

Number 950



**UNIVERSITY OF
CAMBRIDGE**

Computer Laboratory

Probabilistic concurrent game semantics

Hugo Paquet

August 2020

15 JJ Thomson Avenue
Cambridge CB3 0FD
United Kingdom
phone +44 1223 763500
<https://www.cl.cam.ac.uk/>

© 2020 Hugo Paquet

This technical report is based on a dissertation submitted September 2019 by the author for the degree of Doctor of Philosophy to the University of Cambridge, Homerton College.

Technical reports published by the University of Cambridge Computer Laboratory are freely available via the Internet:

<https://www.cl.cam.ac.uk/techreports/>

ISSN 1476-2986

Abstract

This thesis presents a variety of models for probabilistic programming languages in the framework of concurrent games.

Our starting point is the model of concurrent games with symmetry of Castellan, Clairambault and Winskel. We show that they form a symmetric monoidal closed bicategory, and that this can be turned into a cartesian closed bicategory using a linear exponential pseudo-comonad inspired by linear logic.

Then, we enrich this with probability, relying heavily on Winskel's model of probabilistic concurrent strategies. We see that the bicategorical structure is not perturbed by the addition of probability. We apply this model to two probabilistic languages: a probabilistic untyped λ -calculus, and Probabilistic PCF. For the former, we relate the semantics to the probabilistic Nakajima trees of Leventis, thus obtaining a characterisation of observational equivalence for programs in terms of strategies. For the latter, we show a definability result in the spirit of the game semantics tradition. This solves an open problem, as it is notoriously difficult to model Probabilistic PCF with sequential game semantics.

Finally, we introduce a model for measurable game semantics, in which games and strategies come equipped with measure-theoretic structure allowing for an accurate description of computation with continuous data types. The objective of this model is to support computation with arbitrary probability measures on the reals. In the last part of this thesis we see how this can be done by equipping strategies with parametrised families of probability measures (also known as stochastic kernels), and we construct a bicategory of measurable concurrent games and probabilistic measurable strategies.

Acknowledgements

I would first like to thank my supervisor Glynn Winskel for his support throughout this PhD. Glynn was always available for advice, and he constantly showed patience and understanding as I worked on problems I found interesting. Our relationship has become an open and friendly one, for which I am very grateful.

This thesis owes a great deal to Pierre Clairambault: he taught me all about game semantics, and our subsequent collaboration has had a strong influence on me; I hope it will continue. It has been really enjoyable spending time together in Cambridge and in Lyon. I have also benefited greatly from interacting with Simon Castellan, who has been an excellent collaborator and friend over the years. This thesis is a direct continuation of work by Glynn, Pierre and Simon, which would not have been possible without their contribution and guidance.

There are many other people I want to acknowledge. I am very grateful to Martin Hyland and Guy McCusker for examining this thesis. I enjoyed the many discussions I have had with Marcelo Fiore; he has taught me a lot. I am grateful to Christine Tasson for introducing me to semantics, and for remaining a source of encouragement over the years. I have had many inspiring conversations with Tamara von Glehn. I would like to thank Jim Laird for our discussions, and for an enjoyable visit to Bath. I would also like to thank Ohad Kammar, who has been influential in my understanding of probabilistic programming. I leave out many others from the Computer Lab, and from ENS Lyon, whom I have enjoyed spending time with.

Lise Gough has been an incredible source of help with funding and travel. I am also grateful for financial support from Luke Ong for the final months of writing up.

It has always been fun socialising and discussing with other PhD students. I am lucky to have been part of a wonderful cohort of Theory students in Cambridge: Matt Daggitt, Dylan McDermott, Philip Saville, Ian Orton. We have enjoyed countless scientific discussions, lunches and pub nights. I want to thank Philip in particular: he has been the ideal office mate, and without his help I probably would not have embarked on the bicategorical development in this thesis. Thanks also to Aurore Alcolei and Marc de Visme.

I am grateful for my friends in Cambridge, in particular Laura, Ruth, Sam, Max, Bill, and others from Syzygie and the Homerton MCR. I want to thank Aaron, whose presence made the final few weeks of writing up manageable. Finally I would like to thank Manon and my parents for their invaluable support.

Contents

1	Introduction	11
1.1	Background	12
1.1.1	Game semantics	12
1.1.2	Concurrent games based on event structures	13
1.1.3	Probabilistic programs and their semantics	14
1.2	This thesis	15
2	Concurrent games	17
2.1	Basic games and strategies as event structures	17
2.1.1	Partial-order models of processes	17
2.1.2	Games and strategies	19
2.2	Symmetry in concurrent games	21
2.2.1	Symmetry in event structures	21
2.2.2	Arenas	22
2.2.3	Thin strategies	24
2.3	Composing strategies	24
2.3.1	Interaction	24
2.3.2	Strategies as morphisms	25
2.3.2.1	Interaction	26
2.3.2.2	Hiding and composition	26
2.3.3	Causally compatible pairs	27
2.4	Copycat and associativity of composition	28
2.4.1	Identity strategies	28
2.4.2	Copycat is an identity	28
2.4.3	Associativity of composition	29
2.5	A pseudo-double category	29
2.5.1	The pseudo-double category \mathcal{G}	35
2.5.2	Lifting	39
2.5.3	Symmetric monoidal structure	40
2.6	A cartesian closed bicategory	42
2.6.1	Closed structure	43
2.6.2	Finite products	47
2.6.3	A linear exponential pseudo-comonad	48
2.6.4	The cartesian closed Kleisli bicategory	51
2.6.5	Colimits of strategies	54

3	Probability in concurrent games	57
3.1	Probabilistic strategies	57
3.2	Markov strategies	62
3.3	Determinism and race-freeness	64
3.3.1	Deterministic strategies	64
3.3.2	Race-free games	65
3.4	The bicategory PG	65
3.4.1	Composition of probabilistic strategies	65
3.4.2	2-Cells	67
3.5	Rigid maps and push-forward valuations	67
3.6	Bicategorical structure	69
4	Innocence and the untyped λ-calculus	71
4.1	Syntax and semantics of the probabilistic λ -calculus	72
4.1.1	Syntax and operational semantics	72
4.1.2	A reflexive object in PG _!	73
4.1.3	Interpretation of λ^+	76
4.2	Probabilistic innocence	77
4.2.1	Justifiers and visibility	77
4.2.2	Sequential innocence	78
4.3	Adequacy	79
4.3.1	Convergence	79
4.3.2	The relation \triangleleft	80
4.3.3	Closure properties and proof of adequacy	82
4.4	Strategies and Nakajima trees	84
4.4.1	Probabilistic Nakajima trees	84
4.4.2	Nakajima-like strategies	86
4.4.3	Nakajima trees correspond to Nakajima-like strategies	89
4.5	The correspondence theorem	90
4.5.1	Leventis' method for approximating trees	90
4.5.2	Local semantics	94
4.5.3	The strategy $\llbracket M \rrbracket_{\text{nak}}^\Gamma$ as a lub	96
4.5.4	Wrapping up and full abstraction	98
5	Probabilistic PCF	99
5.1	Syntax and operational semantics	99
5.2	Arenas with questions and answers	100
5.3	Semantics of PPCF	102
5.4	Definability	103
6	Measurable concurrent games	107
6.1	Measurable event structures	108
6.1.1	Fibrations of event structures	108
6.1.2	Event structures and presheaves	109
6.1.3	Symmetry in measurable event structures	110
6.1.4	Embedding Ess into Mess	113
6.2	Measurable games and strategies	113

