(x^\sigma \cup \{s_1^+, s_3^+\}) - v(x^\sigma \cup \{s_2^+, s_3^+\}) \\ & + v(x^\sigma \cup \{s_1^+, s_2^+, s_3^+\}), \end{aligned}$$

and the general formula can be written following the *inclusion-exclusion principle* – Glynn called this the *drop condition* as it requires that a quantity, the “drop” obtained as the left hand side term minus the right hand side term, should be positive.

Note that by *receptive*, it suffices to assign probability valuations to $+$ -covered configurations. This makes the compositional structure of probabilistic strategies very easy:

$$v_{\tau \odot \sigma}(x^\tau \odot x^\sigma) = v_\sigma(x^\sigma) \cdot v_\tau(x^\tau) \quad (13.1)$$

for all $x^\tau \odot x^\sigma \in \mathcal{C}^+(\tau \odot \sigma)$; and $v_{\mathfrak{A}}(\mathfrak{A}_x) = 1$ for all $x \in \mathcal{C}(A)$. Probabilistic strategies directly inherit from **CG** the structure of a bicategory (the only difficulty being composition of the drop condition). But it is “only” a bicategory, and we needed much more in order to interpret PCF, so our first task was to merge this with **TCG**.

In doing that, we had the good surprise to find that this needed almost no work at all! In dealing with symmetry, the only condition that must be added to strategies is

$$\text{invariance: for all } x^\sigma \cong_\sigma y^\sigma, \text{ we have } v_\sigma(x^\sigma) = v_\sigma(y^\sigma)$$

ensuring that the probability valuation is invariant under symmetry – the drop condition, for instance, remained unchanged³, forming a relative Seely category that we shall refer to as **ProbNTCG** in the context of this monograph. Then, the same constructions used to model PCF and reported in Chapter 8 work just as well to give a model of probabilistic PCF – this was the starting point of the paper [Castellan et al., 2018b].

Probabilistic innocence. Then in that situation, we looked for a notion of *probabilistic innocence* – note that we are talking about *sequential* probabilistic innocence, as it is already unclear what is parallel *non-deterministic* innocence, see Section 14.2.1. Hence we could build on the sequential non-deterministic innocence of [Castellan et al., 2014]:

Definition 13.2.1. A strategy σ on board A is non-deterministic sequential innocent if:

$$\begin{aligned} \text{negative:} & \text{ for all } s \in \min(\sigma), \text{pol}_\sigma(s) = -, \\ \text{forestal:} & \text{ if } s_1, s_2 \leq_\sigma s, \text{ then } s_1 \leq_\sigma s_2 \text{ or } s_2 \leq_\sigma s_1, \\ \text{locally conflicting:} & \text{ if } s_1 \sim_\sigma s_2, \text{ then } [s_1]_\sigma = [s_2]_\sigma, \\ \text{locally sequential:} & \text{ if } s^- \rightarrow_\sigma s_1^+, s_2^+, \text{ then } s_1^+ \sim_\sigma s_2^+. \end{aligned}$$

³Note in passing the nightmare avoided by using *thin* concurrent games instead of the *saturated* concurrent games of [Castellan et al., 2014]! See Section 7.1.2 for a reminder on the thin / saturated distinction.

This takes as a definition the forestial shape that was derived in this monograph from combined sequentiality and parallel innocence (see Lemma 11.3.2); with the difference that the resulting forest might not be O-branching: Player is now also able to branch. But that Player branching cannot be parallel branching, it must come with a conflict. Thus any configuration is an O-branching forest, where each branch is an independent exploration of a branch of the term, following the intuitions of Section 11.3.1.

Now for a *probabilistic strategy*, how should the probabilistic weight behave with respect to this forestial structure? What we proposed in [Castellan et al., 2018b] is that on top of non-deterministic sequential innocence, we should have

Definition 13.2.2. Consider $\sigma : A$ probabilistic, non-deterministic seq. innocent.

It is **probabilistic innocent** iff for all $x, y \in \mathcal{C}(\sigma)$ such that $x \cup y \in \mathcal{C}(\sigma)$,

$$v_{\sigma}(x \cup y) = \frac{v_{\sigma}(x) \cdot v_{\sigma}(y)}{v_{\sigma}(x \cap y)}.$$

This asserts that two *causally independent* branches of a configuration must also be *probabilistically independent*: the extensions $x \cap y \subseteq x$ and $x \cap y \subseteq y$ must explore distinct branches of the syntax tree, and hence trigger independent probabilistic choices.

Probabilistic innocent strategies form a relative Seely category that supports the interpretation of probabilistic PCF, and enjoy a finite definability result yielding intensional full abstraction for probabilistic PCF. But probabilistic PCF has a simpler, actually fully abstract model: probabilistic coherence spaces [Ehrhard et al., 2018] and hence their backbone, the weighted relational model [Laird et al., 2013].

Our paper had a second result, an interpretation-preserving collapse to the weighted relational model, but we shall come back to that in Section 13.3. In any case, there was clearly a parallel between concurrent games and relational models: just as relational models are made quantitative by attaching to each point of the “web” a coefficient, and concurrent games are made quantitative by attaching to each configuration a coefficient. But the relational supports more elaborate weights than just scalars: in particular, points of the web can be decorated by finite dimensional Hilbert spaces and completely positive maps, to obtain a computationally adequate model of the quantum λ -calculus [Pagani et al., 2014]. This suggested to do the same in concurrent games.

13.2.2 Quantum strategies

This is joint work with Marc de Visme for his PhD (2017-2020), and Glynn Winskel.

In Marc’s PhD thesis, we set as main objective to build a computationally adequate concurrent games model for the quantum λ -calculus; this is a call-by-value language, linear but with an explicit exponential modality, structured via the types

$$A, B ::= \text{qubit} \mid \mathbf{1} \mid A \otimes B \mid A \oplus B \mid \text{list}(A) \mid A \multimap B \mid !(A \multimap B),$$


with e.g. $\mathbb{B} = \mathbf{1} \oplus \mathbf{1}$ defined as syntactic sugar; linearity is of course crucial because qubits *cannot* be duplicated. Terms include the expected term formers for these types,

along with new primitives, used to manipulate quantum data:

$$\mathbf{new} : \mathbb{B} \multimap \text{qubit}, \quad \mathbf{meas} : \text{qubit} \multimap \mathbb{B}, \quad U : \text{qubit}^{\otimes n} \multimap \text{qubit}^{\otimes n},$$


where **new** initializes a new qubit based on a boolean, **meas** measures a qubit (with a probabilistic result), and U imports any n -ary unitary inside the syntax.

Building a concurrent games model for this language posed several challenges. First of course, the language is call-by-value – while there were no earlier concurrent games models of call-by-value languages, it was not too difficult to infer what the corresponding structure should be, looking at traditional games models for call-by-value languages [Abramsky and McCusker, 1997, Honda and Yoshida, 1999] and from the game semantics of tensorial logic in *e.g.* [Melliès and Tabareau, 2007]. In particular, in contrast with **NTCG**, types (ignoring the quantum aspect) should be interpreted as *positive* games, where Player always starts. For instance, the type $\mathbf{1} \multimap \mathbf{1}$ yields the board



where computation starts by Player providing a value – in that written λ^+ , as it carries the information that the program yields an abstraction. This enables Opponent to call the function by providing an argument which can only be the unique value of unit type, which in turn enables Player to terminate with again, the unique value of unit type.

The main challenge was dealing with quantum information. For this, our idea follows the paradigm to separate *quantum data* and *classical control*, putting together ideas from probabilistic concurrent games above, and quantum relational semantics [Pagani et al., 2014]. For *classical control*, qubits are interpreted just like the unit type, so that the plain event structure for $\text{qubit} \multimap \text{qubit}$ is exactly like $\mathbf{1} \multimap \mathbf{1}$ above. The difference comes with *quantum data*, for which each event $a \in A$ of the game comes labelled with a (finite dimensional) Hilbert space $\mathcal{Q}_A(a)$. Events that do not carry any quantum state are annotated with the 1-dimensional Hilbert space, the unit for the tensor of Hilbert spaces, *i.e.* the complex numbers \mathbb{C} . Events that do have a quantum content carry a non-trivial Hilbert space; for instance $\text{qubit} \multimap \text{qubit}$ becomes the game



with Hilbert space annotations in green. Intuitively, playing a move annotated with a Hilbert space H , has a side effect of opening H in the quantum register. Accordingly,




Figure 13.7: Quantum strategy for measurement

any configuration $x \in \mathcal{C}(A)$ yields a finite dimensional Hilbert space

$$\mathcal{Q}_A(x) = \bigotimes_{a \in x} \mathcal{Q}_A(a),$$

the tensor of the Hilbert spaces opened in x . Now the idea is that a strategy $\sigma : A$ is accompanied by a valuation as for probabilistic strategies, except it is *typed*: to each $x^\sigma \in \mathcal{C}(\sigma)$ we associate a *density operator* on the Hilbert space of moves played so far:

$$v_\sigma(x^\sigma) \in \mathbf{DOp}(\mathcal{Q}_A(\partial_\sigma x^\sigma));$$

recall that density operators are the standard representation of *mixed quantum states*, *i.e.* supporting pure states but also convex (probabilistic) sums and measurements. Now, the definition of quantum strategies is analogous to that of probabilistic strategies, except that the valuations are *typed*: the *drop condition* can be defined, though it is with respect to the so-called *Löwner order* between density operators and involves some tracing out to bring all density operators to the same Hilbert space.

At the core of our approach, we use that for a strategy between games $\sigma : A \vdash B$, a quantum valuation on $A \vdash B$, *i.e.* a density operator on a tensor Hilbert space

$$v_\sigma(x^\sigma) \in \mathbf{DOp}(\mathcal{Q}_A(x_A^\sigma) \otimes \mathcal{Q}_B(x_B^\sigma))$$

may be equivalently presented as a *completely positive map*, *i.e.* a morphism

$$v_\sigma^\rightarrow(x^\sigma) \in \mathbf{CPM}(\mathcal{Q}_A(x_A^\sigma), \mathcal{Q}_B(x_B^\sigma)),$$

in the compact closed category **CPM** that we shall not define here, but which should be thought of as taking mixed states on $\mathcal{Q}_A(x_A^\sigma)$ (so, certain operators on $\mathcal{Q}_A(x_A^\sigma)$ to mixed states on $\mathcal{Q}_B(x_B^\sigma)$). This lets us define the quantum valuation of the composition $\tau \odot \sigma$ simply by composition of the quantum valuations of σ and τ respectively, via

$$v_{\tau \odot \sigma}^\rightarrow(x^\tau \odot x^\sigma) = v_\tau^\rightarrow(x^\tau) \circ v_\sigma^\rightarrow(x^\sigma)$$

reminiscent of (13.1) for probabilistic strategies.

As an example, we show in Figure 13.7 the quantum strategies for measurement. Its control flow seems to select a boolean non-deterministically, but each branch is

weighted by the **CPM** map taking a density matrix on \mathbb{C}^2 and maps it to (a 1 by 1 matrix with) the probability of obtaining the corresponding value via a measurement.

With Marc and Glynn, we published a first paper [Clairambault et al., 2019] presenting the model; its main results are computational adequacy, but also an important consequence of the quantum *drop condition*: our quantum valuations can be presented as *superoperators*, those **CPM** maps that are trace non-increasing, they correspond to the physically realizable transformations on mixed quantum states. More precisely, consider $\sigma : A$, $x^\sigma \in \mathcal{C}(\sigma)$ and $\partial_\sigma x^\sigma = x_A$, and split x_A into its positive and negative events $x_A = x_A^- \uplus x_A^+$. Then we may present equivalently $v_\sigma(x^\sigma)$ as

$$v_\sigma^-(x^\sigma) \in \mathbf{CPM}(\mathcal{Q}_A(x_A^-), \mathcal{Q}_A(x_A^+)),$$

which we show is always a superoperator. This was an improvement with respect to the earlier denotational model of the quantum λ -calculus [Pagani et al., 2014], for which there was no such boundedness property – their model construction required valuations to include an infinitary completion of **CPM**, just as the weighted relational model requires an infinitary completion of scalars to work around convergence issues.

In a second paper with Marc [Clairambault and de Visme, 2020], we were able to exploit this boundedness property to adapt in our model the proof method of Ehrhard, Pagani and Tasson for their full abstraction result of probabilistic coherence spaces for Probabilistic PCF [Ehrhard et al., 2018]. This shows that our model is fully abstract for the quantum λ -calculus, modulo a quotient of strategies that amounts to a quantitative relational collapse to the model of [Pagani et al., 2014], which entails that their model was *already* fully abstract. We cover the ideas around this quantitative collapse next.

13.3 Quantitative Relational Collapses

13.3.1 Counting witnesses and quantitative collapse

When working on [Castellan et al., 2018b], one of our goals was to understand the link between probabilistic alternating game semantics [Danos and Harmer, 2000] and probabilistic coherence spaces [Danos and Ehrhard, 2011] via the weighted relational model [Laird et al., 2013]; for that we meant to define an interpretation-preserving functor

$$(-) : \mathbf{ProbNTCG-Vis} \rightarrow \overline{\mathbb{R}}_+ \mathbf{-Rel}$$

refining quantitatively the one of Corollary 10.4.15 – recall that here, $\overline{\mathbb{R}}_+ \mathbf{-Rel}$ is the *weighted relational model* [Laird et al., 2013], here weighted by the semiring $\overline{\mathbb{R}}_+ = \mathbb{R}_+ \uplus \{+\infty\}$ of the completed positive reals. Its objects are sets, and morphisms from A to B are matrices $(\alpha_{a,b})_{a \in A, b \in B}$ where $\alpha_{a,b} \in \overline{\mathbb{R}}_+$ for all $a \in A$ and $b \in B$.

Composition of $\alpha \in \overline{\mathbb{R}}_+ \mathbf{-Rel}(A, B)$ and $\beta \in \overline{\mathbb{R}}_+ \mathbf{-Rel}(B, C)$ is matrix multiplication:

$$(\beta \circ \alpha)_{a,c} = \sum_{b \in B} \alpha_{a,b} \cdot \beta_{b,c}$$

which might not converge if B is infinite – hence the completion to $\overline{\mathbb{R}}_+$.

At the time we were investigating this, it felt like it should be a minor variation of the usual relational collapse of Section 10.4. Recall that for $\sigma : A$, Definition 10.4.3 has

$$(\sigma) = \{x_A \in (A) \mid \exists x^\sigma \in \mathcal{C}^+(\sigma), \partial_\sigma x^\sigma \in x_A\},$$

recording all $x_A \in (A)$ which is reachable in σ , which is witnessed by $x^\sigma \in \mathcal{C}^+(\sigma)$ such that $\partial_\sigma x^\sigma \in x_A$. Intuitively, instead of only recording the *mere existence* of such a witness, we must *sum the weight* of all possible witnesses. Of course that is too naive: as a strategy needs to be receptive to all Opponent reindexings, and as there are usually countably many of those, each “witness” in the sense above usually has countably many distinct but symmetric peers, yielding almost always infinitely many witnesses.

In [Castellan et al., 2018b], we proposed to solve this issue by summing not over all concrete witnesses as above, but over all *symmetry classes* of witnesses:

$$(\sigma)_{x_A} = \sum_{x^\sigma \in \text{wit}_\sigma^\cong(x_A)} v_\sigma(x^\sigma)$$

for $\text{wit}_\sigma^\cong(x_A) = \{x^\sigma \in \mathcal{C}^+(\sigma) \mid \partial_\sigma x^\sigma \in x_A\}$ using that by *invariance*, the valuation v_σ lifts to symmetry classes. This yields a finite coefficient that operates as it should on simple examples. For composition, functoriality amount to proving the equality

$$(\tau \odot \sigma)_{x_A \vdash x_C} = \sum_{x_B \in (B)} (\sigma)_{x_A \vdash x_B} \cdot (\tau)_{x_B \vdash x_C}$$

which seems like it should be a natural consequence of a bijection

$$\text{wit}_{\tau \odot \sigma}^\cong(x_A \vdash x_C) \simeq \sum_{x_B \in (B)} \text{wit}_\sigma^\cong(x_A \vdash x_B) \times \text{wit}_\tau^\cong(x_B \vdash x_C). \quad (13.2)$$


This seems simple enough: from $x^\sigma \in \text{wit}_\sigma^\cong(x_A \vdash x_B)$ and $x^\tau \in \text{wit}_\tau^\cong(x_B \vdash x_C)$, we can pick representatives $x^\sigma \in x^\sigma$, $x^\tau \in x^\tau$ and a symmetry $\theta_B : x_B^\sigma \cong_B x_B^\tau$, and obtain an element of $\text{wit}_{\tau \odot \sigma}^\cong(x_A \vdash x_C)$ as (the symmetry class of) $x^\tau \odot_\theta x^\sigma$ defined following Definition 7.4.6. It can be checked that any $x \in \text{wit}_{\tau \odot \sigma}^\cong(x_A \vdash x_C)$ can be obtained in this ways, and that $x^\sigma \in \text{wit}_\sigma^\cong(x_A \vdash x_B)$ and $x^\tau \in \text{wit}_\tau^\cong(x_B \vdash x_C)$ can be *uniquely* recovered from x . So this bijection seemed to work just fine, as expected!

Counting witnesses. This is how the collapse was written in [Castellan et al., 2018b], and it was the state of my understanding for about two years. But later on, when writing the details of the quantum version of this collapse with Marc, we realized that something was off – the more typed version of the collapse made it more apparent.

The issue is that in the reasoning above, we have defined

$$\text{wit}_\sigma^\cong(x_A \vdash x_B) \times \text{wit}_\tau^\cong(x_B \vdash x_C) \rightarrow \text{wit}_{\tau \odot \sigma}^\cong(x_A \vdash x_C)$$

a function by working on concrete representatives, and we had not checked that the result does not depend on the representative. And of course it does, and it also depends

Figure 13.8: Composition of σ and τ

on the choice of the mediating symmetry θ_B – which in hindsight should have been obvious with some familiarity with the resource calculus, because θ_B somehow corresponds to the non-determinism of reduction in the resource calculus (see Section 13.4). As a matter of fact, (13.2) fails; let us see a counter-example:

Example 13.3.1. Take the board α with one event \mathbf{q}^- , $\kappa_\alpha(\emptyset) = 1$, $\kappa_\alpha(\{\mathbf{q}^-\}) = 0$, and

$$\sigma : !(\alpha \multimap \alpha) \vdash ! \alpha \multimap \alpha, \quad \tau : !(\alpha \multimap \alpha)$$

two strategies as in Figure 13.8. Their assignment of copy indices uses functions

$$f : \mathbb{N} \rightarrow \mathbb{N}, \quad g : \mathbb{N} \rightarrow \mathbb{N}, \quad h : \mathbb{N} \rightarrow \mathbb{N}, \quad k : \mathbb{N}^2 \rightarrow \mathbb{N},$$

whose precise identity is irrelevant. Their composition unfolds as in Figure 13.8, yielding four pairwise conflicting, non-symmetric positive moves.

Enriching non-deterministic PCF with a type α with no constant, the substitution

$$(\lambda x^\alpha. f(f x)) [(\lambda y^\alpha. y \otimes y)/f]$$

gives a perfect syntactic counterpart to the composition in Figure 13.8. As there are two calls to the non-deterministic choice, this reduces to $\lambda x^\alpha. x$, but in four different ways. In the copy indices of each positive move of $\tau \odot \sigma$, one can read back how the two non-deterministic choices were resolved: the upper row corresponds to the first call yielding $\mathbf{q}_{i,f(i)}^+$, the leftmost column to the second call yielding $\mathbf{q}_{i,f(i)}^+$, and so on.

Reading back witnesses from the components and their composition, we have

$$\# \text{wit}_{\tau \odot \sigma}^{\cong} \left(\begin{array}{c} \mathbf{q}^- \\ \vdots \\ \mathbf{q}^+ \end{array} \right) = 4, \quad \# \text{wit}_{\sigma}^{\cong} \left(\begin{array}{cc} \mathbf{q}^- & \mathbf{q}^- \\ \vdots & \vdots \\ \mathbf{q}^+ & \mathbf{q}^+ \end{array} \right) = 1, \quad \# \text{wit}_{\tau}^{\cong} \left(\begin{array}{cc} \mathbf{q}^- & \mathbf{q}^- \\ \vdots & \vdots \\ \mathbf{q}^+ & \mathbf{q}^+ \end{array} \right) = 3;$$

as σ and τ cannot synchronize on any other symmetry class, this contradicts (13.2).

After staring at this example for a while, it seemed that the mistake was in the enumeration of the witnesses for τ . Indeed, up to symmetry, there are exactly *three* configurations of τ corresponding to two calls: (1) both choices may be resolved with $\mathbf{q}_{i,f(i)}^+$ (call this “left-left”); (2) both choices may be resolved with $\mathbf{q}_{i,g(i)}^+$ (call this “right-right”);

and (3) we may have one of each. But arguably, (3) really identifies two cases: (3a) “left-right” and (3b) “right-left”. Those configurations are symmetric, but it turns out that **TCG** is equipped with exactly the right tool to keep them distinct. Imagine fixed, for every $x_A \in \mathbf{A}$, a concrete representative $\underline{x}_A \in x_A$. Then, for $\sigma : A$ we set

$$\text{wit}_\sigma^+(x_A) = \{x^\sigma \in \mathcal{C}^+(\sigma) \mid \partial_\sigma x^\sigma \cong_A^+ \underline{x}_A\}$$

the set of **positive witnesses**⁴ for x_A , *i.e.* those *concrete* configurations of σ for which there exists a *positive* symmetry to the chosen representative \underline{x}_A . The *positivity* requirement morally fixes the identity of each Opponent move, factoring out Opponent reindexing but still letting us see the Player behaviour. We now have

$$\overline{\text{wit}}_\tau^+ \left(\begin{array}{cc} \mathbf{q}_0^- & \mathbf{q}_1^- \\ \vdots & \vdots \\ \mathbf{q}^+ & \mathbf{q}^+ \end{array} \right) = \left\{ \left(\begin{array}{cc} \mathbf{q}_0^- & \mathbf{q}_1^- \\ \vdots & \vdots \\ \mathbf{q}_{0,f(0)}^+ & \mathbf{q}_{1,f(1)}^+ \end{array} \right), \left(\begin{array}{cc} \mathbf{q}_0^- & \mathbf{q}_1^- \\ \vdots & \vdots \\ \mathbf{q}_{0,g(0)}^+ & \mathbf{q}_{1,g(1)}^+ \end{array} \right), \left(\begin{array}{cc} \mathbf{q}_0^- & \mathbf{q}_1^- \\ \vdots & \vdots \\ \mathbf{q}_{0,f(0)}^+ & \mathbf{q}_{1,g(1)}^+ \end{array} \right), \left(\begin{array}{cc} \mathbf{q}_0^- & \mathbf{q}_1^- \\ \vdots & \vdots \\ \mathbf{q}_{0,g(0)}^+ & \mathbf{q}_{1,f(1)}^+ \end{array} \right) \right\}$$

a set of concrete configurations of cardinal 4, solving the mismatch.

In [Clairambault and Paquet, 2021] we detail this, which involves delving into an intricate analysis of symmetries in concurrent games. Thanks to this we are able to correct the interpretation-preserving functor $(-): \mathbf{ProbNTCG-Vis} \rightarrow \overline{\mathbb{R}}_+ \mathbf{-Rel}$ of [Castellan et al., 2018b]: the correct definition turns out to be given by

$$(\sigma)_{x_A} = \sum_{x^\sigma \in \text{wit}_\sigma^+(x_A)} v_\sigma(x^\sigma),$$

summing over *positive witnesses* rather than symmetry classes⁵. Beyond probabilistic weights, we prove this for games and relations weighted by arbitrary semirings of resources, modulo mild hypotheses. An interesting case is the interpretation of PCF in $\overline{\mathbb{N}}\mathbf{-Rel}$, the relational model weighted by natural numbers. If $\vdash M : A$ is a term, to each $a \in \llbracket A \rrbracket$, $\llbracket M \rrbracket_a$ is a natural number, which might be arbitrarily high. So the relational model *counts* something, but what? It turns out from our study that the weighted relational model counts *positive witnesses* in the sense of concurrent game semantics.

Finally, in [Clairambault and de Visme, 2020], we also apply these ideas for the quantitative collapse of quantum strategies onto the quantum relational model.

13.3.2 Generalized species of structure

Interpreting programs in **TCG** allows us – for each point of the web in the sense of relational semantics – to replace a *natural number*, as given by the weighted relational model, with a *set of witnesses*. This is an example of what is increasingly being called *proof-relevant denotational semantics*; which beyond collecting behaviours, tracks, for each possible behaviour, a set of *witnesses*. The main example of such a model is the cartesian closed bicategory of *generalized species of structure* [Fiore et al., 2008].

⁴The cardinal of $\text{wit}_\sigma^+(x_A)$ does not depend on the representative, provided it satisfies a condition called *canonicity*, not covered here, see [Clairambault and Paquet, 2021].

⁵One can sum over symmetry classes, with correcting coefficients: $(\sigma)_{x_A} = \sum_{x^\sigma \in \text{wit}_\sigma^+(x_A)} \frac{\#\mathcal{S}_-(x_A)}{\#\mathcal{S}_-(x^\sigma)} \cdot v_\sigma(x^\sigma)$.

At the time (2021-2022), there was intense activity, in the quantitative semantics community, around species of structure. Roughly speaking, species of structure are a categorification of the relational model. The objects of **Rel**, *i.e.* sets, are replaced with categories. The morphisms of **Rel**, *i.e.* functions $\alpha : A \times B \rightarrow \{0, 1\}$ are replaced with

$$\alpha : A^{\text{op}} \times B \rightarrow \mathbf{Set},$$

i.e. profunctors from A to B . The *bang* of **Rel**, which to any set associates its set of finite multisets, is replaced with the *free symmetric monoidal category*. Altogether, to a program $\vdash M : A$ and some $a \in A$, the interpretation of M associates a set $\llbracket M \rrbracket_a$ – notably, Federico Olimpieri proved [Olimpieri, 2021] that this set could be presented as the set of *derivations* of $M : a$ in a non-idempotent intersection type system.

As two “proof-relevant” denotational models, it was appealing to try to understand the links between concurrent games and generalized species of structure. As generalized species of structure replace the coefficient computed by the weighted relational model with a set, it is tempting to conjecture that the two agree – that the coefficient given by the weighted relational model is the cardinal of the witnesses given by generalized species. But it is easy to see that this is not true: for instance, for any set A we have

$$(\text{id}_A)_{a,a'} = \begin{cases} 1 & \text{if } a = a', \\ 0 & \text{otherwise} \end{cases}$$

for the identity in the weighted relational model, compared with $(\text{id}_A)_{a,a'} = A[a, a']$ given by the hom-functor, the identity in the bicategory of profunctors. We can plainly see that species of structures import morphisms of A as witnesses, whereas the weighted relational model does not. From the perspective of concurrent games these morphisms are nothing but the *symmetries* between configurations, so it seems that witnesses in the sense of generalized species of structure should correspond to witnesses in the sense of concurrent games *along with* explicit symmetries, *i.e.* for a strategy $\sigma : A \vdash B$,

$$(\sigma)(x_A, x_B) = \{(\theta_A^-, x^\sigma, \theta_B^+) \mid \theta_A^- : x_A \cong_A^- x_A^\sigma, x^\sigma \in \mathcal{C}^+(\sigma), \theta_B^+ : x_B^\sigma \cong_B^+ x_B\}.$$

Though defined on configurations, this can be turned into a profunctor

$$(\sigma) : \mathcal{S}(A)^{\text{op}} \times \mathcal{S}(B) \rightarrow \mathbf{Set}$$

from the groupoid $\mathcal{S}(A)$ of symmetries of A , to that on B – the functorial action is an elaboration of Lemma 7.2.7. With significant effort with Hugo Paquet and Federico Olimpieri, we proved in [Clairambault et al., 2023a] that this yields a pseudofunctor from a cartesian closed bicategory of thin concurrent games (with visible strategies, so as to cope with deadlocks) into that of generalized species of structure. It turns out that up to isomorphism, generalized species of structure can be seen as computing positive witnesses in the sense of concurrent games, along with explicit positive symmetries! This is striking, as species of structure were invented between any of those. This also shows that the cardinality mismatch between the weighted relational model and generalized species of structure exactly amounts to the number of positive symmetries.

This is a new step in the relational collapses of games model, refining earlier work in that it is proof-relevant, and operates at the bicategorical level. It also emphasizes that the combinatorial core of proof-relevant models is their handling of symmetries, and there is an incredible wealth of mathematical structures to be worked out here, in which thin concurrent games play a central role. This is also strongly connected to the coefficients appearing in the Taylor expansion of λ -terms, see the next section.

13.4 Game Semantics as Taylor Expansion

This line of work is in collaboration with Lison Blondeau-Patissier, as part of her PhD thesis (2021-2024), and with Lionel Vaux Auclair.

13.4.1 Pointer concurrent games and positional injectivity

As we have seen, thin concurrent games provide an expressive representation of the interactive behaviour of programs, but at the cost of a significant technical complexity – a lot of which comes from copy indexing and uniformity. Could we conceivably simplify this by adopting an “Hyland-Ong” style representation, with pointers? It turns out that we can to an extent, losing some features such as non-deterministic branching, but gaining some such as a compelling connection with the Taylor expansion of λ -terms.

The idea is the following: instead of a concurrent strategy σ , we retain only its set of symmetry classes of (+-covered) configurations, which must hence be defined directly, without appealing to all the technology of Part II. To do that we start with:

Definition 13.4.1. Consider A an arena in the sense of Definition 3.1.3.

An *exploration*⁶ $x \in \mathcal{E}(A)$ of arena A is $x = \langle |x|, \leq_x, \partial_x \rangle$ such that $\langle |x|, \leq_x \rangle$ is a finite forest, and the *display map* $\partial_x : |x| \rightarrow |A|$ is a function satisfying the conditions:

- minimality-respecting: for any $a \in |x|$, a is \leq_x -minimal iff $\partial_x(a)$ is \leq_A -minimal,
- causality-preserving: for all $a_1, a_2 \in |x|$, if $a_1 \rightarrow_x a_2$ then $\partial_x(a_1) \rightarrow_A \partial_x(a_2)$.

If A is a mixed board, then explorations of A are analogous to configurations of A : they serve as concrete representatives for symmetry classes. A **symmetry** between explorations $x, y \in \mathcal{E}(A)$ is a bijection $\varphi : |x| \simeq |y|$ that preserves and reflects all structure. This lets us define **positions** as symmetry classes of explorations – if A is a mixed board without conflict, those are in one-to-one correspondence with $\mathcal{E}_{\cong}(A)$.

But one should not immediately quotient explorations to positions, because doing so blurs out the identity of individual moves, preventing us from enriching it with the *dynamic* causality from the program. Instead those are added on *explorations*:

Definition 13.4.2. An *augmentation* on arena A is a tuple $q = \langle |q|, \leq_q, \leq_{(q)}, \partial_q \rangle$, where

⁶In the published papers, we call those *configurations*. Here, I use *exploration* instead to avoid the collision with the configurations of an event structure.




Figure 13.9: A configuration




Figure 13.10: An isogmentation

$\langle |q| \rangle = \langle |q|, \leq_{\langle |q| \rangle}, \partial_q \rangle \in \mathcal{E}(A)$, and $\langle |q|, \leq_q \rangle$ is a forest⁷ satisfying:

rule-abiding:	for all $a_1, a_2 \in q $, if $a_1 \leq_{\langle q \rangle} a_2$, then $a_1 \leq_q a_2$,
courteous:	for all $a_1 \rightarrow_q a_2$, if $\text{pol}(a_1) = +$ or $\text{pol}(a_2) = -$, then $a_1 \rightarrow_{\langle q \rangle} a_2$,
deterministic:	for all $a^- \rightarrow_q a_1^+$ and $a^- \rightarrow_q a_2^+$, then $a_1 = a_2$,
+ -covered:	for all $a \in q $ maximal in q , we have $\text{pol}(a) = +$,
negative:	for all $a \in \min(q)$, we have $\text{pol}(a) = -$,

we then write $q \in \mathcal{A}(A)$, and call $\langle |q| \rangle \in \mathcal{E}(A)$ the **desequentialization** of q .