6.2.1	Pullbacks in Mess	114
6.2.2	Composition of measurable strategies	116
6.3	The pseudo-double category \mathcal{MG}	118
6.3.1	Measurable copycat	119
6.3.2	2-Cells	120
6.3.3	Structural isomorphisms	123
6.4	Categorical properties	126
6.4.1	Symmetric monoidal closed structure	126
6.4.2	Products	127
6.4.3	A pseudo-comonad	127
7	Probability in measurable games	129
7.1	Probability theory	129
7.1.1	Measures and kernels	129
7.1.2	Products	130
7.1.3	Standard Borel spaces	130
7.2	Probability in measurable event structures with symmetry	131
7.2.1	Conditional valuations	131
7.3	Probability in measurable strategies	133
7.3.1	Measurable race-freeness	133
7.3.2	Conditional valuations on measurable essps	133
7.3.3	Interaction of probabilistic strategies	134
7.3.4	The drop condition in the interaction	138
7.3.5	Composition of probabilistic strategies	141
7.3.6	Probabilistic copycat	142
7.4	The bicategory of probabilistic measurable strategies	143
7.4.1	The pseudo-double category \mathcal{PMG}	143
7.4.2	Monoidal structure	144
7.4.3	The cartesian closed bicategory	146
7.4.4	The discrete sub-bicategory	146
8	Conclusion	147

Chapter 1

Introduction

It is exactly fifty years ago that Dana Scott [Sco69] introduced domain theory. In [Sco70] he put forward his ideas for a “mathematical theory of computation”, laying down the foundations for the field of denotational semantics. Scott advocates the use of continuous functions between certain partially ordered sets as an interpretation for programs, following some principles: the mathematical representation should be independent of any implementation details, but it must be *sound* (meaning that the abstract representation should give an accurate account of program behaviour) and designed *compositionally*: the semantics of a program must be obtained from the semantics of its parts.

Since then, denotational semantics has evolved far beyond Scott’s domain-theoretic model. A variety of models for programs is now available, reflecting the need to support a wide range of programming features and study many aspects of program behaviour.

To relate and unify these models, and to build more, the general framework of *category theory* offers much help: in a category, the notion of composition is central, so that in a majority of cases, presenting a semantic model as a category provides it with a sound structural basis. In this situation, the interpretation of common program constructs (function application, recursion, conditionals, *etc.*) is done by applying general categorical principles guaranteeing correctness.

This thesis fits in a line of research in denotational semantics concerned with the interpretation of programs as strategies in certain two-player games: this is known as *game semantics*. More specifically, our results are about a theory of *concurrent games* introduced in [RW11] and extensively developed in the past decade, most notably by Glynn Winskel, Pierre Clairambault and Simon Castellan.

Our investigation into concurrent games is in the context of *probabilistic computation*. With an eye towards building models for probabilistic programs, we build on the theory of probabilistic strategies put forward by Winskel [Win13a] and propose some extensions; our contributions include several generalisations of existing concurrent games models. We carry out some applications to programming language semantics, but in this thesis the main point of concern is the mathematical development of the models, and ensuring they enjoy the necessary categorical properties for future application to the semantics of probabilistic programs. Following a growing trend

Opponent: “What is the value of this program?”
 Player: “What is the value of its argument f ?”
 Opponent: “ f is a function. What is its argument?”
 Player: “3.”
 Opponent: “Then the value of f is n .”
 Player: “Then the value of the program is n .”

The *strategy* representing the term contains all possible plays in an execution, accounting for all possible execution contexts.

One key aspect should be emphasised: the higher-order term is reduced to an exchange of first-order messages, and at no point is it necessary to build a function space. This is due to the *intensionality* of the game model. Models in domain theory are called *extensional*. (Our goal is not to draw a full comparison here, simply to highlight this as an advantage of game models over domain theory ones; though of course this is mitigated by the very technical nature of games.)

After the initial success obtained with PCF, the field of game semantics was pushed much further. Games proved to be widely applicable, and we only cite a few of the many computational features accounted for: first-order references [AM96], general references [AHM98], control operators [Lai97], call-by-value computation [HY97, AM97], nondeterminism [HM99], and probability [DH02].

1.1.2 Concurrent games based on event structures

Rideau and Winskel’s original *concurrent games* model appears in [RW11]. In this work they show how *event structures* [NPW79], a fundamental model for concurrent processes, can be used as the basis for a general games model in which *plays* are no longer the central object. Their viewpoint is not the same as that of [HO00, AJM00] and their successors, and the new model is different in many ways:

- It is *linear*: no move can be played more than once;
- It is *causal*: event structures make explicit the causal dependency relation between moves;
- The resulting mathematical framework is a *bicategory*, rather than a category; this is a technical point, but arguably, the mathematical universe is richer and has finer structure.
- Explicit branching point: in a concurrent strategy, every nondeterministic choice made by a Player is represented explicitly. This is not the case in traditional game semantics, since in any given play only the chosen path is recorded.

This last point is important to us, because, as pointed out by Harmer [Har99], the absence of explicit branching information makes the modelling of nondeterministic and probabilistic computation less modular. (Another solution to this problem is to use presheaves over plays, see [TO15, TO14].)

The linearity requirement, on the other hand, was an obstacle to the development of game semantics models for languages of the kind discussed above. So in order to fully exploit the advantages of concurrent games, Castellan, Clairambault and Winskel introduced *games with symmetry* [CCW14], based on event structures with symmetry [Win07]. This significant generalisation allowed the same authors to carve out a model of *concurrent Hyland-Ong* games in [CCW15], in which, for instance, the original definability argument for PCF can be re-enacted. But the games of [CCW15] are much more general, and thus concurrent games already have far-reaching applications. Examples include the theory of programs with parallel features [CC16, CCW17], non-angelic nondeterminism [CCHW18] and quantum programming [CdVW19].

Independently of the above, and soon after the seminal paper [RW11], Winskel developed a theory of probabilistic strategies [Win13a] on concurrent games. By exploiting a connection between event structures and domain theory, he defines a notion of *probabilistic event structures* generalising previous attempts at probabilistic models for concurrency [VW04, AB06]. The result is a model of concurrent probabilistic strategies exhibiting the four characteristics listed above.

This thesis sets out to bring together this probabilistic model and the work on games with symmetry of Castellan et al., with the objective of modelling general programming languages with probabilistic features. Furthermore, we are concerned with the support of *continuous* probability distributions, which are not readily supported by the model of [Win13a].

These objectives are not new; there is a large body of research on the denotational semantics of probabilistic programs. In what follows we mention a few approaches.

1.1.3 Probabilistic programs and their semantics

A basic example of a probabilistic program is given by a single call to a random primitive **coin** which, say, returns one of the boolean values **tt** or **ff**, each with probability $\frac{1}{2}$. One can then build more sophisticated programs: **if coin then 2 else 3**, $\lambda f.f$ **coin**, *etc.*; soon it becomes tricky to correctly reason about their operational behaviour.

The earliest examples of probabilistic programming languages are found in work by Saheb-Djaromi [SD78, Sah80] and Kozen [Koz79]. Jones and Plotkin [JP89, Jon90] studied semantics for a probabilistic language in the setting of domain theory, via a monad called the *probabilistic powerdomain*. The approach is successful but not entirely satisfying – see [JT98] for a detailed account. Recently other domain-theoretic models have been defined and extensively studied by Goubault-Larrecq [Gou15, Gou19]

The above models are of the *extensional* kind: a program is interpreted as a function from an input space to the space of probability distributions on a output space, or a variation thereof. There are also a variety of *intensional* approaches, among them Danos and Harmer’s *probabilistic game semantics* [DH02], an obvious precursor to the present work. This is a probabilistic version of Hyland and Ong’s game model, which is fully abstract (in the sense of Plotkin [Plo77], *i.e.* without the

need for a quotient) for an extension of PCF with probabilistic choice and first-order references called *Probabilistic Algol*.

Closely related, also, is the line of research stemming from the relational and coherent semantics of linear logic [Gir88, Gir87] often referred to as *quantitative semantics*. Some pioneering work is by Lamarche [Lam92], a general approach is developed in [LMMP13], and there is a series of papers focusing on probabilistic computation, including [DE11, EPT11, ETP14]. Note that [ETP14] contains a striking full abstraction result for a probabilistic extension of PCF *without* references.

Recent developments and *probabilistic programming*. The past few years have seen a rise in the development of models for probabilistic programs. This is largely due to the development of a set of tools and methods for statistical modelling referred to as *probabilistic programming* [GHNR14]. The idea is for a statistical model to be encoded as a probabilistic program, whose variables correspond to the random variables in the model. The advantages of this approach for practitioners are, on one hand, access to a wide range of features (data structures, conditionals, recursion, higher-order functions, etc.) making for a succinct encoding of sophisticated models, and on the other hand, access to built-in inference engines, giving approximate solutions to inference problems at runtime. Implementing such systems efficiently is difficult, however. The expressive power of programming languages, albeit elegant, makes well-known inference methods insufficient in general. Thus, extending those methods to arbitrary probabilistic programs is the subject of active research, which triggered the need for more theoretical guarantees. This has been very stimulating for the semantics community: see for instance [BDLGS16, SYW⁺16, Sta17, HKS17, VKS19].

We now detail the contributions of each chapter.

1.2 This thesis

In Chapter 2, we give an introduction to concurrent games, and to the games with symmetry as first presented in [CCW15]. Our presentation is inspired by that in the journal versions [CCRW17, CCW19]. Most of the material in this chapter is obtained directly from the above papers, though in Sections 2.5 and 2.6 we offer a much more thorough description of the bicategorical structure of the model than was previously available.

In Chapter 3, we introduce Winskel’s probabilistic strategies and enrich them with symmetry. The extension is not particularly problematic and the results of [Win13b], found in full details in a set of unpublished notes [Win], provide most of the tools. We introduce Markov strategies, which are needed in the next chapter, show a “push-forward” property for maps of probabilistic strategies, and discuss some categorical properties.

In Chapter 4, we describe an application to the semantics of a probabilistic extension of the untyped λ -calculus. We define *sequential innocent* probabilistic strategies, and prove they are stable under composition. We show an adequacy result

for the model, and compare the semantics of the language to a notion of probabilistic *Nakajima trees* introduced by Leventis [Lev16]. This work was done in collaboration with Pierre Clairambault, see the conference paper [CP18].

In Chapter 5 we discuss another application, this time to a probabilistic extension of PCF. We exploit the innocence condition of the previous chapter to prove a definability result. The results of this chapter, together with a study of the connections between probabilistic concurrent games and the probabilistic relational model of [LMMP13], are joint work with Simon Castellan, Pierre Clairambault and Glynn Winskel, see [CCPW18].

In Chapter 6, we discuss a generalisation of the bicategory of concurrent games with symmetry of Chapter 2. In this generalised model, games and strategies come equipped with *measure-theoretic* structure, allowing for an accurate description of computation with continuous data types, *e.g.* real numbers. The objective of this model is to support computation with arbitrary probability measures on the reals. This is a relatively technical mathematical development: we give in full details the construction of a bicategory of *measurable games and strategies*.

In Chapter 7, we enrich measurable games with probability using (basic) tools from measure theory. The result is a generalisation of the model of Chapter 3, which can in principle be used to model, say, a version of PCF with a real number type and continuous distributions. (We leave this important application as further work.)

The content of Chapters 6 and 7 originates in joint work with Glynn Winskel, presented in [PW18]. But this work was done in games *without* symmetry; the present development is so far unpublished.

We conclude in Chapter 8.

Chapter 2

Concurrent games

The starting point of this thesis is a framework for games and strategies known as *concurrent games*. Originally appearing in work by Rideau and Winskel [RW11], theory and applications of concurrent games have been substantially developed by Castellan, Clairambault and Winskel (see for example [CCW14, CCW15]). The extended versions [CCRW17, CCW19] have largely inspired the presentation given here.

The purpose of this chapter is to give a formal account of the construction of the model. This will be useful in later chapters where this construction is extended with probability.

The chapter is organised as follows. In Section 2.1 we introduce event structures and the basic notion of games and strategies based on them. In Section 2.2 we see that for our purposes additional structure is needed, *symmetry* in event structures. In Section 2.3 we see how to compose strategies, and in Section 2.4 we discuss the associativity of composition and identity strategies. In the remaining two sections, 2.5 and 2.6, we discuss some categorical properties of the model.


2.1 Basic games and strategies as event structures

We start with a brief introduction to *event structures*, which play a fundamental role in the development of the model; we will see that both games and strategies are instances of event structures.

An event structure consists in a partially ordered set of events equipped with extra structure indicating which of the events are compatible. We first discuss how this can be used to model nondeterministic processes.


2.1.1 Partial-order models of processes

Consider the set containing three events a , b , and c under the partial order \leq generated by $a \leq c$ and $b \leq c$. This may be represented as follows:



As the graphical representation suggests, the relation \leq should be interpreted as a dependency relation: the event c can only occur once both events a and b have occurred. In this example a and b are said to be *concurrent*, meaning that they are causally independent and in an execution of the process one may find both events occurring.

Now take the same set of events but with the order \leq instead defined by $a \leq c$ only. Suppose that a and b are known to be incompatible events: this is pictured as



with the wiggly line expressing this incompatibility. These two relations on events (dependency and incompatibility) are combined in an elegant way in *event structures*, originally introduced in [NPW79]. In order to better convey the idea, we first define a class of event structures whose definition is more intuitive. This is the class of event structures *with binary conflict*, which contains most examples of event structures found in this thesis:

Definition 2.1. An **event structure with binary conflict** is a tuple $(E, \leq, \#)$ consisting of a set of event E , a partial order \leq on E , and an irreflexive *conflict* relation $\#$ on E satisfying the following two axioms:

- for every $e \in E$, the set $[e] = \{e' \in E \mid e' \leq e\}$ is finite; and
- if $e \# e'$ and $e' \leq e''$ then $e \# e''$.

The set $[e]$ is the **causal history** of event e ; we also write $[e]$ for $[e] \setminus \{e\}$. The first condition in the definition is the requirement that every event should be accessible in finite time. The second says that conflict is hereditary: in the previous example, b and c are incompatible, though this was kept implicit in the diagram. Indeed we only draw $e \rightsquigarrow e'$ when e and e' are in **minimal conflict**: $e \# e'$ and for any $e_0 \leq e$ and $e_1 \leq e'$, either $e_0 = e$ and $e_1 = e'$ or $\neg(e_0 \# e_1)$. As seen in the above figures, in order to depict the dependency relation \leq it is sufficient to draw **immediate causality** $e \rightarrow e'$, defined as $e < e'$ with nothing in between. (The order \leq is then recovered as the reflexive, transitive closure.)

Event structures with binary conflict can be generalised to account for situations in which, say, three events are pairwise compatible but mutually incompatible. Instead of a conflict relation we take the dual: a *consistency* relation indicating which subsets of events may occur together.