An augmentation equips an exploration of the game with the causal dependency due to the program. In particular, if $\sigma : A$ is a causal strategy on a mixed board A , then any $x^\sigma \in \mathcal{E}(\sigma)$ gives rise to an augmentation $q \in \mathcal{A}(A)$ on \underline{A} by setting $|q| = x^\sigma$, $s_1 \leq_q s_2$ iff $s_1 \leq_\sigma s_2$, $s_1 \leq_{\langle |q| \rangle} s_2$ iff $\partial_\sigma s_1 \leq_A \partial_\sigma s_2$, and for all $s \in |q|$, $\partial_q s = \text{bl}_A(\partial_\sigma s)$.

In thin concurrent games, all events have a distinct identity and yield a distinct image in the game with distinct copy indices, as illustrated in Figure 13.9; this rigidity is then tamed by 2-cells mediating between strategies. In contrast, in pointer concurrent games, individual events have no identity; just like elements of a multiset have no individual identity. This is illustrated in Figure 13.10, but it is not properly captured yet by Definition 13.4.2 – it remains to quotient out the identity of events. For that, a **symmetry** between augmentations q and p is a bijection $\varphi : |q| \simeq |p|$ preserving and reflecting all structure. Finally, an **isogmentation** is a symmetry class of augmentations; and then *strategies* should morally be certain aggregations of isogmentations.

Our original motivation in developing these definitions was not really to define a fully-fledged variant of concurrent games, but rather to obtain the tools we needed to prove an *injectivity* theorem [Blondeau-Patissier and Clairambault, 2021] along the following lines. Consider $\sigma, \tau : A$ two total finite innocent strategies in the Hyland-Ong sense. Each play $s \in \sigma$ reaches a *position* obtained by simply forgetting the chronological order; so that for each strategy we can retain its set of positions only. Now the theorem is that if σ and τ reach the same positions, then they are equal. Now proving

⁷For now, we have only explored pointer concurrent games in a sequential innocent setting – of course, there is no reason why the methodology would not extend further.

that involves a *pumping* technique similar to that used in traditional injectivity results in the linear logic literature [de Carvalho and de Falco, 2012], which needed understanding the effect of Opponent duplications on reached positions – and in turn, this invited a representation of plays factoring out Opponent’s scheduling, *i.e.* isogmentations⁸.

13.4.2 Isogmentations and the resource calculus

It is natural to see in isogmentations a variation of *resource terms*, the terms of the finitary calculus called *resource calculus* used as target for the Taylor expansion of λ -terms [Ehrhard and Regnier, 2008] – this observation was nothing new, the similarity between resource terms and plays up to Melliès’ homotopy had already been remarked before [Melliès, 2006, Tsukada and Ong, 2016]. For instance, along this similarity, the isogmentation in Figure 13.10 could be naturally be written syntactically as

$$\lambda f^{o \rightarrow o}. f [\lambda x^o. x, \lambda x^o. x] [\lambda x^o. x].$$

More precisely, there is a bijection between isogmentations on (the arena for) a simple type A , and β -normal resource terms in a simply-typed extensional resource calculus – moreover the correspondence is straightforward, closely following the standard finite definability argument for innocent strategies (see Theorem 3.2.14) – but this correspondence does not account for non-normal terms and for the dynamics of resource terms.

To account for these dynamics, one difficulty is that isogmentations do not easily compose. Indeed, we expect isogmentations $q \in \mathcal{A}(A \vdash B)$ and $p \in \mathcal{A}(B \vdash C)$ to be *composable* if, writing $\langle q \rangle = x_A \vdash x_B$ and $\langle p \rangle = y_B \vdash y_C$ for the corresponding positions, we have $x_B = y_B$ – so that we may write $\langle q \rangle = x_A \vdash x_B$ and $\langle p \rangle = x_B \vdash x_C$. In that case, up to symmetry we can pick concrete representatives for q and p which share the same events on B , letting us synchronize them. However, this idea bumps against an issue similar to that exposed in Example 13.3.1: the resulting composition depends on the representative chosen... To eliminate this dependency on the representative, we first define a composition $p \odot_\theta q$ through a symmetry $\theta : x_B \cong_B y_B$, and then set

$$p \odot q = \sum_{\theta : x_B \cong_B y_B} p \odot_\theta q$$

a *formal sum* of isogmentations. This observation – that the composition of isogmentations yields a formal sum of isogmentations, rather than a single isogmentation – is reminiscent of the fact that substitution in the resource calculus also yields a sum

$$s\langle [u_1, \dots, u_n]/x \rangle = \sum_{\pi \in S(n)} s\langle u_{\pi(1)}/x_1 \rangle \dots \langle u_{\pi(n)}/x_n \rangle$$

if x has n occurrences in s (the result is 0 otherwise) – where x_1, \dots, x_n is an enumeration of all occurrences of x in t . This prompted us to extend the correspondence between isogmentations and normal resource terms to the *dynamics* of the resource calculus.

⁸In other terms, isogmentations are plays up to Melliès’ homotopy equivalence on plays [Melliès, 2006].

This is what we achieve in [Blondeau-Patissier et al., 2023b]: we construct an interpretation of (simply-typed, extensional) resource terms as **strategies**, *i.e.* formal sums of isogmentations weighted by non-negative rational coefficients. For normal forms, this coincides with the bijection between normal resource terms and isogmentations mentioned above. But for non-normal terms, we show that the interpretation is invariant under reduction – also accounting for the coefficients arising throughout reduction. The proof of this result is structured by a new categorical model we call a **resource category**. Resource categories axiomatize the structure available in our category of pointer concurrent games. Much of their structure resembles notions of differential categories [Blute et al., 2020], but without an exponential modality.

13.4.3 Taylor expansion, extensional resource terms

The correspondence between normal resource terms and isogmentations invites the conjecture that the game semantics of a λ -term *is* the normal form of its Taylor expansion.

But the Taylor expansion typically targets the *pure* λ -calculus, for which this is not quite true: game semantics is an extensional model, while the Taylor expansion of a λ -term follows its *Böhm tree* (*i.e.*, the normal form of the Taylor expansion of a λ -term M is the Taylor expansion of its Böhm tree, as proved in [Ehrhard and Regnier, 2003]). So what we did instead is to devise a variant of the untyped resource calculus, the *extensional resource calculus*, obtained by seeking a syntax for the isogmentations on the universal arena of [Ker et al., 2002]. The extensional resource calculus has a nice syntactic theory, and is the target of an *extensional Taylor expansion* that captures \mathcal{H}^* [Blondeau-Patissier et al., 2023a]. At the date of writing, the correspondence of this extensional Taylor expansion with pointer concurrent games is still work in progress.

13.5 Operational Concurrent Games

This line of work is in collaboration with Simon Castellan.

Operationalizing a denotational model. As hopefully demonstrated in this monograph, thin concurrent games offer a very expressive intensional presentation of the interactive behaviours of programs. The work presented here already accounts for complex programming features, and there is no obstacle in principle to extend it further.

Yet, something is puzzling. More than any other denotational semantics, game semantics has a strong operational flavour. To the trained game semanticist, each play or configuration of the interpretation of a program matches an operational execution trace: it is this informal connection that guided our introduction to game semantics in Chapter 3. Yet, actually *proving* that a concrete play or configuration is in the interpretation is discouragingly – and embarrassingly – tricky: one must unfold the definition of the interpretation of a λ -term in a cartesian closed category, which is itself obtained as the Kleisli category of a model of intuitionistic linear logic. Overall, the interpretation of a program unfolds to a large algebraic expression in the language of the categorical model




Figure 13.11: Is this a valid configuration for this program?

– for instance, the program of Figure 13.11 (partially) unfolds to an expression like

$$\text{seqo} \left((\text{evo}((\llbracket F \rrbracket \otimes \llbracket x := !y; F(y := 1) \rrbracket^!) \circ (\delta \otimes \text{id} \otimes \text{id}))) \otimes \llbracket !x \rrbracket \right) \circ (\text{id} \otimes \delta \otimes \text{id}) \circ (\text{id} \otimes \text{cell} \otimes \text{cell}),$$

and answering the question in Figure 13.11 means putting forward an analogous configuration for each of the components of this expression, and proving they globally aggregate to the configuration claimed: this is not problematic in principle, but not efficient or even tractable in practice, so it is an obstacle to using game semantics to reason on concrete programs. Game semantics papers are full of examples, yet I feel confident in claiming that none but a few trivial ones have ever been proved formally⁹.

It seems we must choose between an interpretation that is compositional and one that it obtained directly operationally from the program. Or, we could try to devise an alternative – operational – method to generate the denotational object. Connecting operational and game semantics in that way is far from a new idea. Though it was clear from early work on game semantics [Danos et al., 1996] that interaction between strategies was related to execution by abstract machines, the idea to generate strategies by operational means was – to our knowledge – first suggested by Laird in his trace semantics for higher-order references [Laird, 2007] (with the explicit connection with game semantics later worked out by Jaber [Jaber, 2015]). In the 2010s, several works were proposed blurring the lines between operational and game semantics [Ghica and Tzevelekos, 2012, Levy and Staton, 2014], typically via LTSs dealing with open programs by sending and receiving messages from the environment; for these systems, compositionality is a theorem rather than a definition. But these developments do not yield strategies in the sense of previously established models. They also only deal with sequential deterministic programs, and it seems hard to extend LTS-based techniques to give a causal account of the execution of higher-order concurrent programs.

⁹This observation is meant to apply to all the game semantics literature, not just concurrent games.

But in fact, a powerful connection between operational and denotational semantics was already established significantly before the above, although not then presented as such: it is Baillot’s result [Baillot, 1999] that Girard’s *Geometry of Interaction* for IMELL generates the strategy obtained as its interpretation in so-called Abramsky-Jagadeesan-Malacaria (AJM) games [Abramsky et al., 2000]. We would like a similar correspondence for concurrent games as presented in this monograph. But how so? Geometry of Interaction is based on linear logic proof nets and Baillot’s theorem relies on the rewriting theory of proof nets, but there are no such tools for IA_{\parallel} .

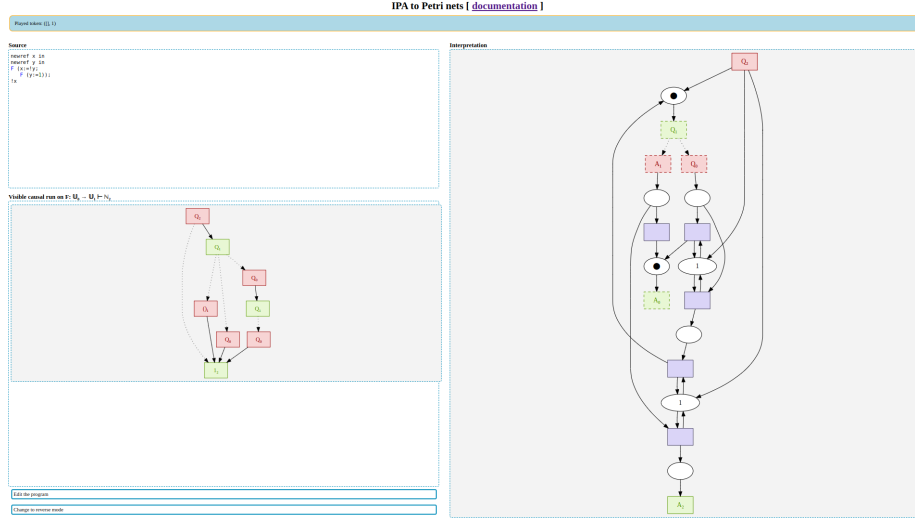
Geometry of Interaction via Petri nets. In [Castellan and Clairambault, 2023] we ask the question: can we provide an analogue of Baillot’s theorem for IA_{\parallel} , *i.e.* a notion of execution for IA_{\parallel} programs that generates, just as a side effect of execution, the *same* concurrent strategy than that computed denotationally?

Our idea is the following: while there are no proof nets for IA_{\parallel} , a proof net *equipped with the GoI token game* is an instance of an object that is very well-known in the verification and concurrency communities: a *coloured Petri net*. So we sidestep proof nets and aim to translate IA_{\parallel} programs directly into Petri nets. More precisely, we devise a notion we call a *Petri strategy*: a Petri strategy on game A is a (finite) coloured Petri net equipped with special transitions called *negative* and *positive* (on top of the old transitions now dubbed *neutral*). While neutral transitions have the same familiar mechanics as usual, negative and positive transitions are different: negative transitions have no precondition, instead triggered by an Opponent move in A and introducing a new token into the net. Dually, positive transitions have no postcondition, instead playing a positive move in A . Now, Petri nets satisfying adequate conditions support a notion of *unfolding* into event structures: each run has an implicit causal structure and those can be aggregated into an event structure. Accordingly, a *Petri strategy* is additionally required to unfold to a valid concurrent strategy. Finally, Petri strategies σ and τ can be easily composed, by putting them side by side and merging adequately matching negative and positive transitions. Altogether, we obtain a compositional interpretation of IA_{\parallel} into Petri strategies and an unfolding to concurrent strategies, such that

$$\begin{array}{ccc}
 \text{IA}_{\parallel} & & \\
 \downarrow \llbracket - \rrbracket & \searrow \llbracket - \rrbracket & \\
 \text{PetriStrat} & \xrightarrow{\text{Unf}} & \text{CG}
 \end{array}$$

commutes. The interpretation of IA_{\parallel} into Petri nets is linear time, and can be regarded as a syntactic translation – a compilation – to an intermediate representation.

This is implemented and available at <https://ipatopetrinets.github.io/>. A screenshot of the application appears in Figure 13.12: on the right hand side we have the Petri strategy, with which the user can interact by clicking on available transitions. On the upper left corner is the program, and below is the configuration of the corresponding concurrent strategy reconstructed by the unfolding. The configuration is that of Figure 13.11 which, by our main result, proves that this is indeed a configuration of

Figure 13.12: $IA_{//}$ to Petri nets

the concurrent strategy interpreting this program. Thanks to the causal nature of our interpretation, there is also a *reverse mode* that let us *undo* events, not necessarily in the opposite order in which they were played.

13.6 Miscellaneous

Finally, in this section, I give an account of other works that stand a bit on the side: though interesting, they are not part of the main thrust of our work on concurrent games.


13.6.1 Full Abstraction for Parallel-Or

This is joint work with Simon Castellan and Glynn Winskel.

In Chapter 10, I have presented our notion of parallel innocence, culminating in Chapter 12 to finite definability and intensional full abstraction for $PCF_{//}$. But in fact, our original objective was *not* $PCF_{//}$, but PCF plus *parallel-or*, the famous parallel primitive

$$\begin{aligned} \mathbf{por\ ff\ ff} &= \mathbf{ff} \\ \mathbf{por\ tt\ \perp} &= \mathbf{tt} \\ \mathbf{por\ \perp\ tt} &= \mathbf{tt} \end{aligned}$$

introduced by Plotkin in order to obtain a finite definability result with respect to Scott domains and continuous functions [Plotkin, 1977]. A central motivation in the original development of game semantics was precisely to ban **por** and obtain a cartesian closed category of sequential functions, and thus it felt deliciously ironic to attempt to close the

Figure 13.13: The concurrent strategy $\text{por} : \mathbf{B}_1 \multimap \mathbf{B}_2 \multimap \mathbf{B}_3$

loop and capture the intensional behaviour of programs from $\text{PCF} + \text{por}$. Unfortunately, back in 2014, we bumped against a tricky problem and could not quite make this work.