This extra generality comes at the cost of a slightly more involved definition, but makes some of the abstract theory more natural.

Definition 2.2. An **event structure** is a tuple (E, \leq, Con) consisting of a set E , a partial order \leq on E , and a nonempty set Con of subsets of E , satisfying the following axioms:

- (*Finite causes*) for every $e \in E$, the set $[e] = \{e' \in E \mid e' \leq e\}$ is finite;
- for every $e \in E$, $\{e\} \in \text{Con}$;
- if $X \in \text{Con}$ and $Y \subseteq X$ then $Y \in \text{Con}$; and
- if $X \in \text{Con}$, $e \in X$, and $e' \leq e$, then $X \cup \{e'\} \in \text{Con}$.

A **configuration** of an event structure E is a subset of events $x \subseteq E$ which is consistent ($x \in \text{Con}$) and down-closed (if $e \in x$ and $e' \leq e$ then $e' \in x$). The set of all *finite* configurations of E is denoted by $\mathcal{C}(E)$. It is a partial order under inclusion \subseteq , with least element the empty configuration \emptyset .

When $x, y \in \mathcal{C}(E)$, we write $x \prec y$ and say y **covers** x when there is $e \in E$ such that $y = x \cup \{e\}$; we may also write $x \prec^e y$. The axioms of event structures ensure the following important fact: given $x \in \mathcal{C}(E)$ there exists a **covering chain** for x , *i.e.* a sequence x_1, \dots, x_n of configurations such that $\emptyset \prec x_1 \prec \dots \prec x_n \prec x$. Covering chains are not necessarily unique.

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

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

with componentwise causality (and polarity). The consistent sets are those of the form $\parallel_{i \in I} X_i$ for I , such that $X_i \in \text{Con}_{A_i}$ for each $i \in I$. (Observe that any configuration $x \in \mathcal{C}(\parallel_{i \in I} A_i)$ corresponds to $\parallel_{i \in I} x_i$ where each $x_i \in \mathcal{C}(A_i)$, and at most finitely many x_i are nonempty).

2.1.2 Games and strategies

In concurrent games, event structures are used to represent both games and strategies — so in programming languages, types and terms. Events correspond to moves, and we use a polarity function to distinguish between moves of each player. An **event structure with polarity** (esp) is an event structure (E, \leq, Con) equipped with a polarity function $\text{pol} : E \rightarrow \{+, -\}$. The polarity $\text{pol}(e)$ indicates which of the two players is responsible for the move e : **Player** (+) or **Opponent** (-).

We will often keep the data implicit and refer to an esp $(E, \leq, \text{Con}, \text{pol})$ simply as E . Moreover we use the following notation: if $x, y \in \mathcal{C}(E)$ and $x \subseteq y$, we write $x \subseteq^+ y$ (resp. $x \subseteq^- y$) if every $e \in y \setminus x$ has positive (resp. negative) polarity. The relations \prec^+ and \prec^- are defined analogously.

Definition 2.3. A **game** is an event structure with polarity.

So a game can be seen as a set of available moves for each player, along with compatibility and causal constraints between them.

Informally, a *strategy* should specify Player's behaviour at any point in the game, according to the moves Opponent chooses to play. (Note, by *strategy* we always mean a strategy for Player.) We encode this as another event structure with polarity S whose events are in some sense *labelled* by those of the game, via a map $\sigma : S \rightarrow A$. To ensure that the strategy S obeys the causal constraints imposed by the game, the map σ will be required to satisfy the axioms of a *map of event structures*.

Definition 2.4. If E and E' are event structures, a **map of event structures** $f : E \rightarrow E'$ is a function on events which

- *preserves configurations*: for every $x \in \mathcal{C}(E)$, the direct image $fx \in \mathcal{C}(E')$; and
- *is locally injective*: for every $x \in \mathcal{C}(E)$ and $e, e' \in x$, if $f(e) = f(e')$ then $e = e'$.

If E and E' have polarity, f must additionally preserve it.

The local injectivity condition gives for each $x \in \mathcal{C}(E)$ a bijection $x \cong fx$.

Although maps of event structures are defined on events, their action on configurations is sometimes easier to describe. So we begin by observing the following property:

Lemma 2.5. *Suppose $f, g : E \rightarrow E'$ are maps of event structures. If $fx = gx$ for every $x \in \mathcal{C}(E)$, then $f = g$.*

Proof. Let $e \in E$. By assumption, $f[e] = g[e]$ and $f[e] = g[e]$. Because the restrictions of f and g to $[e]$ are injective, $f(e) = f[e] \setminus f[e] = g[e] \setminus g[e] = g(e)$. \square

It is also helpful to know when a function of configurations is generated by a map of event structures:

Lemma 2.6. *Let $f : \mathcal{C}(E) \rightarrow \mathcal{C}(E')$ be a monotone function which preserves cardinality ($|f(x)| = |x|$) and preserves unions ($f(x \cup y) = fx \cup fy$). Then, there is a map of event structures $E \rightarrow E'$ whose action on configurations corresponds to f .*

The proof is straightforward and can be found in [Cas17]. We can now define concurrent strategies:

Definition 2.7. For a game A , a **strategy on A** consists of an event structure with polarity S , together with a map of event structures $\sigma : S \rightarrow A$ which is

- **receptive**: if $x \in \mathcal{C}(S)$ and $\sigma x \subseteq^- y$ for some $y \in \mathcal{C}(A)$, there exists a unique $x' \in \mathcal{C}(S)$ such that $x \subseteq x'$ and $\sigma x' = y$;
- **courteous**: if $e, e' \in S$ are such that $e \rightarrow e'$ and $\sigma(e) \not\vdash \sigma(e')$, then $\text{pol}(e) = -$ and $\text{pol}(e') = +$.

In words, the *courtesy* axiom says that a strategy may only specify additional causal dependencies of Player moves on Opponent moves. The *receptivity* axiom states that at any stage Player must be prepared to let Opponent play the moves that the game A makes available to them.

Although sufficient in some contexts, for the purposes of this thesis games and strategies as presented above must be enriched with extra structure. The issue is with *duplication*: we expand on this in what follows.

2.2 Symmetry in concurrent games

In order to model a programming language with no restriction on resource usage (specifically, a non-*affine* language), it will be important to allow for moves of the game to be played more than once. This is not possible with the above because strategies must be locally injective.

One could consider an expanded version of the game where each move exists in many copies, but this naive approach fails in ensuring that Player's behaviour is uniform with respect to Opponent's particular choice of copy.

This issue is addressed by considering games and strategies based on *event structures with symmetry* [Win07], which extend event structures with extra information expressing when configurations should be considered equivalent. The extra structure and associated axioms will ensure that strategies are uniform.

We first define event structures with symmetry, and then see how the games and strategies above can in turn be extended with symmetry.

2.2.1 Symmetry in event structures

There are multiple presentations of event structures with symmetry [Win07]; the following best suits our purposes:

Definition 2.8. An **isomorphism family** on an event structure E is a set \cong_E of bijections $\theta : x \cong y$, for $x, y \in \mathcal{C}(E)$, satisfying the following axioms

(*Groupoid*) \cong_E contains all identity bijections, and is closed under composition and inverse;

(*Restriction*) if $(\theta : x \cong y) \in \cong_E$ and $x' \subseteq x$, then there exists $y' \subseteq y$ and $(\theta' : x' \cong y') \in \cong_E$ such that $\theta' \subseteq \theta$ (where \subseteq is inclusion of graphs); and

(*Extension*) if $(\theta : x \cong y) \in \cong_E$ and $x \subseteq x'$, then there exists $y' \in \mathcal{C}(E)$ with $y \subseteq y'$, and $(\theta' : x' \cong y') \in \cong_E$, such that $\theta \subseteq \theta'$ (note that θ' is not necessarily unique).

Additionally, if E has polarity, we ask that each $\theta \in \cong_E$ preserve it.

Definition 2.9. An **event structure with symmetry (and polarity)** is a pair $\mathcal{E} = (E, \cong_E)$ of an event structure (with polarity) and an isomorphism family on it.

We use $\mathcal{E}, \mathcal{A}, \mathcal{B}, \mathcal{S}, \mathcal{T}$, *etc.* to denote event structures with symmetry; it is understood that E, A, B, S, T , *etc.* denote the underlying event structures. To indicate that a bijection $\theta : x \cong y$ is in \cong_E , we will write $\theta : x \cong_E y$. We do not distinguish between the bijection θ , and its *graph* $\{(e, \theta(e)) \mid e \in x\}$; in particular the set operations \subseteq, \subset, \cup , *etc.* extend to bijections; and because elements of \cong_E preserve polarity when it is there, the polarised extensions \subseteq^+, \subset^+ , *etc.* are also well-defined.

Example 2.10. A process featuring countably many copies of a single event can be modelled as an event structure with symmetry \mathcal{E} as follows: $E = \mathbb{N}$; causality is trivial ($i \leq j$ iff $i = j$); all subsets are consistent; and the isomorphism family \cong_E contains all bijections $\theta : x \cong y$ between finite configurations of E (i.e. finite subsets of \mathbb{N}).

More generally, suppose \mathcal{A} is an event structure with symmetry and polarity. An expanded version of it is given by the *esp* $!A$ which has underlying *esp* $!A = \parallel_{i \in \omega} \mathcal{A}$. Its isomorphism family $\cong_{!A}$ contains those $\theta : \parallel_{i \in \omega} x_i \cong \parallel_{i \in \omega} y_i$ for which there is $\pi : \omega \cong \omega$ and $\theta_i : x_i \cong_{A_i} x_{\pi(i)}$ such that $\theta(i, a) = (\pi(i), \theta_i(a))$ for each $(i, a) \in !A$. (Lemma 2.15 below states that this satisfies the axioms for an isomorphism family.)

Definition 2.11. An map $\mathcal{E} \rightarrow \mathcal{E}'$ of event structures with symmetry is a map $f : E \rightarrow E'$ which additionally preserves symmetry: if $\theta : x \cong_E y$ then $f\theta : fx \cong_{E'} fy$.

2.2.2 Arenas

A general model of games and strategies with symmetry may then be defined, following [CCW14]. There, a game can be an arbitrary event structure with symmetry, and uniformity of strategies is enforced via a *saturation* condition.

The work in this thesis follows an alternative approach, introduced in [CCW15], in which games are event structures with symmetry equipped with two further isomorphism families, containing bijections only affecting moves of one of the two players. This concept appears in [CCW15], and in more details in [CCW19], under the name of *thin concurrent game*. We choose this approach over the saturated strategies of [CCW14], as it will prove much more accommodating to probability.

Definition 2.12. A **game with symmetry** is a tuple $\mathcal{A} = (A, \cong_A, \cong_A^+, \cong_A^-)$, where A is a game and \cong_A, \cong_A^+ , and \cong_A^- are three isomorphism families on A such that

- the families \cong_A^+ and \cong_A^- are sub-families of \cong_A ;
- if $\theta \in \cong_A^- \cap \cong_A^+$, then θ is the identity bijection on some $x \in \mathcal{C}(A)$;
- if $\theta \in \cong_A^-$ and $\theta \subseteq^- \theta'$ for some $\theta' \in \cong_A$ then $\theta' \in \cong_A^-$; and
- if $\theta \in \cong_A^+$ and $\theta \subseteq^+ \theta'$ for some $\theta' \in \cong_A$ then $\theta' \in \cong_A^+$.

We extend the notation for symmetric configurations: if $\theta : x \cong y$ is in the isomorphism \cong_A^+ , we write $\theta : x \cong_A^+ y$. The following property of games with symmetry is easily derived from the axioms:

Lemma 2.13. *Let \mathcal{A} be a game with symmetry and let $\theta : x \cong_A^+ y$. Suppose $x \subseteq^- x' \in \mathcal{C}(A)$. Then there is a unique extension $\theta \subseteq \theta' : x' \cong_A^+ y'$.*

Proof. Suppose there are extensions $\theta' : x' \cong_A^+ y'$ and $\theta'' : x' \cong_A^+ y''$. Then the bijection $\theta'' \circ \theta'^{-1} : y' \cong y''$ is in \cong_A^+ , but it is also a negative extension of $\text{id}_y \in \cong_A^-$, so necessarily in \cong_A^- by the third axiom of Definition 2.12. Thus by the second axiom $\theta'' \circ \theta'^{-1}$ is an identity bijection, which implies the result. \square

We will be particularly interested in the following class of games:

Definition 2.14. An **arena** is a game with symmetry \mathcal{A} whose underlying *esp* A is

- **forest-shaped:** if $a \leq b$ and $c \leq b$ then either $a \leq c$ or $c \leq a$; and
- **alternating:** if $a \rightarrow b$ then $\text{pol}(a) = -\text{pol}(b)$.

Say an arena \mathcal{A} is **negative** (resp. **positive**) if all its initial moves have negative (resp. positive) polarity. If \mathcal{A} is either negative or positive we say it is **polarised**.

The classic example of a game with symmetry is obtained via the construction of $!A$ in Example 2.10, given an essp \mathcal{A} . The associated isomorphism family $\cong_{!A}$ allows for an arbitrary choice of *copy index* for the moves of \mathcal{A} . When \mathcal{A} is a polarised arena, there are two sub-families $\cong_{!A}^+$ and $\cong_{!A}^-$ defined as follows:

- $\cong_{!A}^-$ contains bijections $\theta : x \cong_{\cong_{!A}} y$ such that for each $i \in \omega$ the bijection $\theta_i : x_i \cong y_{\pi(i)}$ is in the isomorphism family $\cong_{\bar{A}}$, and moreover if $\pi(i) \neq i$, then all initial moves of x_i have negative polarity.
- $\cong_{!A}^+$ contains bijections $\theta : x \cong_{\cong_{!A}} y$ such that for each $i \in \omega$ the bijection $\theta_i : x_i \cong y_{\pi(i)}$ is in the isomorphism family $\cong_{\bar{A}}^+$, and moreover if $\pi(i) \neq i$, then all initial moves of x_i have positive polarity.

Lemma 2.15. *When \mathcal{A} is a polarised arena, the tuple $(!A, \cong_{!A}, \cong_{!A}^-, \cong_{!A}^+)$ is an arena with same polarity.*

Proof. We first check that $\cong_{!A}, \cong_{!A}^-$, and $\cong_{!A}^+$ are isomorphism families. We check the axioms for the first one.

(Groupoid) If $x \in \mathcal{C}(!A)$ then the identity bijection $\theta : x \cong x$ is in $\cong_{!A}$ since (taking π to be the identity on ω) all $\theta_i : x_i \cong x_i$ are in \cong_A . If $\theta : x \cong_{!A} y$, then we have bijections $\theta_i^{-1} : y_i \cong_A x_{\pi^{-1}(i)}$ for each i and therefore $\theta^{-1} : y \cong_{!A} x$. If $\theta : x \cong_{!A} y$ and $\varphi : y \cong_{!A} z$ with reindexing bijections π_θ and π_φ respectively, we can take $\pi = \pi_\varphi \circ \pi_\theta$ and let $(\varphi \circ \theta)_i : x_i \cong z_{\pi(i)}$, clearly in \cong_A for each i . So $\varphi \circ \theta \in \cong_{!A}$.