Let us describe the problem. There is a natural concurrent strategy for **por**, described in Figure 13.13. When prompted by Opponent, the strategy plays two causally independent moves, prompting the evaluation of its two arguments in parallel. If both arguments return **ff**, then the strategy merges those two threads and returns **ff**. However, it suffices that *one* argument returns **tt** for the overall computation to return **tt**. What may come as a surprise is that although the extensional behaviour of **por** is functional (and thus deterministic), its intensional behaviour as given by this strategy is not. Indeed, each argument returning **tt** must cause an independent instance of the answer **tt** at toplevel. But those must conflict, since they match the same move in the game! In other words,


por tt tt

triggers an (invisible) race as the two threads compete to provide the same answer – operationally, **por tt tt** can be proved via two distinct rules, also witnessing this race.

How to formulate a notion of determinism accepting **por**, but refusing an actual non-deterministic choice? We could not figure out how to solve this back in 2014, hence reverted in [Castellan et al., 2015] to giving a model giving account of (causally deterministic) parallel evaluation in plain PCF.

Disjunctive causality. In [Castellan et al., 2017b] we solved that, constructing an intensionally fully abstract model for (an affine version of) PCF with **por**. The core of our solution is a variant of **CG** in which the semantics of **por** is *deterministic*.

The reasons why **por** is deterministic run quite deep: it boils down to the inability of event structures to express *disjunctive causality*, i.e. causal patterns such as



where the *same* event 3 may be caused *either* by 1 or by 2. But in fact, this is not a limitation of event structures in general, only of event structures of the kind we consider in this monograph, which are known as *prime event structures* in the literature. There are alternatives: for instance, *general event structures* [Winskel, 1980] allow events to be enabled in several distinct ways. We give an equivalent notion [Winskel, 1986]:

Definition 13.6.1. A **configuration family** is a pair $(|A|, A)$ (often just written A) where $|A|$ is a set of events, and A is a set of **configurations**, finite subsets of $|A|$ satisfying:

- rooted: $\emptyset \in A$
- complete: for $x, y \in A$, if $x \uparrow y$ (they are **compatible**, i.e. there exists $z \in A$ such that $x \subseteq z$ and $y \subseteq z$), then $x \cup y \in A$.
- coincidence-free: if $x \in A$, for all distinct $e_1, e_2 \in x$ there exists $y \subseteq x$ such that $y \in A$ and $e_1 \in y \Leftrightarrow e_2 \notin y$.


Configuration families do support disjunctive causality: the causal pattern of (13.3) is expressible with the configuration family with events $\{1, 2, 3\}$ and the configurations

$$\{\emptyset, \{1\}, \{2\}, \{1, 3\}, \{2, 3\}, \{1, 2, 3\}\},$$

and accordingly, *parallel-or* can be formalized as a configuration family that is *deterministic* in a suitable sense. So – why is this monograph entirely phrased in terms of configuration families, rather than prime event structures? The issue with configuration families is that they do not support *hiding* as is required in the composition of strategies:

Proposition 13.6.2. *Configuration families do not support hiding. There is a configuration family A and $V \subseteq |A|$ s.t. $\{x \cap V \mid x \in A\}$ is not a configuration family.*

Proof. Set $|A| = \{a, b, c, d\}$, and configurations those specified by the diagram



with $V = \{b, c, d\}$, i.e. we hide the event a . Then, the set of configurations obtained by hiding fails completeness: it contains $\{b\}, \{d\}, \{b, c, d\}$ but not $\{b, d\}$. \square

Prime event structures support hiding because each event in the event structure post-hiding comes with a *unique, minimal* causal history in the event structure pre-hiding, a property that is lost here¹⁰. The failure of hiding comes from our inability to assign canonically, to any configuration of the hidden structure, a *witness* pre-hiding.

¹⁰One could object here that perhaps we could build concurrent games on something even more permissive than configuration families. But we must retain *coincidence-freeness*, because they are how we express the absence of causal loops (which distinguish concurrent games from relational semantics as illustrated in Section 10.4): any configuration should be reachable adding events one by one. Rooted coincidence-free families do support hiding, but they yield a non-associative composition of strategies.

Observational determinism. The key idea of our contribution is that we can recover hiding on configuration families, by requiring an additional *determinism* hypothesis:

Definition 13.6.3. A *deterministic configuration family with polarities (dcfp)* is a configuration family A with $\text{pol}_A : |A| \rightarrow \{-, 0, +\}$, satisfying the additional condition:

disjunctive deterministic: for all $x \subseteq_A^{\geq 0} y$ and $x \subseteq z$, we have $y \cup z \in A$.

where $x \subseteq_A^{\geq 0} y$ means that $x \subseteq y$ and $\text{pol}_A(y \setminus x) \subseteq \{0, +\}$.

Here, the events with null polarities are those that we wish to hide, and indeed:

Proposition 13.6.4. If A is a dcfp, then its *hiding* A_\downarrow , with events the set $V = \{a \in A \mid \text{pol}(a) \neq 0\}$ and configurations those $x \cap V$ for $x \in A$, is still a dcfp.

The key idea is as follows: a configuration $x \in A_\downarrow$ may have many different witnesses, *i.e.* configurations $y \in A$ such that $y \cap V = x$ – in general forming a family $(y_i)_{i \in I}$. But thanks to disjunctive determinism, all these witnesses are *compatible*, and

$$y = \bigcup_{i \in I} y_i \in A$$

is then a *canonical witness* for x . As in prime event structures, it is this ability to assign canonically to each $x \in A_\downarrow$ a witness in A that makes hiding go through. The situation here is somewhat dual to that for prime event structures: the canonical witness is the union of all witnesses, instead of the intersection as in the construction of **CG**.

Replaying the construction of **CG** based on the ideas above yields a compact closed category **Disj** of **disjunctive deterministic strategies**, where **por** admits a deterministic interpretation. In [Castellan et al., 2017b] we then proceed to build a hybrid model pairing **Disj** and **CG**, so that we have access to both disjunctive determinism and parallel innocence, letting us prove a finite definability result and intensional full abstraction.

13.6.2 A Semantic Proof of Herbrand's Theorem

This is joint work with Aurore Alcolei, Martin Hyland and Glynn Winskel.

In a formula of (first-order) predicate logic, there is a clear causal dependency between instance of quantifiers that may be represented as a concurrent game, *e.g.*

$$\llbracket \forall x (\exists y P(x, y) \vee \forall z \neg P(x, z)) \rrbracket = \begin{array}{c} \forall x \\ \swarrow \quad \searrow \\ \exists y \quad \forall z \end{array}$$

though this representation ignores for now all propositional connectives and atomic predicates. Then, a proof intuitively induces a further dependency between instances

it has countably many copies of each quantifier – though it is often more convenient to only draw the part of the game where only existential quantifiers are replicated.

For instance, the strategy for the usual proof of the Drinker's formula is:

$$\left[\begin{array}{c} \frac{}{\vdash \neg P(c), P(y), \neg P(y) \vee \forall y P(y)} \\ \frac{}{\vdash \neg P(c), P(y), \exists x (\neg P(x) \vee \forall y P(y))} \\ \frac{}{\vdash \neg P(c), \forall y P(y), \exists x (\neg P(x) \vee \forall y P(y))} \\ \frac{}{\vdash \neg P(c) \vee \forall y P(y), \exists x (\neg P(x) \vee \forall y P(y))} \\ \frac{}{\vdash \exists x (\neg P(x) \vee \forall y P(y)), \exists x (\neg P(x) \vee \forall y P(y))} \\ \hline \vdash \exists x (\neg P(x) \vee \forall y P(y)) \end{array} \right] = \begin{array}{c} \exists x^c \quad \exists x^y \\ \downarrow \quad \nearrow \\ \forall y \quad \forall y \end{array}$$

assuming we have at least one constant symbol c – here the existential quantifier has two copies, witnessing the fact that the proof has a contraction. Instantiating quantifiers following this strategy, we obtain the quantifier-free formula

$$(\neg P(c) \vee P(y)) \vee (\neg P(y) \vee P(z))$$

which is a tautology – here as there are two instantiations for the existential quantifier, there are two copies of the subsequent sub-formula, linked with a disjunction.

In [Alcolei et al., 2018] we provide a denotational semantics for first-order classical proofs following this idea. For an existential formula $\exists x P(x)$, the interpretation of a proof corresponds exactly to a set of closed first-order terms t_1, \dots, t_n such that $P(t_1) \vee \dots \vee P(t_n)$ is a propositional tautology, recovering Herbrand's theorem. For more general formulas, we essentially recover Miller's expansion trees [Miller, 1987].

This work is original in that of all the developments considered in this monograph, this is the only one for which the games considered are not polarized (all minimal moves may not have the same polarity) and where causality is not alternating.

13.6.3 Revisiting Games Models of MALL

As mentioned in Section 6.5, (thin) concurrent games belong to a family of games model rejecting the hypothesis that plays should be totally chronologically ordered, initiated in the seminal paper [Abramsky and Melliès, 1999] and used in particular to construct a fully complete model for multiplicative additive linear logic.

Abramsky and Melliès' take on games is rather radical: their games have neither *moves* nor *plays*, but only a partial order of *positions*. More precisely, they are *dI-domains*, *i.e.* directed-complete, bounded-complete partial order satisfying two further axioms. If D is a dI-domain, we write D^\top its extension with a top element \top – it follows that D^\top is a complete lattice. If E is an event structure (so that this automatically applies to a game A in the sense of Definition 6.1.3), then the domain $\mathcal{C}^\infty(E)$ of potentially infinite configurations is such a dI-domain. If A is a game, then $\mathcal{C}^\infty(A)$ is a new presentation of that game retaining only positions, forgetting individual moves.

Accordingly, a *strategy* in the sense of Abramsky and Melliès cannot play individual moves and acts on positions instead. Applied to a position x , a strategy σ *saturates* it, returning a new position that intuitively incorporates all moves that σ is prepared to play in x (or \top if the strategy is undefined on the current position). More precisely, strategies are *continuous, stable closure operators*¹¹ – recall $f : D \rightarrow D'$ between dl-domains is **stable** iff for $x, y \in D$, if x, y are bounded then $f(x \wedge y) = f(x) \wedge f(y)$.

Definition 13.6.5. A *closure-strategy* on dl-domain D , written $\sigma : D$, is a **continuous stable closure operator** on D^\top , i.e. a continuous function $\sigma : D^\top \rightarrow D^\top$ which is:

- extensive: for all $x \in D^\top$, $x \leq \sigma(x)$,
- idempotent: for all $x \in D^\top$, $\sigma(\sigma(x)) = \sigma(x)$,
- stable: there is a stable function $f : D \rightarrow D$ satisfying that for all $x \in D$ such that $\sigma(x) \neq \top$, $\sigma(x) = x \vee f(x)$.

Stable implies *extensive*, however we state it separately as *extensive* and *idempotent* together they define a *closure operator*, a standard notion independently of stability.

The strategies may be composed; and Abramsky and Melliès construct a category **ClosOp** of dl-domains and closure-strategies, using which they build a fully complete model of linear logic. Naively, it might seem full completeness follows by “fixing” the non-associativity of composition in Blass games [Blass, 1992], and thus reinstating the equational laws expected of a model of MALL. In truth, the situation is more subtle: Abramsky and Melliès (implicitly) construct a model of a *polarized* version of MALL. It is then the *positionality* of strategies in **ClosOp** that allows the depolarization to MALL to induce a congruence, reinstating the desired equational laws and obtaining full completeness, as made more explicit in later work by Melliès [Melliès, 2005].

In [Clairambault, 2023], I revisit this story in the context of concurrent games on event structures. A lot of that paper consists in a synthesis of old ideas, already existing in the literature (though with a different technical underpinning).

From event structures to closure operators. But this article also has original contributions, with in particular the first proper treatment of the relationship between move-based game semantics and Abramsky and Melliès’ model. Though the intuition seems clear, the definition is not so obvious (several incorrect attempts appear in the literature).

Given a (causally deterministic) concurrent strategy $\sigma : A$, we must define

$$\text{Clos}(\sigma) : \mathcal{C}^\infty(A) \rightarrow \mathcal{C}^\infty(A)$$

a closure-strategy. Given $x_A \in \mathcal{C}^\infty(A)$, by causal determinism, there is at most one $x^\sigma \in \mathcal{C}^\infty(\sigma)$ such that $\partial_\sigma x^\sigma = x_A$. If we can find such x^σ , then we may saturate it by taking a maximal $x^\sigma \subseteq^+ y^\sigma$, and set $\text{Clos}(\sigma)(x_A) = y_A^\sigma$. But what should we do if there is *no* $x^\sigma \in \mathcal{C}^\infty(\sigma)$ such that $\partial_\sigma x^\sigma = x_A$? As it turns out, the correct answer is to play all positive moves of σ enabled in x_A , *even though* x_A might have positive moves that σ will never play. If doing so yields a conflict, then we return \top .

¹¹Some additional conditions appear in the course of [Abramsky and Melliès, 1999], omitted here.

Chapter 14

Conclusions

14.1 The Work Done so Far

What do the following works have in common?

- Hyland and Ong’s fully abstract games model for PCF [Hyland and Ong, 2000],
- Abramsky and McCusker’s fully abstract games model for IA based on alternating non-innocent strategies [Abramsky and McCusker, 1996],
- Abramsky and Melliès’ fully complete model for MALL based on closure operators [Abramsky and Melliès, 1999],
- The relational model and its syntactic presentation, non-idempotent intersection types [Girard, 1988],
- Ghica and Murawski’s fully abstract model for IA_{\parallel} [Ghica et al., 2006],
- The weighted relational model [Laird et al., 2013],
- Alternating probabilistic game semantics [Danos and Harmer, 2000],
- Generalized species of structure [Fiore et al., 2008],
- The Taylor expansion of λ -terms [Ehrhard and Regnier, 2008],
- The GoI token machines [Danos and Regnier, 1996, Mackie, 1995].

The results presented in this monograph, accounting both for the work presented in details in Parts II and III and that presented more synthetically in Chapter 13, show that all these models are essentially obtained from concurrent games via forgetful operations: they are presentations of certain aspects of concurrent strategies. This list is of course a tiny fragment of denotational semantics, but this already gives a sign that we are getting closer to the *full interactive behaviour* mentioned in the introduction.

Our point is certainly not that this makes all these models obsolete. Concurrent games are too intensional for many purposes, and syntactic methods such as non-idempotent intersection types or Taylor expansion, by working directly on syntax, avoid the overhead brought by the categorical machinery of denotational semantics. More generally, different purposes call for different tools; concurrent games unifies these models but with the price of a high technical complexity which is not always desirable.

In contrast, we argue that this work strengthens all these models by inscribing them within a common semantic landscape and making them part of a joint mathematical theory. Many questions on their strengths, limitations and potential extensions can be attacked under the lense of concurrent games: for one example, how should non-idempotent intersection types and Taylor expansion of programs best deal with branching effects (such as non-determinism and parallelism) or quantitative effects (such as probabilities)? Concurrent strategies support all this and extend Taylor expansion conservatively, so the question above is reduced to the following: can concurrent games be described more syntactically, in the style of intersection types or Taylor expansion?

14.2 Open Problems

Next, here are some open problems around the theory presented in this monograph. Those are not necessarily long-term research objectives or directions (though some could be), but rather technical difficulties encountered during the technical developments, that we have not been able to solve and have had to avoid or work around.

14.2.1 Non-deterministic parallel innocence

In Chapter 10 we have developed the theory of *parallel innocent* causal strategies; but our notion of innocence (Definitions 10.1.8 and 10.1.11) smuggles in *causal determinism*, i.e. the absence of minimal conflicts between Player moves. But parallel innocence is intended to ban state and more generally side-channel interference between threads, there is no reason why it would necessarily enforce determinism. For instance, what are the strategies generated by PCF_{\parallel} plus a non-deterministic choice operator \Diamond ?

First of all, in the case of *sequential innocence* there is no problem, as spelled out in Definition 13.2.1: the problem appears only for non-deterministic *parallel* innocence. First, the conditions of Definition 13.2.1 do not work as-is: *locally sequential* is wrong for obvious reasons (as it forbids parallel branching); *locally conflicting* forbids parallel- or (see Figure 13.13) which should certainly be parallel innocent.