(Restriction) Let $\theta : x \cong_{!A} y$ and $x' \subseteq x$. If $x = \bigcup_{i \in \omega} \{i\} \times x_i$ then we can write $x' = \bigcup_{i \in \omega} \{i\} \times x'_i$, where for all i we have $x'_i \subseteq x_i$. The restriction of each $\theta_i : x_i \cong y_{\pi(i)}$ to x'_i is a bijection $\theta'_i : x'_i \cong_A y'_{\pi(i)}$, and so with the same reindexing π we have a bijection $\theta' : x' \cong_{!A} y'$, where $y' = \bigcup_{i \in \omega} \{i\} \times y'_i$.

The same argument can be used to check the (Extension) axiom, and given that every bijection in an isomorphism family must preserve polarity, it is straightforward to do the same verifications for $\cong_{!A}^-$ and $\cong_{!A}^+$. We show now that $(!A, \cong_{!A}, \cong_{!A}^-, \cong_{!A}^+)$ satisfies the axioms of a game with symmetry.

The isomorphism families $\cong_{!A}^-$ and $\cong_{!A}^+$ are subsets of $\cong_{!A}$ by definition. Suppose $\theta : x \cong y$ is in $\cong_{!A}^- \cap \cong_{!A}^+$ and let $(i, a) \in x$. By definition $\theta(i, a) = (\pi(i), \theta_i a)$ for some $\theta_i \in \cong_{\bar{A}}^- \cap \cong_{\bar{A}}^+$, which must be the identity given that \mathcal{A} is a game with symmetry. Moreover if $i \neq \pi(i)$, then the initial moves of x_i have both negative and positive polarity. But x_i is nonempty since $(i, a) \in x$, so we must have $i = \pi(i)$ and θ is the identity bijection.

Suppose now that $\theta : x \cong_{\cong_{!A}^-} y$ and $\theta : x' \cong_{\cong_{!A}^-} y'$ with reindexing π and π' respectively, such that $\theta \subseteq^- \theta'$. For $i \in \omega$, we have that $\theta_i \subseteq^- \theta'_i$ and therefore $\theta'_i \in \cong_{\bar{A}}^-$. Suppose $\pi'(i) \neq i$. If x_i has any initial moves, then $\pi(i) = \pi'(i)$ and so the initial moves of x_i must be negative since $\theta \in \cong_{!A}^-$, and $x'_i \setminus x_i$ only has negative moves by assumption. \square

2.2.3 Thin strategies

To ensure uniformity, a strategy on a game with symmetry \mathcal{A} is defined as a map $\sigma : \mathcal{S} \rightarrow \mathcal{A}$ where \mathcal{S} too is equipped with symmetry. There are extra conditions, which we discuss below:

Definition 2.16. Given an arena \mathcal{A} , a **strategy on \mathcal{A}** consists of an event structure with symmetry and polarity \mathcal{S} , and a map $\sigma : \mathcal{S} \rightarrow (A, \cong_A)$, such that:

- the underlying map $\sigma : S \rightarrow A$ is a strategy in the sense of Definition 2.7;
- σ is **\sim -receptive**: if $\theta : x \cong_S y$ and $\sigma\theta \subseteq^- \psi$ then there exists a unique $\theta' \in \cong_S$ such that $\theta \subseteq \theta'$ and $\sigma\theta' = \psi$; and
- σ is **thin**: if $x \in \mathcal{C}(S)$ and $\text{id}_x \subseteq^+ \theta$ for some $\theta \in \cong_S$, then $\theta = \text{id}_{x'}$ for some $x' \in \mathcal{C}(S)$.

The \sim -receptivity condition is very natural: informally, because σ is a strategy for Player, the symmetry of Opponent moves in the strategy should be canonically induced from the symmetry of the game. Thinness is more subtle: it says that there is no non-trivial “positive symmetry” in a strategy. This does *not* mean that Player is not allowed to play several copies of the same move, only that those should not be recorded as symmetric in \mathcal{S} .

Thin strategies, introduced in [CCW15], can be considered “up to positive symmetry”; in other words, it is possible to consider two strategies equivalent if they only differ in Player’s choice of copy indices. The key insight of [CCW15] is that thinness makes this equivalence a congruence for composition of strategies. In this thesis, this issue arises in the construction of a bicategory of games and thin strategies (specifically, when defining *horizontal composition* of 2-cells).

The next section shows how strategies compose.

2.3 Composing strategies

2.3.1 Interaction

We start with an important definition: the **dual** A^\perp of an event structure with polarity A is the same event structure with polarity function reversed: $\text{pol}_{A^\perp} = -\text{pol}_A$. If $\mathcal{A} = (A, \cong_A, \cong_A^-, \cong_A^+)$ is an arena, we define $\mathcal{A}^\perp = (A^\perp, \cong_A, \cong_A^+, \cong_A^-)$, *i.e.* the roles of \cong_A^+ and \cong_A^- are swapped.

Suppose \mathcal{A} is a fixed arena. Consider a strategy $\sigma : \mathcal{S} \rightarrow \mathcal{A}$, and a strategy $\tau : \mathcal{T} \rightarrow \mathcal{A}^\perp$ on the dual arena. Just as σ specifies the behaviour of Player, τ can be thought of as a strategy for Opponent, sometimes called a **counterstrategy**.

We consider the situation where Player follows strategy σ and Opponent strategy τ . The goal of this section is to use constructions on event structures to understand this interaction.

In particular, the pullback

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

always exists in the category of event structures with symmetry, where we abuse the notation slightly: \mathcal{A} refers to the ess (A, \cong_A) without the extra structure \cong_A^- and \cong_A^+ , and all polarity information is ignored (otherwise $\tau : \mathcal{T} \rightarrow \mathcal{A}$ would not be a valid map).

The ess $\mathcal{S} \wedge \mathcal{T}$ models the result of the interaction of strategies \mathcal{S} and \mathcal{T} . Every configuration of $\mathcal{S} \wedge \mathcal{T}$ corresponds to the synchronisation of a configuration of \mathcal{S} and a configuration of \mathcal{T} . Thus σ and τ must agree on a set of moves, and the causal constraints imposed on Player and Opponent must be compatible.

Formally, suppose $x \in \mathcal{C}(\mathcal{S})$ and $y \in \mathcal{C}(\mathcal{T})$ are such that $\sigma x = \tau y$. Then, by the local injectivity condition on maps of event structures, so that there is a bijection $\theta : x \cong \sigma x = \tau y \cong y$. We say that such a bijection is **secured** if the transitive closure of the relation \trianglelefteq defined for $(s, t), (s', t') \in \theta$ as

$$(s, t) \trianglelefteq (s', t') \text{ iff } s \leq_S s' \text{ or } t \leq_T t'$$

is also anti-symmetric.

An event structure is fully determined by its set of configurations [Win86]. Thus the event structure with symmetry $\mathcal{S} \wedge \mathcal{T}$ is characterised by the following result:

Lemma 2.17 ([CCW19]). *Configurations of $\mathcal{S} \wedge \mathcal{T}$ correspond to pairs $(x, y) \in \mathcal{C}(\mathcal{S}) \times \mathcal{C}(\mathcal{T})$ such that $\sigma x = \tau y$ and the composite bijection $x \cong \sigma x = \tau y \cong y$ is secured. The isomorphism family $\cong_{\mathcal{S} \wedge \mathcal{T}}$ comprises, for every $z, z' \in \mathcal{C}(\mathcal{S} \wedge \mathcal{T})$, the bijections $\theta : z \cong z'$ such that $\Pi_1 \theta : \Pi_1 z \cong_S \Pi_1 z'$ and $\Pi_2 \theta : \Pi_2 z \cong_T \Pi_2 z'$.*

2.3.2 Strategies as morphisms

For arenas \mathcal{A} and \mathcal{B} , a **strategy from \mathcal{A} to \mathcal{B}** is a strategy on the arena $\mathcal{A}^\perp \parallel \mathcal{B}$. Let \mathcal{C} be a third arena, and suppose $\sigma : \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$ and $\tau : \mathcal{T} \rightarrow \mathcal{B}^\perp \parallel \mathcal{C}$ are strategies. The *composition* of σ and τ will arise after the two strategies synchronise on the game \mathcal{B} , and this synchronisation is hidden so as to leave only \mathcal{A} and \mathcal{C} open for playing.

2.3.2.1 Interaction

We form the pullback

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

in which all polarity information is ignored, since in particular τ and σ disagree on the polarity assigned to events of B . We denote by $\mathcal{T} \circledast \mathcal{S}$ the pullback $(\mathcal{S} \parallel \mathcal{C}) \wedge (\mathcal{A} \parallel \mathcal{T})$, called the **interaction** of σ and τ . The composite map $\mathcal{T} \circledast \mathcal{S} \rightarrow \mathcal{A} \parallel \mathcal{B} \parallel \mathcal{C}$ induced by the pullback diagram is written $\tau \circledast \sigma$.

If $e \in \mathcal{T} \circledast \mathcal{S}$ is mapped to the A - or the C -component of $\mathcal{A} \parallel \mathcal{B} \parallel \mathcal{C}$, then its polarity is determined as follows: we take it to be $\text{pol}_{A^\perp}((\tau \circledast \sigma)e)$ or $\text{pol}_C((\tau \circledast \sigma)e)$, accordingly. For events mapped to the B -component, the polarity is ambiguous because σ and τ disagree. But we can recover a form of polarity by assigning **neutral** polarity to those events.

In addition to the negative/neutral/positive description of events of $\mathcal{T} \circledast \mathcal{S}$, the following terminology will be useful: we say $e \in \mathcal{T} \circledast \mathcal{S}$ is a σ -**action** if $\Pi_1 e$ is a positive event of \mathcal{S} , and a τ -**action** if $\Pi_2 e$ is a positive event of \mathcal{T} . No event is both a σ -action and a τ -action, and negative events of $\mathcal{T} \circledast \mathcal{S}$ are neither of the two.

2.3.2.2 Hiding and composition

To compose strategies τ and σ we need to hide the synchronisation events. Accordingly we will hide from $\mathcal{T} \circledast \mathcal{S}$ the events with neutral polarity. The remaining events are called **visible**.

To do this we use the following general construction on event structures with symmetry [CCW19]. Let \mathcal{E} be an esp and $V \subseteq E$ a subset of events closed under symmetry, so if $\theta : x \cong_E y$ and $e \in x \cap V$, then $\theta(e) \in V$. Then $E \downarrow V$ is the esp with events V , and causality, consistency and polarity directly induced from E . Importantly, a configuration $x \in \mathcal{C}(E \downarrow V)$ has a **unique witness** $[x] \in \mathcal{C}(E)$ obtained as the down-closure of x in E . The isomorphism family $\cong_{E \downarrow V}$ is the set

$$\{\theta : x \cong y \mid x, y \in \mathcal{C}(E \downarrow V) \text{ and } \exists \theta' : [x] \cong_E [y] \text{ with } \theta \subseteq \theta'\}$$

giving an esp $\mathcal{E} \downarrow V = (E \downarrow V, \cong_{E \downarrow V})$.

So from the map $\tau \circledast \sigma : \mathcal{T} \circledast \mathcal{S} \rightarrow \mathcal{A} \parallel \mathcal{B} \parallel \mathcal{C}$, to obtain a map to $\mathcal{A} \parallel \mathcal{C}$ we first consider $(\mathcal{T} \circledast \mathcal{S}) \downarrow V$, where V contains \mathcal{A} -moves and \mathcal{C} -moves only:

$$V = \{e \in \mathcal{T} \circledast \mathcal{S} \mid (\tau \circledast \sigma)(e) = (1, a) \text{ with } a \in \mathcal{A}, \text{ or } (3, c) \text{ with } c \in \mathcal{C}\}.$$

We denote $(\mathcal{T} \circledast \mathcal{S}) \downarrow V$ by $\mathcal{T} \odot \mathcal{S}$. Restricting the map $\tau \circledast \sigma$ to $\mathcal{T} \odot \mathcal{S}$ gives a map of event structures with symmetry $\tau \odot \sigma : \mathcal{T} \odot \mathcal{S} \rightarrow \mathcal{A} \parallel \mathcal{C}$. Polarity can be recovered, unambiguously: $\text{pol}_{\mathcal{T} \odot \mathcal{S}}(e) \triangleq \text{pol}_{\mathcal{A}^\perp \parallel \mathcal{C}}((\tau \odot \sigma)(e))$.

The map $\tau \odot \sigma : \mathcal{T} \odot \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{C}$ satisfies all the axioms for a strategy from \mathcal{A} to \mathcal{C} [CCW15]; it is the **composition** of τ and σ .

2.3.3 Causally compatible pairs

Before continuing with the bicategorical development, we devote the rest of this section to a more detailed description of the essps $\mathcal{T} \circledast \mathcal{S}$ and $\mathcal{T} \odot \mathcal{S}$. The results we present will be useful in the rest of the thesis.

By Lemma 2.17, configurations of $T \circledast S$ correspond to pairs $(x_S \parallel x_C, x_A \parallel x_T) \in \mathcal{C}(S \parallel C) \times \mathcal{C}(A \parallel T)$ for which the bijection

$$x_S \parallel x_C \cong \sigma x_S \parallel x_C = x_A \parallel \tau x_T \cong x_A \parallel x_T$$

is secured. Since x_A and x_C are determined by x_S and x_T (respectively), the pairs (x_S, x_T) suffice to characterise the configurations of $T \circledast S$.

Definition 2.18. Configurations $x_S \in \mathcal{C}(S)$ and $x_T \in \mathcal{C}(T)$ are said to be **causally compatible** when they are matching, *i.e.* there is $x_A \parallel x_B \parallel x_C \in \mathcal{C}(A \parallel B \parallel C)$ such that $\sigma x_S = x_A \parallel x_B$ and $\tau x_T = x_B \parallel x_C$, and the induced bijection $x_S \parallel x_C \cong \sigma x_S \parallel x_C = x_A \parallel \tau x_T \cong x_A \parallel x_T$ is secured.

So, configurations of $\mathcal{C}(T \circledast S)$ correspond to causally compatible pairs, and we write $x_T \circledast x_S$ for the configuration corresponding to the pair (x_S, x_T) . This notation, borrowed from [CC16], will significantly help the technical development. Note however that \circledast is *not* a total operator on configurations: writing $x_T \circledast x_S$ assumes that x_S and x_T are causally compatible.