Recall that parallel innocence started with *pre-innocence* (Definition 10.1.8), which follows the idea that Player cannot causally merge threads forked by Opponent. It is natural to attempt to follow the same idea for conflict, and define:

Definition 14.2.1. For A a board, a visible $\sigma : A$ is **#-pre-innocent** iff

#-pre-innocence: If $m_1^+ \rightsquigarrow_{\sigma} m_2^+ \in \sigma$ and $\rho_1 \rightarrow m_1^+, \rho_2 \rightarrow m_2^+ \in \text{gcc}(\sigma)$,
then $\min(\rho_1) = \min(\rho_2)$ and their least distinct moves are positive.





Figure 14.1: The subtlety of non-deterministic parallel innocence

In other words, *Player cannot impose conflict between threads forked by Opponent*.

It is quite intuitive that strategies obtained from PCF_{\parallel} plus \odot must satisfy this, as \odot can only be used to generate conflict between moves that share the same causal predecessors. And for quite some time, we were convinced we had a proof that this (in the presence of visibility) was stable under composition. With that condition, it is then tempting to define parallel innocence without causal determinism but with $\#$ -pre-innocence, so that it also covers non-deterministic parallel innocence.

However, $\#$ -pre-innocence is not stable under composition. The strategy

$$(o \multimap o) \vdash o \multimap o \multimap (o \multimap o \multimap o) \multimap o$$



is visible, pre-innocent and $\#$ -pre-innocent, but its composition with the identity on o is not, as the minimal conflict is transported to the two final positive moves, though the gccs leading to them are forked by Opponent. Therefore the open question is:

Open question 14.2.2. Define a condition of non-deterministic parallel innocence.

Non-deterministic parallel innocence should compose, and support the interpretation of PCF_{\parallel} with \odot , and of parallel-or. Additionally, non-deterministic parallel innocent strategies which are causally deterministic should be parallel innocent in the sense of this monograph. Non-deterministic parallel innocent strategies which are sequential should be non-deterministic sequential innocent as in Definition 13.2.1.

In other word, the problem is to disentangle non-determinism and parallelism.

Note that non-deterministic parallel innocence is not going to be conservative with respect to deterministic parallel innocence (which probably means that we do not have

deterministic parallel innocence quite right either) as illustrated in Figure 14.1. By composing a deterministic parallel innocent strategy with a sequential non-deterministic innocent strategy (that is obviously definable in a non-deterministic linear λ -calculus with non-deterministic choice), we break #-pre-innocence. Something is wrong with the deterministic parallel innocent strategy of Figure 14.1, but it seems that it cannot be banned with simple, local forbidden patterns as in pre-innocence and #-pre-innocence; our best guess is a global correctness criterion in the style of linear logic proof nets.

Once equipped with a satisfactory notion of non-deterministic innocence, it would be natural to further seek a corresponding extension of the sequential probabilistic innocence of Definition 13.2.2 to a full parallel probabilistic innocence.

14.2.2 Proper determinism

An issue strongly related to proper non-deterministic parallel innocence, is that of determinism. As we have seen, \mathbf{IA}_{\parallel} is a genuinely non-deterministic language, due to the combination of parallelism and state causing races in the memory. Accordingly, in Part III, our conditions of *parallel innocence* and *sequentiality* each had to reimpose determinism; but this was done in radically different ways.

Recall indeed *causal determinism*, a part of parallel innocence:

Definition 10.1.1. Consider $\sigma : A \vdash B$ a causal strategy.

We say σ is **causally deterministic** if for all $s \sim_{\sigma} s'$ in σ , $\text{pol}(s) = \text{pol}(s') = -$.

And recall also *sequential determinism*, a part of sequentiality:

Definition 11.1.4. A strategy $\sigma \in \mathbf{NTCG}(A, B)$ is **sequential** if:

- pointed: if $s \in \sigma$, there is a unique $\text{init}(s) \leq_{\sigma} s$ minimal in σ ,
- sequential determinism: if $tn_1^+, tn_2^+ \in \Downarrow\text{-}\mathcal{L}(\sigma)$, then $n_1 = n_2$,
- sequential visibility: if $s \in \Downarrow\text{-}\mathcal{L}(\sigma)$, writing $\ulcorner s \urcorner = s_1 \dots s_n$ and $1 \leq j \leq n$,
if s_j is non-minimal then there is $1 \leq i < j$ s.t. $j(s_j) = s_i$,

where the justifier $j(m)$ of $m \in \sigma$ refers to Definition 10.2.1.

Open question 14.2.3. Define a proper notion of determinism, unifying these two.

Note that such a proper determinism *should not be compositional* – indeed both \mathbf{IA} and \mathbf{PCF}_{\parallel} are deterministic languages, but their union is not. However, proper determinism should become compositional in conjunction with either non-deterministic parallel innocence and non-deterministic sequentiality.

As a reasonable starting point, it might be that in the presence of (an appropriate notion of) non-deterministic parallel innocence, *sequential determinism* implies causal determinism. Indeed, proving that amounts to showing that all Player immediate conflicts are reachable in an alternating linearization, which seems plausible, if conflicting patterns are anything like the *globules* of Definition 10.5.1.

14.2.3 All well-bracketings

Rather than the clear statement of an open problem, this is a description of an unfortunate state of affairs in this monograph: namely, the status of well-bracketing – understood broadly as a constraint on the function call/return discipline. In this monograph, we have seen four different mechanisms akin to well-bracketing.

Firstly, of course, we have seen the traditional **chronological well-bracketing** of Section 3.2.2, explicitly requiring that calls and returns, ordered chronologically, are... well-bracketed. This is of course very natural, though restricted to sequential plays.

Secondly, in non-alternating games we have seen **logical well-bracketing** (Section 5.1.1), also used in concurrent games when constructing \rightarrow -**Strat** (Section 9.3.1). Logical well-bracketing asks that a function may only call its arguments *before* it returns, and may only return once all called arguments have returned. A logically well-bracketed play may not be chronologically well-bracketed even if alternating (Figure 5.3).

Next, in the definition of **NTCG** we have deployed *winning conditions* (Section 8.2.2) imported from [Melliès and Tabareau, 2007]. This **positional well-bracketing**, avoiding the explicit distinction of Questions and Answers, treats well-bracketing as a linearity constraint: a call may only be answered once. The *winning* mechanism enforces that winning strategies are never the first to break this linearity constraint, which hits surprisingly close to chronological well-bracketing at least for sequential innocent total strategies. In general, positional well-bracketing does not subsume the other more interactive notions, and *e.g.* cannot guarantee finite definability result. On the other hand, it is by far the best way to formulate the relational collapse of Section 10.4.

Finally, we have seen *globularity*, which is a **causal well-bracketing**: it ensures that the forking of merging of threads also respects the call/return discipline for parallel innocent strategies, in that a strategy may not merge threads with calls still pending since their fork. Again this is not guaranteed by our previous notions, but is nevertheless required for the definability result of Chapter 12.

This profusion of subtly different notions of well-bracketing is messy and unsatisfactory. It feels that they could be unified in principle, but at the cost of requiring importing even more intensional information onto strategies, such as the *concurrency pointers* of [Laird, 2001b] to distinguish causal links coming from the term structure from those going through side-channel communication – this may allow us to extend causal well-bracketing beyond parallel innocent strategies. But does this actually work? And is it worth making the model even more complex? We leave this as an exercise to the reader.

14.2.4 Observing causality

Typically, concurrent games models are not fully abstract; as argued in Section 1.6 the strength of concurrent games is not to capture observational equivalence, but to record intensional information. Nevertheless, it is natural to wonder under which circumstances this causal information might be observed. From full abstraction results such as Theorem 5.1.12, this seems to boil down to asking when causal strategies may be entirely recovered (up to positive iso) from their non-alternating plays.

As presented in Section 13.1.1, the causal information *cannot* be reconstructed from plays in general. But we think it might be under the right circumstances: if each move can be replicated and the strategy is parallel innocent, then in principle causality may be observed by duplicating Opponent moves and observing the triggered duplications of Player moves. Indeed in examples, it is usually easy to define plays that separate distinct parallel innocent strategies. This brings us to the conjecture:

Conjecture 14.2.4. *Consider A a mixed board such that no moves are answers, and $\sigma, \tau : A$ two (possibly finite) parallel innocent strategies.*

If $\mathcal{U}\text{-Unf}(\sigma) = \mathcal{U}\text{-Unf}(\tau)$, then $\sigma \approx \tau$.

If this was the case, it would likely entail that two terms in PCF_{\parallel} yield *isomorphic* strategies if and only if they cannot be distinguished by contexts from IA_{\parallel} extended with a control operator. However, one should not underestimate this problem: given a non-alternating play s , there may in principle be several parts of the strategy causing it – thus it is hard to pin down what one individual play teaches us about a strategy.

Perhaps the right path to that conjecture is via another:

Conjecture 14.2.5. *Consider A a mixed board such that no moves are answers, and $\sigma, \tau : A$ two (possibly finite) parallel innocent strategies.*

If $(\sigma) = (\tau)$, then $\sigma \approx \tau$.

In [Blondeau-Patissier and Clairambault, 2021], we have already proved this for finite sequential innocent strategies. It would imply Conjecture 14.2.4, since the relational collapse can be extracted from the plays. The proof method for the sequential innocent case seems to extend in principle, but it becomes very technical and so far we have not managed to work out the details.

14.3 Perspectives and Future Directions

While the open problems above may tie loose ends and offer good avenues for short-term technical developments, they do not give long-term overarching research directions for the future. While we believe that the work reported here gives valuable steps in the direction presented in our introduction, the long-term goal is to use concurrent games to help make denotational semantics a connected mathematical theory. This suggests a *wealth* of research directions for the future, and we now introduce some of them.

The cube of branching structure. The core contribution of this monograph is the *disentangling of parallelism and state*, leading to a semantic square in the style of that presented in Section 3.3.3. This fits together parallelism and state in a single model, but why only those? After all, there are multiple other programming features studied in the game semantics literature – to cite only a few: control operators, exceptions, higher-order references, coroutines, probabilistic choice, etc. In principle, all of those semantics could be reformulated within concurrent games, but it does not seem such a good research strategy to do it exhaustively for all of them – if only because such a massive undertaking would prevent us from investigating other interesting questions.




Figure 14.2: The cube of branching structures

So we introduce the following criterion. Some of the features above change the rules following which control flows through the program, but do not change the *geometry* of the control flow – this includes control operators, exceptions, higher-order references and coroutines: adding any of those to a sequential deterministic programming language yields a sequential deterministic programming language.

In contrast, some programming features heavily affect the geometric shape of execution: in this monograph, we have investigated in depth *parallelism*, and also *non-determinism* arising from the combination of parallelism and state. *Non-determinism* makes execution a tree, *parallelism* makes it a partial order, and the two together turn execution into an *event structure*. This makes concurrent games based on event structures perfect for this combination; we argue it is their ability to represent explicitly this *branching structure* that made possible the results presented in Part III. We expect giving a proper treatment of a new branching structure is a more fundamental problem than dealing with other programming features: after all, most of the non-branching effects mentioned above did *not* require changing much of the basic sequential deterministic game semantics canvas; unlike the results presented in this monograph.

There is a major branching structure that is not properly covered within this monograph: *probabilistic choice*. By itself, it differs from non-deterministic choice by the addition of quantitative weights representing probabilities (as described in Section 13.2). But more fundamentally, it is the branching structure in the presence of all three of parallelism, non-determinism and probabilistic choice that poses the real challenge. This combination of branching effects is embodied by the cube in Figure 14.2.

Its left hand face is that studied in Part III. Its front face seems easy to handle working from the results of Section 13.2. The real challenge is in dealing with the all three effects together, which requires us to elaborate concurrent strategies to deal with the combination of non-determinism and probabilistic choice – perhaps based on de Visme’s *mixed event structures* [de Visme, 2019]. But besides having to develop a new theory of con-

current games, constructing a semantic cube to match Figure 14.2 means we *must* solve the open problems of Section 14.2.1 and give proper notions of probabilistic and non-deterministic parallel innocence. So this is a long-term objective that will require many developments before we can properly tackle it.

Call-by-value and evaluation order. A major element missing from all the development presented in this monograph, is evaluation order. All the languages considered here are call-by-name. This is mainly for historical reasons, because we connect to earlier work in the literature which is set in call-by-name. There is no obstacle, in principle, to replaying the results presented here in call-by-value: in particular, it has been known for a long time how to build games models for call-by-value languages [Abramsky and McCusker, 1997, Honda and Yoshida, 1999]. Our point here is not to redo all the previous work in call-by-value, as this would be a major undertaking offering comparatively little conceptual insight – though some of it does need to be redone.

There are several frameworks for languages and models encompassing both call-by-name and call-by-value: the main example is of course Levy’s call-by-push-value [Levy, 1999] which focuses on the impact of evaluation order on effects, but also others such as the *bang calculus* (inspired from Girard’s translations of intuitionistic logic into linear logic), which instead captures evaluation order via its impact on resource consumption [Ehrhard and Guerrieri, 2016]. Or finally, system L [Curien et al., 2016], encompassing both. This latter system, to our knowledge, does not yet have a concrete games model, which is unfortunate since it is the family of models that gives a concrete account of both effects and resources. A concurrent games analysis of such systems could be enlightening and also help us understand how *parallelism* fits into this picture.

In fact, there is currently increasing research activity revisiting classical topics (such as the pure λ -calculus, solvability, Böhm trees etc) in call-by-value. It seems that adequate notions of infinite normal forms in call-by-value include permutative equivalences inducing the structure of a DAG, rather than simply a tree [Accattoli et al., 2023], meaning that it might naturally yield a parallel innocent concurrent strategy. A natural question is that of full abstraction for the pure λ -calculus in call-by-value, which is still missing; drawing inspirations from the results of this monograph this also offers enticing areas of investigation, connecting concurrent games with Taylor expansion [Kerinec et al., 2020] and intersection types [Arrial et al., 2023] in call-by-value.

Finally, a major objective is to extend our geometry of interaction multi-token machine presented in Section 13.5 to a call-by-value language – this is a prerequisite before we can seriously investigate if it might be useful in practice. We have preliminary ideas about how such a machine should work exploiting subtle memoization mechanisms in Petri nets, but much ground work is needed: in call-by-name, the validity of the machine rests on the adequate concurrent games model of \mathbf{IA}_{\parallel} presented in Chapter 9, and no such model exists yet for a higher-order call-by-value stateful concurrent language.

Taylor expansion and intersection types. As reported in Section 13.4.3, for the pure λ -calculus, concurrent games offer a representation of programs very close to that proposed by the Taylor expansion of λ -terms [Ehrhard and Regnier, 2008] – more pre-

cisely, concurrent strategies form what could be regarded as a very rigid Taylor expansion. With respect to game semantics, the Taylor expansion of λ -terms has the advantage of being a syntactic method. This makes it much easier to deploy, and with a much tighter connection with the original syntax. In return, concurrent games are much more general, supporting with in the same framework many additional programming features including non-determinism, concurrency, probabilistic choice, shared memory, etc.

This suggests a number of natural questions: to what extent could concurrent games models be presented syntactically beyond the pure λ -calculus? How does it compare with existing extensions of the Taylor expansion, such as for non-deterministic choice [Olimpieri and Vaux Auclair, 2022] or probabilistic choice [Dal Lago and Leventis, 2019]? Can we use concurrent games as a guide to construct intersection type systems or Taylor expansions for effectful languages (such as \mathbf{IA}_{\parallel}), giving us the expressiveness of concurrent games models without all the accompanying technical complications?

Categorification. Games models lack uniformity; while they often formalize similar ideas and follow a common methodology, different models usually rest on very different technical underpinnings. The work presented in this monograph attempts to cope with this state of affairs by linking the models together, providing functorial bridges that make explicit how they relate to concurrent games. But a complementary methodology is to capture the construction of various games models as instances of one common categorical construction; notable works in this direction include [Eberhart and Hirschowitz, 2018] and more recently, Melliès’ *template games* [Melliès, 2019a, Melliès, 2019b]. However, for the time being, concurrent games evade this categorical unification. Melliès’ original template games lack mechanisms to eliminate deadlocking synchronizations – a new, more elaborate version [Melliès, 2021] does capture the *interaction* of concurrent strategies (see Proposition 6.2.13), but not the composition, for subtle reasons. So the question remains open: how should template games be generalized so as to encompass concurrent games?

A strongly related topic is the proper categorical treatment of *symmetries* in concurrent games. Recently [Clairambault and Forest, 2023], we constructed a cartesian closed bicategory of certain *spans of groupoids*, reproducing categorically much of the structure of symmetries in thin concurrent games (but, like template games, also failing to account for deadlocks in compositions). It is our conviction that a proper categorification of game semantics should rest on thin spans of groupoids. The link with template games is fascinating, as the two models use completely different mechanisms to account for replication of resources – thin spans are *thin* as in thin concurrent games, while template games are *saturated* (as in Section 7.1.2). Another fundamental question is the link with distributors and generalized species of structure; we give some elements in [Clairambault et al., 2023b] but this should be completed to account for inclusion of configurations and not just symmetries, and should also include a proper treatment of the linear bicategorical (or double-categorical) structure.

Thin spans of groupoids really are *thin concurrent games without concurrent games*. It seems that many of our passed achievements with thin concurrent games (typically, the enrichment with quantitative valuations and the corresponding relational collapses,

up to generalized species of structure) are actually independent from the basic game mechanisms, and should really be carried out in thin spans of groupoids instead. We should spell this out and investigate extensions made possible by the fact that we are no longer hindered by the intricate combinatorics of strategy composition.

Qualitative concurrent games. Thin concurrent games are a *quantitative* model, in the sense that they record explicitly the *multiplicity* of resource usage – even more, each individual move has a separate identity, formalized as a copy index. In categorified relational models such as distributors, it is possible to change the exponential modality to make it *qualitative*; this is done by changing the *morphisms* in $!A$ so as to allow contraction and weakening. In principle, the same could be done in thin concurrent games, by generalizing *symmetries* to encompass more general relations between configurations that may not preserve multiplicity. There are fundamental questions to solve regarding how such relations behave with respect to polarity. A first concrete objective would be to construct such a generalization of thin concurrent games, and deduce a new, more concrete proof of Ehrhard’s *extensional collapse* theorem [Ehrhard, 2012] linking qualitative and quantitative models.

One interest of such a concrete proof is that we expect it to generalize in a situation with *infinite multiplicities*. More precisely, we aim to generalize the relational collapse presented in this monograph to a collapse to the *infinitary relational model* [Grellois and Melliès, 2015], also accounting for infinite configurations. Paired with an infinitary version of Ehrhard’s extensional collapse, we expect this is the key ingredient to a purely semantic proof of the decidability of the higher-order model-checking problem, which would be a beautiful illustration of the strength of thin concurrent games.

Bibliography

- [Abramsky, 2003] Abramsky, S. (2003). Sequentiality vs. concurrency in games and logic. *Math. Struct. Comput. Sci.*, 13(4):531–565.
- [Abramsky et al., 1998] Abramsky, S., Honda, K., and McCusker, G. (1998). A fully abstract game semantics for general references. In *Thirteenth Annual IEEE Symposium on Logic in Computer Science, Indianapolis, Indiana, USA, June 21-24, 1998*, pages 334–344. IEEE Computer Society.
- [Abramsky et al., 2000] Abramsky, S., Jagadeesan, R., and Malacaria, P. (2000). Full abstraction for PCF. *Inf. Comput.*, 163(2):409–470.
- [Abramsky and McCusker, 1996] Abramsky, S. and McCusker, G. (1996). Linearity, sharing and state: a fully abstract game semantics for idealized algol with active expressions. *Electron. Notes Theor. Comput. Sci.*, 3:2–14.
- [Abramsky and McCusker, 1997] Abramsky, S. and McCusker, G. (1997). Call-by-value games. In *CSL*, volume 1414 of *Lecture Notes in Computer Science*, pages 1–17. Springer.
- [Abramsky and Melliès, 1999] Abramsky, S. and Melliès, P. (1999). Concurrent games and full completeness. In *LICS*, pages 431–442. IEEE Computer Society.
- [Accattoli et al., 2023] Accattoli, B., Faggian, C., and Lancelot, A. (2023). Normal form bisimulations by value. *CoRR*, abs/2303.08161.
- [Alcolei et al., 2018] Alcolei, A., Clairambault, P., Hyland, M., and Winskel, G. (2018). The true concurrency of herbrand’s theorem. In *CSL*, volume 119 of *LIPIcs*, pages 5:1–5:22. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.
- [Alcolei et al., 2019] Alcolei, A., Clairambault, P., and Laurent, O. (2019). Resource-tracking concurrent games. In *FoSSaCS*, volume 11425 of *Lecture Notes in Computer Science*, pages 27–44. Springer.
- [Altenkirch et al., 2010] Altenkirch, T., Chapman, J., and Uustalu, T. (2010). Monads need not be endofunctors. In *International Conference on Foundations of Software Science and Computational Structures*, pages 297–311. Springer.
- [Arnold and Niwiński, 2001] Arnold, A. and Niwiński, D. (2001). *Rudiments of μ -calculus*. Elsevier.

- [Arrial et al., 2023] Arrial, V., Guerrieri, G., and Kesner, D. (2023). Quantitative inhabitation for different lambda calculi in a unifying framework. *Proc. ACM Program. Lang.*, 7(POPL):1483–1513.
- [Baillot, 1999] Baillot, P. (1999). *Approches dynamiques en sémantique de la logique linéaire: jeux et géométrie de l'interaction*. PhD thesis, Aix-Marseille 2.
- [Baillot et al., 1997a] Baillot, P., Danos, V., Ehrhard, T., and Regnier, L. (1997a). Believe it or not, ajm's games model is a model of classical linear logic. In *LICS*, pages 68–75. IEEE Computer Society.
- [Baillot et al., 1997b] Baillot, P., Danos, V., Ehrhard, T., and Regnier, L. (1997b). Timeless games. In *Computer Science Logic, 11th International Workshop, CSL '97, Annual Conference of the EACSL, Aarhus, Denmark, August 23-29, 1997, Selected Papers*, pages 56–77.
- [Blass, 1992] Blass, A. (1992). A game semantics for linear logic. *Ann. Pure Appl. Log.*, 56(1-3):183–220.
- [Blondeau-Patissier and Clairambault, 2021] Blondeau-Patissier, L. and Clairambault, P. (2021). Positional injectivity for innocent strategies. In Kobayashi, N., editor, *6th International Conference on Formal Structures for Computation and Deduction, FSCD 2021, July 17-24, 2021, Buenos Aires, Argentina (Virtual Conference)*, volume 195 of *LIPICs*, pages 17:1–17:22. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.
- [Blondeau-Patissier et al., 2023a] Blondeau-Patissier, L., Clairambault, P., and Vaux Auclair, L. (2023a). Extensional taylor expansion. *CoRR*, abs/2305.08489.
- [Blondeau-Patissier et al., 2023b] Blondeau-Patissier, L., Clairambault, P., and Vaux Auclair, L. (2023b). Strategies as resource terms, and their categorical semantics. In Gaboardi, M. and van Raamsdonk, F., editors, *8th International Conference on Formal Structures for Computation and Deduction, FSCD 2023, July 3-6, 2023, Rome, Italy*, volume 260 of *LIPICs*, pages 13:1–13:22. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.
- [Blute et al., 2020] Blute, R., Cockett, J. R. B., Lemay, J. P., and Seely, R. A. G. (2020). Differential categories revisited. *Appl. Categorical Struct.*, 28(2):171–235.
- [Boudes, 2009] Boudes, P. (2009). Thick subtrees, games and experiments. In *Typed Lambda Calculi and Applications, 9th International Conference, TLCA 2009, Brasilia, Brazil, July 1-3, 2009. Proceedings*, pages 65–79.
- [Calderon and McCusker, 2010] Calderon, A. C. and McCusker, G. (2010). Understanding game semantics through coherence spaces. *Electr. Notes Theor. Comput. Sci.*, 265:231–244.
- [Cartwright et al., 1994] Cartwright, R., Curien, P., and Felleisen, M. (1994). Fully abstract semantics for observably sequential languages. *Inf. Comput.*, 111(2):297–401.

- [Castellan, 2015] Castellan, S. (2015). La stratégie de la fourchette. In *Vingt-sixièmes Journées Francophones des Langages Applicatifs (JFLA 2015)*.
- [Castellan and Clairambault, 2016] Castellan, S. and Clairambault, P. (2016). Causality vs. interleavings in concurrent game semantics. In *CONCUR*, volume 59 of *LIPICs*, pages 32:1–32:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.
- [Castellan and Clairambault, 2021] Castellan, S. and Clairambault, P. (2021). Disentangling parallelism and interference in game semantics. Submitted.
- [Castellan and Clairambault, 2023] Castellan, S. and Clairambault, P. (2023). The geometry of causality: Multi-token geometry of interaction and its causal unfolding. *Proc. ACM Program. Lang.*, 7(POPL):689–717.
- [Castellan et al., 2018a] Castellan, S., Clairambault, P., Hayman, J., and Winskel, G. (2018a). Non-angelic concurrent game semantics. In Baier, C. and Lago, U. D., 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.
- [Castellan et al., 2018b] Castellan, S., Clairambault, P., Paquet, H., and Winskel, G. (2018b). The concurrent game semantics of probabilistic PCF. In *LICS*, pages 215–224. ACM.
- [Castellan et al., 2017a] Castellan, S., Clairambault, P., Rideau, S., and Winskel, G. (2017a). Games and strategies as event structures. *Log. Methods Comput. Sci.*, 13(3).
- [Castellan et al., 2014] Castellan, S., Clairambault, P., and Winskel, G. (2014). Symmetry in concurrent games. In Henzinger, T. A. and Miller, D., 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.
- [Castellan et al., 2015] Castellan, S., Clairambault, P., and Winskel, G. (2015). 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.
- [Castellan et al., 2017b] Castellan, S., Clairambault, P., and Winskel, G. (2017b). Observably deterministic concurrent strategies and intensional full abstraction for parallel-or. In *FSCD*, volume 84 of *LIPICs*, pages 12:1–12:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.
- [Castellan et al., 2019] Castellan, S., Clairambault, P., and Winskel, G. (2019). Thin games with symmetry and concurrent hyland-ong games. *Log. Methods Comput. Sci.*, 15(1).

- [Chappe et al., 2023] Chappe, N., He, P., Henrio, L., Zakowski, Y., and Zdancewic, S. (2023). Choice trees: Representing nondeterministic, recursive, and impure programs in coq. *Proc. ACM Program. Lang.*, 7(POPL):1770–1800.
- [Clairambault, 2010] Clairambault, P. (2010). *Logique et Interaction : une Étude Sémantique de la Totalité. (Logic and Interaction : a Semantic Study of Totality)*. PhD thesis, Paris Diderot University, France.
- [Clairambault, 2023] Clairambault, P. (2023). A tale of additives and concurrency in game semantics. In *Samson Abramsky on Logic and Structure in Computer Science and Beyond*, pages 363–414. Springer.
- [Clairambault and de Visme, 2020] Clairambault, P. and de Visme, M. (2020). Full abstraction for the quantum lambda-calculus. *Proc. ACM Program. Lang.*, 4(POPL):63:1–63:28.
- [Clairambault et al., 2019] Clairambault, P., de Visme, M., and Winskel, G. (2019). Game semantics for quantum programming. *Proc. ACM Program. Lang.*, 3(POPL):32:1–32:29.
- [Clairambault and Forest, 2023] Clairambault, P. and Forest, S. (2023). The cartesian closed bicategory of thin spans of groupoids. In *LICS*, pages 1–13.
- [Clairambault et al., 2012] Clairambault, P., Gutierrez, J., and Winskel, G. (2012). The winning ways of concurrent games. In *LICS*, pages 235–244. IEEE Computer Society.
- [Clairambault and Harmer, 2010] Clairambault, P. and Harmer, R. (2010). Totality in arena games. *Ann. Pure Appl. Logic*, 161(5):673–689.
- [Clairambault et al., 2023a] Clairambault, P., Olimpieri, F., and Paquet, H. (2023a). From thin concurrent games to generalized species of structure. In *Proceedings of the 38th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2023, Boston, US, June 26–29, 2023*. To appear.
- [Clairambault et al., 2023b] Clairambault, P., Olimpieri, F., and Paquet, H. (2023b). From thin concurrent games to generalized species of structures. In *LICS*, pages 1–14.
- [Clairambault and Paquet, 2021] Clairambault, P. and Paquet, H. (2021). The quantitative collapse of concurrent games with symmetry. *CoRR*, abs/2107.03155.
- [Coquand, 1995] Coquand, T. (1995). A semantics of evidence for classical arithmetic. *J. Symb. Log.*, 60(1):325–337.
- [Curien, 2007] Curien, P. (2007). Definability and full abstraction. In *Computation, Meaning, and Logic*, volume 172 of *Electronic Notes in Theoretical Computer Science*, pages 301–310. Elsevier.
- [Curien and Faggian, 2005] Curien, P. and Faggian, C. (2005). L-nets, strategies and proof-nets. In *CSL*, volume 3634 of *Lecture Notes in Computer Science*, pages 167–183. Springer.

- [Curien et al., 2016] Curien, P., Fiore, M. P., and Munch-Maccagnoni, G. (2016). A theory of effects and resources: adjunction models and polarised calculi. In *POPL*, pages 44–56. ACM.
- [Dal Lago and Leventis, 2019] Dal Lago, U. and Leventis, T. (2019). On the Taylor expansion of probabilistic lambda-terms. In Geuvers, H., editor, *4th International Conference on Formal Structures for Computation and Deduction, FSCD 2019, June 24–30, 2019, Dortmund, Germany*, volume 131 of *LIPICs*, pages 13:1–13:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.
- [Danos and Ehrhard, 2011] Danos, V. and Ehrhard, T. (2011). Probabilistic coherence spaces as a model of higher-order probabilistic computation. *Inf. Comput.*, 209(6):966–991.
- [Danos and Harmer, 2000] Danos, V. and Harmer, R. (2000). Probabilistic game semantics. In *LICS*, pages 204–213. IEEE Computer Society.
- [Danos et al., 1996] Danos, V., Herbelin, H., and Regnier, L. (1996). 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.
- [Danos and Regnier, 1996] Danos, V. and Regnier, L. (1996). Reversible, irreversible and optimal lambda-machines. In *Linear Logic Tokyo Meeting*, volume 3 of *Electronic Notes in Theoretical Computer Science*, pages 40–60. Elsevier.
- [de Carvalho and de Falco, 2012] de Carvalho, D. and de Falco, L. T. (2012). The relational model is injective for multiplicative exponential linear logic (without weakenings). *Ann. Pure Appl. Log.*, 163(9):1210–1236.
- [de Visme, 2019] de Visme, M. (2019). Event structures for mixed choice. In Fokink, W. J. and van Glabbeek, R., editors, *30th International Conference on Concurrency Theory, CONCUR 2019, August 27–30, 2019, Amsterdam, the Netherlands*, volume 140 of *LIPICs*, pages 11:1–11:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.
- [Dixon et al., 2021a] Dixon, A., Lazic, R., Murawski, A. S., and Walukiewicz, I. (2021a). Leafy automata for higher-order concurrency. In *FoSSaCS*, volume 12650 of *Lecture Notes in Computer Science*, pages 184–204. Springer.
- [Dixon et al., 2021b] Dixon, A., Lazic, R., Murawski, A. S., and Walukiewicz, I. (2021b). Verifying higher-order concurrency with data automata. In *LICS*, pages 1–13. IEEE.
- [Eberhart and Hirschowitz, 2018] Eberhart, C. and Hirschowitz, T. (2018). What’s in a game?: A theory of game models. In Dawar, A. and Grädel, E., editors, *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2018, Oxford, UK, July 09–12, 2018*, pages 374–383. ACM.

- [Eberhart et al., 2017] Eberhart, C., Hirschowitz, T., and Seiller, T. (2017). An intensionally fully-abstract sheaf model for π (expanded version). *Log. Methods Comput. Sci.*, 13(4).
- [Ehrhard, 2012] Ehrhard, T. (2012). The scott model of linear logic is the extensional collapse of its relational model. *Theor. Comput. Sci.*, 424:20–45.
- [Ehrhard and Guerrieri, 2016] Ehrhard, T. and Guerrieri, G. (2016). The bang calculus: an untyped lambda-calculus generalizing call-by-name and call-by-value. In *PPDP*, pages 174–187. ACM.
- [Ehrhard et al., 2018] Ehrhard, T., Pagani, M., and Tasson, C. (2018). Full abstraction for probabilistic PCF. *J. ACM*, 65(4):23:1–23:44.
- [Ehrhard and Regnier, 2003] Ehrhard, T. and Regnier, L. (2003). The differential lambda-calculus. *Theor. Comput. Sci.*, 309(1-3):1–41.
- [Ehrhard and Regnier, 2008] Ehrhard, T. and Regnier, L. (2008). Uniformity and the taylor expansion of ordinary lambda-terms. *Theor. Comput. Sci.*, 403(2-3):347–372.
- [Faggian and Hyland, 2002] Faggian, C. and Hyland, M. (2002). Designs, disputes and strategies. In *CSL*, volume 2471 of *Lecture Notes in Computer Science*, pages 442–457. Springer.
- [Faggian and Maurel, 2005] Faggian, C. and Maurel, F. (2005). Ludics nets, a game model of concurrent interaction. In *LICS*, pages 376–385. IEEE Computer Society.
- [Faggian and Piccolo, 2009] Faggian, C. and Piccolo, M. (2009). Partial orders, event structures and linear strategies. In *TLCA*, volume 5608 of *Lecture Notes in Computer Science*, pages 95–111. Springer.
- [Fiore et al., 2008] Fiore, M., Gambino, N., Hyland, M., and Winskel, G. (2008). The cartesian closed bicategory of generalised species of structures. *Journal of the London Mathematical Society*, 77(1):203–220.
- [Gabbay and Ghica, 2012] Gabbay, M. and Ghica, D. R. (2012). Game semantics in the nominal model. In *MFPS*, volume 286 of *Electronic Notes in Theoretical Computer Science*, pages 173–189. Elsevier.
- [Ghica, 2005] Ghica, D. R. (2005). Slot games: a quantitative model of computation. In Palsberg, J. and Abadi, M., editors, *Proceedings of the 32nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2005, Long Beach, California, USA, January 12-14, 2005*, pages 85–97. ACM.
- [Ghica and McCusker, 2003] Ghica, D. R. and McCusker, G. (2003). The regular-language semantics of second-order idealized Algol. *Theor. Comput. Sci.*, 309(1-3):469–502.
- [Ghica and Murawski, 2004] Ghica, D. R. and Murawski, A. S. (2004). Angelic semantics of fine-grained concurrency. In *FoSSaCS*, volume 2987 of *Lecture Notes in Computer Science*, pages 211–225. Springer.

- [Ghica and Murawski, 2008] Ghica, D. R. and Murawski, A. S. (2008). Angelic semantics of fine-grained concurrency. *Ann. Pure Appl. Log.*, 151(2-3):89–114.
- [Ghica et al., 2006] Ghica, D. R., Murawski, A. S., and Ong, C. L. (2006). Syntactic control of concurrency. *Theor. Comput. Sci.*, 350(2-3):234–251.
- [Ghica and Tzevelekos, 2012] Ghica, D. R. and Tzevelekos, N. (2012). A system-level game semantics. In Berger, U. and Mislove, M. W., 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.
- [Girard, 1987] Girard, J. (1987). Linear logic. *Theor. Comput. Sci.*, 50:1–102.
- [Girard, 1988] Girard, J. (1988). Normal functors, power series and λ -calculus. *Ann. Pure Appl. Log.*, 37(2):129–177.
- [Girard, 2001] Girard, J. (2001). Locus solum: From the rules of logic to the logic of rules. *Math. Struct. Comput. Sci.*, 11(3):301–506.
- [Girard, 1989] Girard, J.-Y. (1989). Geometry of Interaction 1: Interpretation of System F. In *Studies in Logic and the Foundations of Mathematics*, volume 127, pages 221–260. Elsevier.
- [Goyet, 2013] Goyet, A. (2013). The lambda lambda-bar calculus: a dual calculus for unconstrained strategies. In *POPL*, pages 155–166. ACM.
- [Grellois and Melliès, 2015] Grellois, C. and Melliès, P. (2015). An infinitary model of linear logic. In Pitts, A. M., 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 41–55. Springer.
- [Griffin, 1990] Griffin, T. (1990). A formulae-as-types notion of control. In Allen, F. E., 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.
- [Harmer, 1999] Harmer, R. (1999). *Games and full abstraction for non-deterministic languages*. PhD thesis, Imperial College London, UK.
- [Harmer, 2004] Harmer, R. (2004). Innocent game semantics. *Lecture notes*, 2007.
- [Harmer et al., 2007] Harmer, R., Hyland, M., and Melliès, P. (2007). Categorical combinatorics for innocent strategies. In *LICS*, pages 379–388. IEEE Computer Society.
- [Harmer and McCusker, 1999] Harmer, R. and McCusker, G. (1999). 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.

- [Honda and Yoshida, 1999] Honda, K. and Yoshida, N. (1999). Game-theoretic analysis of call-by-value computation. *Theor. Comput. Sci.*, 221(1-2):393–456.
- [Hyland and Ong, 2000] Hyland, J. M. E. and Ong, C. L. (2000). On full abstraction for PCF: I, II, and III. *Inf. Comput.*, 163(2):285–408.
- [Hyland, 1997] Hyland, M. (1997). Game semantics. *Semantics and Logics of Computation*, 14:131.
- [Hyland and Schalk, 2002] Hyland, M. and Schalk, A. (2002). Games on graphs and sequentially realizable functionals. In *17th IEEE Symposium on Logic in Computer Science (LICS 2002), 22-25 July 2002, Copenhagen, Denmark, Proceedings*, pages 257–264. IEEE Computer Society.
- [Jaber, 2015] Jaber, G. (2015). Operational nominal game semantics. In Pitts, A. M., 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.
- [Jacq and Melliès, 2018] Jacq, C. and Melliès, P. (2018). Categorical combinatorics for non deterministic strategies on simple games. In Baier, C. and Lago, U. D., 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 39–70. Springer.
- [Joyal, 1977] Joyal, A. (1977). Remarques sur la théorie des jeux à deux personnes. *Gazette des sciences mathématiques du Québec*, 1(4):46–52.
- [Kelly and Laplaza, 1980] Kelly, G. M. and Laplaza, M. L. (1980). Coherence for compact closed categories. *Journal of pure and applied algebra*, 19:193–213.
- [Ker et al., 2002] Ker, A. D., Nickau, H., and Ong, C. L. (2002). Innocent game models of untyped lambda-calculus. *Theor. Comput. Sci.*, 272(1-2):247–292.
- [Kerinec et al., 2020] Kerinec, A., Manzonetto, G., and Pagani, M. (2020). Revisiting call-by-value böhm trees in light of their taylor expansion. *Log. Methods Comput. Sci.*, 16(3).
- [Koenig and Shao, 2020] Koenig, J. and Shao, Z. (2020). Refinement-based game semantics for certified abstraction layers. In *LICS*, pages 633–647. ACM.
- [Krebbers et al., 2014] Krebbers, R., Leroy, X., and Wiedijk, F. (2014). Formal C semantics: CompCert and the C standard. In Klein, G. and Gamboa, R., editors, *Interactive Theorem Proving - 5th International Conference, ITP 2014, Held as Part of the Vienna Summer of Logic, VSL 2014, Vienna, Austria, July 14-17, 2014. Proceedings*, volume 8558 of *Lecture Notes in Computer Science*, pages 543–548. Springer.

- [Laird, 1997] Laird, J. (1997). Full abstraction for functional languages with control. In *LICS*, pages 58–67. IEEE Computer Society.
- [Laird, 2001a] Laird, J. (2001a). A fully abstract game semantics of local exceptions. In *LICS*, pages 105–114. IEEE Computer Society.
- [Laird, 2001b] Laird, J. (2001b). A game semantics of idealized CSP. In *MFPS*, volume 45 of *Electronic Notes in Theoretical Computer Science*, pages 232–257. Elsevier.
- [Laird, 2002] Laird, J. (2002). A categorical semantics of higher order store. In *CTCS*, volume 69 of *Electronic Notes in Theoretical Computer Science*, pages 209–226. Elsevier.
- [Laird, 2007] Laird, J. (2007). A fully abstract trace semantics for general references. In *ICALP*, volume 4596 of *Lecture Notes in Computer Science*, pages 667–679. Springer.
- [Laird et al., 2013] Laird, J., Manzonetto, G., McCusker, G., and Pagani, M. (2013). Weighted relational models of typed lambda-calculi. In *LICS*, pages 301–310. IEEE Computer Society.
- [Lambek and Scott, 1988] Lambek, J. and Scott, P. J. (1988). *Introduction to higher-order categorical logic*, volume 7. Cambridge University Press.
- [Laurent, 2002] Laurent, O. (2002). Polarized games. In *LICS*, page 265. IEEE Computer Society.
- [Leroy, 2009] Leroy, X. (2009). Formal verification of a realistic compiler. *Commun. ACM*, 52(7):107–115.
- [Levy, 1999] Levy, P. B. (1999). Call-by-push-value: A subsuming paradigm. In *TLCA*, volume 1581 of *Lecture Notes in Computer Science*, pages 228–242. Springer.
- [Levy, 2014] Levy, P. B. (2014). Transition systems over games. Talk at the *Chocola* seminar.
- [Levy and Staton, 2014] Levy, P. B. and Staton, S. (2014). 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.
- [Loader, 2001] Loader, R. (2001). Finitary PCF is not decidable. *Theor. Comput. Sci.*, 266(1-2):341–364.
- [Mackie, 1995] Mackie, I. (1995). The geometry of interaction machine. In *POPL*, pages 198–208. ACM Press.
- [McCusker, 2003] McCusker, G. (2003). On the semantics of the bad-variable constructor in algol-like languages. In *MFPS*, volume 83 of *Electronic Notes in Theoretical Computer Science*, pages 169–186. Elsevier.

- [Melliès, 2004a] Melliès, P. (2004a). Asynchronous games 2: The true concurrency of innocence. In *CONCUR*, volume 3170 of *Lecture Notes in Computer Science*, pages 448–465. Springer.
- [Melliès, 2004b] Melliès, P. (2004b). Asynchronous games 3 an innocent model of linear logic. In *CTCS*, volume 122 of *Electronic Notes in Theoretical Computer Science*, pages 171–192. Elsevier.
- [Melliès, 2005] Melliès, P. (2005). Asynchronous games 4: A fully complete model of propositional linear logic. In *LICS*, pages 386–395. IEEE Computer Society.
- [Melliès, 2006] Melliès, P. (2006). Asynchronous games 2: The true concurrency of innocence. *Theor. Comput. Sci.*, 358(2-3):200–228.
- [Melliès, 2019a] Melliès, P. (2019a). Categorical combinatorics of scheduling and synchronization in game semantics. *Proc. ACM Program. Lang.*, 3(POPL):23:1–23:30.
- [Melliès, 2019b] Melliès, P. (2019b). Template games and differential linear logic. In *34th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2019, Vancouver, BC, Canada, June 24-27, 2019*, pages 1–13. IEEE.
- [Melliès, 2021] Melliès, P. (2021). Asynchronous template games and the gray tensor product of 2-categories. In *36th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2021, Rome, Italy, June 29 - July 2, 2021*, pages 1–13. IEEE.
- [Melliès and Mimram, 2007] Melliès, P. and Mimram, S. (2007). Asynchronous games: Innocence without alternation. In *CONCUR*, volume 4703 of *Lecture Notes in Computer Science*, pages 395–411. Springer.
- [Melliès and Tabareau, 2007] Melliès, P. and Tabareau, N. (2007). Resource modalities in game semantics. In *LICS*, pages 389–398. IEEE Computer Society.
- [Melliès, 2003] Melliès, P.-A. (2003). Asynchronous games 1: Uniformity by group invariance.
- [Melliès, 2009] Melliès, P.-A. (2009). Categorical semantics of linear logic. *Panoramas et synthèses*, 27:15–215.
- [Miller, 1987] Miller, D. A. (1987). A compact representation of proofs. *Stud Logica*, 46(4):347–370.
- [Mimram, 2008] Mimram, S. (2008). *Sémantique des jeux asynchrones et réécriture 2-dimensionnelle. (Asynchronous Game Semantics and 2-dimensional Rewriting Systems)*. PhD thesis, Paris Diderot University, France.
- [Moggi, 1991] Moggi, E. (1991). Notions of computation and monads. *Inf. Comput.*, 93(1):55–92.
- [Murawski, 2007] Murawski, A. S. (2007). Bad variables under control. In *CSL*, volume 4646 of *Lecture Notes in Computer Science*, pages 558–572. Springer.

- [Murawski, 2010] Murawski, A. S. (2010). Full abstraction without synchronization primitives. In *MFPS*, volume 265 of *Electronic Notes in Theoretical Computer Science*, pages 423–436. Elsevier.
- [Murawski and Tzevelekos, 2013] Murawski, A. S. and Tzevelekos, N. (2013). Deconstructing general references via game semantics. In *FoSSaCS*, volume 7794 of *Lecture Notes in Computer Science*, pages 241–256. Springer.
- [New et al., 2016] New, M. S., Bowman, W. J., and Ahmed, A. (2016). Fully abstract compilation via universal embedding. In Garrigue, J., Keller, G., and Sumii, E., editors, *Proceedings of the 21st ACM SIGPLAN International Conference on Functional Programming, ICFP 2016, Nara, Japan, September 18-22, 2016*, pages 103–116. ACM.
- [Nickau, 1994] Nickau, H. (1994). Hereditarily sequential functionals. In Nerode, A. and Matiyasevich, Y. V., editors, *Proceedings of the Symposium on Logical Foundations of Computer Science: Logic at St. Petersburg*, volume 813 of *Lecture Notes in Computer Science*, pages 253–264. Springer Verlag.
- [O’Hearn, 2015] O’Hearn, P. W. (2015). From categorical logic to facebook engineering. In *LICS*, pages 17–20. IEEE Computer Society.
- [Olimpieri, 2021] Olimpieri, F. (2021). Intersection type distributors. In *36th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2021, Rome, Italy, June 29 - July 2, 2021*, pages 1–15. IEEE.
- [Olimpieri and Vaux Auclair, 2022] Olimpieri, F. and Vaux Auclair, L. (2022). On the taylor expansion of λ -terms and the groupoid structure of their rigid approximants. *Log. Methods Comput. Sci.*, 18(1).
- [Pagani et al., 2014] Pagani, M., Selinger, P., and Valiron, B. (2014). Applying quantitative semantics to higher-order quantum computing. In *POPL*, pages 647–658. ACM.
- [Paquet, 2020] Paquet, H. (2020). Probabilistic concurrent game semantics. Technical report, University of Cambridge, Computer Laboratory.
- [Paquet, 2023] Paquet, H. (2023). Bi-invariance for uniform strategies on event structures. *Electronic Notes in Theoretical Informatics and Computer Science*, 1.
- [Plotkin, 1977] Plotkin, G. D. (1977). LCF considered as a programming language. *Theor. Comput. Sci.*, 5(3):223–255.
- [Plotkin, 1981] Plotkin, G. D. (1981). Post-graduate lecture notes in advanced domain theory (incorporating the \pisa notes"). *Dept. of Computer Science, Univ. of Edinburgh*.
- [Rideau and Winskel, 2011] Rideau, S. and Winskel, G. (2011). 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.

- [Sakayori and Tsukada, 2017] Sakayori, K. and Tsukada, T. (2017). A truly concurrent game model of the asynchronous π -calculus. In *FoSSaCS*, volume 10203 of *Lecture Notes in Computer Science*, pages 389–406.
- [Saunders-Evans and Winskel, 2006] Saunders-Evans, L. and Winskel, G. (2006). Event structure spans for nondeterministic dataflow. In *EXPRESS*, volume 175 of *Electronic Notes in Theoretical Computer Science*, pages 109–129. Elsevier.
- [Scott and Strachey, 1971] Scott, D. S. and Strachey, C. (1971). *Toward a mathematical semantics for computer languages*, volume 1. Oxford University Computing Laboratory, Programming Research Group Oxford.
- [Selinger, 1997] Selinger, P. (1997). First-order axioms for asynchrony. In Mazurkiewicz, A. W. and Winkowski, J., editors, *CONCUR '97: Concurrency Theory, 8th International Conference, Warsaw, Poland, July 1-4, 1997, Proceedings*, volume 1243 of *Lecture Notes in Computer Science*, pages 376–390. Springer.
- [Stewart et al., 2015] Stewart, G., Beringer, L., Cuellar, S., and Appel, A. W. (2015). Compositional compcert. In *POPL*, pages 275–287. ACM.
- [Tsukada and Ong, 2015] Tsukada, T. and Ong, C. L. (2015). Nondeterminism in game semantics via sheaves. In *30th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2015, Kyoto, Japan, July 6-10, 2015*, pages 220–231. IEEE Computer Society.
- [Tsukada and Ong, 2016] Tsukada, T. and Ong, C. L. (2016). Plays as resource terms via non-idempotent intersection types. In Grohe, M., Koskinen, E., and Shankar, N., 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 237–246. ACM.
- [Turi and Plotkin, 1997] Turi, D. and Plotkin, G. D. (1997). Towards a mathematical operational semantics. In *LICS*, pages 280–291. IEEE Computer Society.
- [Tzevelekos, 2009] Tzevelekos, N. (2009). Full abstraction for nominal general references. *Log. Methods Comput. Sci.*, 5(3).
- [Vale et al., 2023] Vale, A. O., Shao, Z., and Chen, Y. (2023). A compositional theory of linearizability. *Proc. ACM Program. Lang.*, 7(POPL):1089–1120.
- [Winskel, 1980] Winskel, G. (1980). *Events in computation*. PhD Thesis, Edinburgh University.
- [Winskel, 1982] Winskel, G. (1982). Event structure semantics for CCS and related languages. In Nielsen, M. and Schmidt, E. M., editors, *Automata, Languages and Programming, 9th Colloquium, Aarhus, Denmark, July 12-16, 1982, Proceedings*, volume 140 of *Lecture Notes in Computer Science*, pages 561–576. Springer.
- [Winskel, 1986] Winskel, G. (1986). Event structures. In Brauer, W., Reisig, W., and Rozenberg, G., editors, *Petri Nets: Central Models and Their Properties, Advances in Petri Nets 1986, Part II, Proceedings of an Advanced Course, Bad Honnef, Germany, 8-19 September 1986*, volume 255 of *Lecture Notes in Computer Science*, pages 325–392. Springer.

- [Winskel, 1993] Winskel, G. (1993). *The formal semantics of programming languages: an introduction*. MIT press.
- [Winskel, 2007] Winskel, G. (2007). Event structures with symmetry. *Electron. Notes Theor. Comput. Sci.*, 172:611–652.
- [Winskel, 2013a] Winskel, G. (2013a). Distributed probabilistic and quantum strategies. In *MFPS*, volume 298 of *Electronic Notes in Theoretical Computer Science*, pages 403–425. Elsevier.
- [Winskel, 2013b] Winskel, G. (2013b). Strategies as profunctors. In Pfenning, F., 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 418–433. Springer.
- [Xia et al., 2020] Xia, L., Zakowski, Y., He, P., Hur, C., Malecha, G., Pierce, B. C., and Zdancewic, S. (2020). Interaction trees: representing recursive and impure programs in coq. *Proc. ACM Program. Lang.*, 4(POPL):51:1–51:32.

Part V

Appendices

Appendix A

Complements on Relative Seely Categories

Like many other works in denotational semantics, the developments appearing in this monograph target programming languages with a non-linear type system, whose models are certain cartesian closed categories. But as is also common in denotational semantics, our cartesian closed categories are constructed following a linear decomposition, from a model of intuitionistic linear logic. A standard categorical structure used for this purpose is that of a *Seely category* [Melliès, 2009]: it is a symmetric monoidal closed category C also equipped with a cartesian structure along with a comonad $!$ on C and

$$!A \otimes !B \cong !(A \& B), \quad 1 \cong !\top,$$

natural isos called the *Seely isomorphisms* satisfying a few coherence conditions.

It is possible to build a genuine Seely category of games and concurrent strategies, and this is done for instance in [Castellan and Clairambault, 2021]. But the construction is inelegant: as argued in Section 8.2, handling types such as $o \multimap (o \otimes o)$ with a tensor on the right hand side complicates the shape of games and weakens the connection with the relational model. So we shall ban such types and work not with Seely categories, but with *relative* Seely categories, defined and developed here¹.

In Section A.1, we first define relative Seely category, then construct the cartesian closed Kleisli category of a Seely category. Then, in Section A.2 we shall define functors between relative Seely categories, and show how they lift to the Kleisli category.

¹More precisely, the monograph builds on *relative Seely \sim -category*, where every law holds up to an equivalence relation \sim , which is also preserved by all constructions on morphisms – but we omit it in this appendix as it is completely orthogonal to the development.

A.1 Definition and Kleisli Category

Definition. It seems natural to first restate the definition from Section 8.2.1:

Definition 8.2.1. A *relative Seely category* is a symmetric monoidal category $(C, \otimes, 1)$ equipped with a full subcategory C_s together with the following data and axioms:

- C_s has finite products $(\&, \top)$ preserved by the inclusion functor $J : C_s \hookrightarrow C$.
- For every $B \in C$ there is a functor $B \multimap - : C_s \rightarrow C_s$, such that there is

$$\Lambda(-) : C(A \otimes B, S) \simeq C(A, B \multimap S).$$

a bijection natural in $A \in C$ and $S \in C_s$.

- There is a J -relative comonad $! : C_s \rightarrow C$: we have, for every $S \in C_s$, an object $!S \in C$ and a *dereliction* morphism $\text{der}_S : !S \rightarrow S$, and for every $\sigma : !S \rightarrow T$, a *promotion* $\sigma^! : !S \rightarrow !T$, subject to three axioms [Altenkirch et al., 2010]:

$$\begin{aligned} \text{der}_T \circ \sigma^! &= \sigma & (\sigma : !S \rightarrow T) \\ \text{der}_S^! &= \text{id}_{!S} & (S \in C_s) \\ (\tau \circ \sigma^!)^! &= \tau^! \circ \sigma^! & (\sigma : !S \rightarrow T, \tau : T \rightarrow U), \end{aligned}$$

which make $! : C_s \rightarrow C$ a functor, via $!\sigma = (\sigma \circ \text{der}_S)^!$ for $\sigma : S \rightarrow T$.

- The functor $! : C_s \rightarrow C$ is symmetric strong monoidal $(C_s, \&, \top) \rightarrow (C, \otimes, 1)$, so

$$m_0 : 1 \rightarrow !\top \quad m_{S,T} : !S \otimes !T \rightarrow !(S \& T)$$

are natural isos, additional compatible with promotion: the diagram

$$\begin{array}{ccc} !\Gamma & \xrightarrow{\langle f, g \rangle^!} & !(S \& T) \\ \langle \text{der}, \text{der} \rangle^! \downarrow & & \downarrow m^{-1} \\ !(\Gamma \& \Gamma) & & !S \otimes !T \\ & \searrow m^{-1} \quad \nearrow f^! \otimes g^! & \\ & !\Gamma \otimes !\Gamma & \end{array}$$

commutes for all $\Gamma, S, T \in C_s$, $f \in C(!\Gamma, S)$, $g \in C(!\Gamma, T)$.

As we shall annotate with $!$ the components of the constructed Kleisli category, we prefer to write the promotion as σ^\dagger from now on (rather than $\sigma^!$), to avoid the collision.

Kleisli category. We define the Kleisli category $C_!$ as follows. Its *objects* are simply those of C_s . A *morphism* from S to T is $f : !S \rightarrow T$ in C . The *identity* on object S is $\text{id}_S^\dagger = \text{der}_S : !S \rightarrow S$. The *composition* of $\sigma \in C_!(S, T)$ and $\tau \in C_!(T, U)$ is

$$\tau \circ^\dagger \sigma = \tau \circ \sigma^\dagger.$$

It is completely straightforward that this data defines a category; the equations to be verified directly boil down to the axioms of relative comonads above. Thus we have:

Proposition A.1.1. The data above forms a category $C_!$, the *Kleisli category* of $!$.

Cartesian closed structure. Obviously, \top is still terminal in $C_!$. If $S, T \in C_s$, their product is $S \& T \in C_s$. The corresponding projections are (in C):

$$\begin{aligned} \pi_S^! &= \pi_S \circ \text{der}_{S \& T} & : & \quad !(S \& T) \rightarrow S \\ \pi_T^! &= \pi_T \circ \text{der}_{S \& T} & : & \quad !(S \& T) \rightarrow T, \end{aligned}$$

and the pairing of $\sigma : !\Gamma \rightarrow S$ and $\tau : !\Gamma \rightarrow T$ is $\langle \sigma, \tau \rangle_! = \langle \sigma, \tau \rangle$. It is easy to check:

Proposition A.1.2. *This data equips $C_!$ with a cartesian structure.*

Likewise, given $S, T \in C_s$, their arrow is $S \Rightarrow T = !S \multimap T$. The evaluation is

$$\text{ev}_{S,T}^{\Rightarrow} = !((!S \multimap T) \& S) \xrightarrow{m_{!S \multimap T, S}^{-1}} !(S \multimap T) \otimes !S \xrightarrow{\text{der}_{!S \multimap T} \otimes !S} (!S \multimap T) \otimes !S \xrightarrow{\text{ev}_{!S, T}} T$$

and given $\sigma : !(S \& T) \rightarrow U$, its currying is $\Lambda_{S,T,U}^! = \Lambda_{!S, !T, U}(\sigma \circ m_{S,T})$ which is in $C_!(S, !T \multimap U)$ as required. Again, it is an elementary verification that:

Theorem A.1.3. *This data equips $C_!$ with a cartesian closed structure.*

A.2 Relative Seely Functors

Theorem A.1.3 is our main tool to construct cartesian closed categories in this monograph; but we also need an accompanying tool to construct cartesian closed functors from an appropriate notion of functor between relative Seely categories.

Definition. Let us start by defining relative Seely functors.

Definition A.2.1. *Let C, D be relative Seely categories. A **relative Seely functor** $C \rightarrow D$ is a functor $F : C \rightarrow D$ which restricts to $F : C_s \rightarrow D_s$, equipped with:*

- For every $A, B \in C$, natural isomorphisms morphisms

$$\begin{aligned} t_{A,B}^{\otimes} &: FA \otimes FB \cong F(A \otimes B) \\ t^1 &: 1 \cong F1; \end{aligned}$$

making (F, t^{\otimes}, t^1) a strong symmetric monoidal functor $(C, \otimes, 1) \rightarrow (D, \otimes, 1)$;

- for every $S, T \in C_s$, isomorphisms

$$\begin{aligned} t_{S,T}^{\&} &: FS \& FT \cong F(S \& T) \\ t^{\top} &: \top \cong F\top; \end{aligned}$$

- For every $A \in C$ and $S \in C_s$, an isomorphism

$$t_{A,S}^{\multimap} : FA \multimap FS \cong F(A \multimap S);$$

- For every $S \in C_s$, an isomorphism

$$t_S^! : !FS \cong F!S;$$

satisfying the coherence axioms of Figure A.1, and the naturality-like condition that for every $S, T \in \mathcal{C}_s$ and $\sigma : !S \rightarrow T$, the following diagram commutes:

$$\begin{array}{ccc} !FS & \xrightarrow{(F\sigma \circ t_S^!)^\dagger} & !FT \\ t_S^! \downarrow & & \downarrow t_T^! \\ F!S & \xrightarrow{F(\sigma^\dagger)} & F!T \end{array}$$

Lifting to the Kleisli categories. Given a relative Seely functor $F : \mathcal{C} \rightarrow \mathcal{D}$, we define its **lifting** to a functor $F_! : \mathcal{C}_! \rightarrow \mathcal{D}_!$ that we shall prove to be cartesian closed.

On objects, we simply set $F_!(S) = FS$. On morphisms, for $\sigma : !S \rightarrow T$, we set

$$F_!\sigma = F\sigma \circ t_S^! : !FS \rightarrow FT.$$

From this data, a routine calculation ensures that we have:

Proposition A.2.2. *The above data yields a functor $F_! : \mathcal{C}_! \rightarrow \mathcal{D}_!$.*

Next, we show that this functor is cartesian. By definition, this means that FT is terminal – which is clear from t^\top – and that for all $S, T \in \mathcal{C}_!$, the canonical map

$$\langle F_!\pi_S^!, F_!\pi_T^! \rangle_! \in \mathcal{D}_!(F(S \& T), FS \& FT)$$

obtained via the cartesian structures of $\mathcal{C}_!$ and $\mathcal{D}_!$, is invertible. Now using our coherence laws, this morphism simplifies to $\langle F\pi_S \circ \text{der}_{S,T}, F\pi_T \circ \text{der}_{S,T} \rangle : !F(S \& T) \rightarrow FS \& FT$. And indeed, the inverse we seek can be constructed as

$$r_{S,T}^\& = t_{S,T}^\& \circ \text{der}_{FS \& FT} : !(FS \& FT) \rightarrow F(S \& T);$$

that it is an inverse is a direct verification via our coherence laws, which ensure that $t_{S,T}^\&$ is an inverse in \mathcal{D} to the canonical map $\langle F(\pi_S), F(\pi_T) \rangle : F(S \& T) \rightarrow FS \& FT$ given by the cartesian closed structures in \mathcal{C} and \mathcal{D} .

Likewise, proving that $F_!$ is cartesian closed means showing, for all $S, T \in \mathcal{C}_!$, that

$$\Lambda^!(F_!(\text{ev}_{S,T}^\Rightarrow) \circ t_{S,T}^\Rightarrow) \in \mathcal{C}_!(F(S \Rightarrow T), FS \Rightarrow FT)$$

obtained via the cartesian closed structures of $\mathcal{C}_!$ and $\mathcal{D}_!$, has an inverse. But via our coherence rules, this morphism simplifies to the composition in \mathcal{D} :

$$!F(!S \multimap T) \xrightarrow{\text{der}_{F(!S \multimap T)}} F(!S \multimap T) \xrightarrow{(t_{S,T}^\multimap)^{-1}} F(!S) \multimap FT \xrightarrow{t_S^! \multimap FT} !FS \multimap FT,$$

where we use that as usual, relative closure informs a contravariant functor – $\multimap S : \mathcal{D}^{\text{op}} \rightarrow \mathcal{D}$ for all $S \in \mathcal{D}_s$, defined as $\sigma \multimap S = \Lambda(\text{ev}_{B,S} \circ ((B \multimap S) \otimes \sigma))$ for any $\sigma : A \rightarrow B$ in \mathcal{D} . Now, the morphism above has an inverse in $\mathcal{D}_!$, constructed as

$$r_{S,T}^\Rightarrow = t_{S,T}^\multimap \circ ((t_S^!)^{-1} \multimap FT) \circ \text{der}_{FS \Rightarrow FT} : !(FS \multimap FT) \rightarrow F(!S \multimap T)$$

concluding the proof of the following statement:




Figure A.1: Coherence diagrams for relative Seely functors

Theorem A.2.3. *Consider $F : \mathcal{C} \rightarrow \mathcal{D}$ a relative Seely functor.*

Then, the functor $F_! : \mathcal{C}_! \rightarrow \mathcal{D}_!$ constructed above is cartesian closed.

All the axioms of relative Seely functors get used in the proof, except for the last coherence diagram of Figure A.1, which appears unnecessary to ensure that a relative Seely functor lifts to a cartesian closed functor between the Kleisli categories. We include it nonetheless, as it seems likely to be necessary if we ever wish to show that relative Seely functors preserve the interpretation of a linear/non-linear type system.