Elements of the isomorphism family $\cong_{T \circledast S}$ can also be seen as matching pairs:

Lemma 2.19. *Bijections $\theta : x_T \circledast x_S \cong y_T \circledast y_S$ in $\cong_{T \circledast S}$ correspond to pairs (θ_S, θ_T) , where $\theta_S : x_S \cong_S y_S$ and $\theta_T : x_T \cong_T y_T$ and the two are matching, *i.e.* $\sigma \theta_S \parallel \theta_C = \theta_A \parallel \tau \theta_T$ for some $\theta_A \in \cong_A, \theta_C \in \cong_C$. We write $\theta = \theta_T \circledast \theta_S$.*

Proof. By Lemma 2.17, a bijection $\theta : x_T \circledast x_S \cong y_T \circledast y_S$ is in $\cong_{T \circledast S}$ only if $\Pi_1 \theta \in \cong_{S \parallel C}$ and $\Pi_2 \theta \in \cong_{A \parallel T}$. Writing $\Pi_1 \theta = \theta_S \parallel \theta_C$ and $\Pi_2 \theta = \theta_A \parallel \theta_T$, the bijections θ_S and θ_T are matching, since $\sigma \theta_S \parallel \theta_C = (\tau \circledast \sigma) \theta = \theta_A \parallel \tau \theta_T$.

Conversely, given matching $\theta_S : x_S \cong_S y_S$ and $\theta_T : x_T \cong_T y_T$, we take $\theta : x_T \circledast x_S \cong_{T \circledast S} y_T \circledast y_S$ to be the bijection

$$x_T \circledast x_S \cong x_S \parallel x_C \stackrel{\varphi}{\cong} y_A \parallel y_T \cong y_T \circledast y_S$$

where φ stands for either

$$x_S \parallel x_C \cong x_A \parallel x_T \stackrel{\theta_A \parallel \theta_T}{\cong} y_A \parallel y_T \quad \text{or} \quad x_S \parallel x_C \stackrel{\theta_S \parallel \theta_C}{\cong} y_S \parallel y_C \cong y_A \parallel y_T.$$

(The two bijections are equal by the matching requirement.) It is routine to check that $\Pi_1 \theta = \theta_S \parallel \theta_C$ and $\Pi_2 \theta = \theta_A \parallel \theta_T$. \square

Lemma 2.20. *Configurations of $T \odot S$ correspond to causally compatible pairs $(x_S, x_T) \in \mathcal{C}(S) \times \mathcal{C}(T)$ which are **minimal** among those with the same projection to A and C .*

We write $x_T \odot x_S \in \mathcal{C}(T \odot S)$ for the configuration corresponding to (x_S, x_T) . As with $x_T \circledast x_S$, this notation assumes that x_S and x_T are causally compatible and that the pair is minimal.

2.4 Copycat and associativity of composition

2.4.1 Identity strategies

The following strategy, known as **copycat**, will act as identity on a game \mathcal{A} . It comprises an essp $\mathbb{C}_{\mathcal{A}}$ and a map

$$\alpha_{\mathcal{A}} : \mathbb{C}_{\mathcal{A}} \rightarrow \mathcal{A}^{\perp} \parallel \mathcal{A}.$$

As the name suggests, copycat might informally be described as the strategy on $\mathcal{A}^{\perp} \parallel \mathcal{A}$ in which Player reproduces on each side Opponent's behaviour on the other.

Formally, the events, polarity and consistency of $\mathbb{C}_{\mathcal{A}}$ are those of $\mathcal{A}^{\perp} \parallel \mathcal{A}$, and the causality is that of $\mathcal{A}^{\perp} \parallel \mathcal{A}$ enriched with the pairs $\{((a, 1), (a, 2)) \mid a \in A \text{ and } \text{pol}_A(a) = +\} \cup \{((a, 2), (a, 1)) \mid \text{pol}_A(a) = -\}$.

Note that because $\mathbb{C}_{\mathcal{A}}$ is just $\mathcal{A}^{\perp} \parallel \mathcal{A}$ with added causal constraints, configurations of $\mathbb{C}_{\mathcal{A}}$ can be seen as a subset of those of $\mathcal{A}^{\perp} \parallel \mathcal{A}$. The following characterisation will be useful:

Lemma 2.21. $\mathcal{C}(\mathbb{C}_{\mathcal{A}}) = \{x_1 \parallel x_2 \in \mathcal{C}(A \parallel A) \mid x_1 \supseteq^+ x_1 \cap x_2 \subseteq^- x_2\}$.

The isomorphism family $\cong_{\mathbb{C}_{\mathcal{A}}}$ contains bijections $\theta_1 \parallel \theta_2 : x_1 \parallel x_2 \cong x'_1 \parallel x'_2$ with $\theta_1 : x_1 \cong_{A^{\perp}} x'_1$, $\theta_2 : x_2 \cong_A x'_2$, and $\theta_1^{\perp} \supseteq^+ \theta_1^{\perp} \cap \theta_2 \subseteq^- \theta_2$.

Because of this characterisation, the relation $x_1 \sqsubseteq x_2$ defined as $x_1 \supseteq^+ x_1 \cap x_2 \subseteq^- x_2$ between configurations of A plays a significant role. It is a partial order, called the **Scott order**.

2.4.2 Copycat is an identity

As it tends to be the case when composition is given by a universal construction, there is no strict notion of identity in concurrent games. Indeed games and strategies are not a category, but a *bicategory*, whose construction we will formally give in the next section. Here we simply discuss in which *weaker* sense the copycat strategy is the identity strategy on a game.

In general for a strategy $\sigma : \mathcal{S} \rightarrow \mathcal{A}^{\perp} \parallel \mathcal{B}$, there are isomorphisms of essps $\lambda_{\sigma} : \mathbb{C}_{\mathcal{B}} \odot \mathcal{S} \cong \mathcal{S}$ and $\rho_{\sigma} : \mathcal{S} \odot \mathbb{C}_{\mathcal{A}} \cong \mathcal{S}$ making the following diagrams commute:

$$\begin{array}{ccc} \mathbb{C}_{\mathcal{B}} \odot \mathcal{S} & \xrightarrow{\lambda_{\sigma}} & \mathcal{S} \\ \alpha_{\mathcal{S}} \odot \sigma \searrow & & \swarrow \sigma \\ & \mathcal{A}^{\perp} \parallel \mathcal{B} & \end{array} \qquad \begin{array}{ccc} \mathcal{S} \odot \mathbb{C}_{\mathcal{A}} & \xrightarrow{\rho_{\sigma}} & \mathcal{S} \\ \sigma \odot \alpha_{\mathcal{A}} \searrow & & \swarrow \sigma \\ & \mathcal{A}^{\perp} \parallel \mathcal{B} & \end{array}$$

Showing the existence of these isomorphisms is not an easy task [CCRW17], and requires a careful analysis of the interactions $\mathbb{C}_{\mathcal{B}} \circledast \mathcal{S}$ and $\mathcal{S} \circledast \mathbb{C}_{\mathcal{A}}$. We do not reproduce the full proof here; however we give sufficient detail that similar proofs can be carried out in the probabilistic setting.

We consider only the composition $\sigma \odot \alpha_{\mathcal{A}}$: the argument for $\alpha_{\mathcal{B}} \odot \sigma$ is symmetric, and anyway follows from the first case since σ can also be seen as a strategy from \mathcal{B}^{\perp} to \mathcal{A}^{\perp} , and from this point of view pre-composition is post-composition.

Recall that configurations of \mathbb{C}_A are of the form $y_A \parallel x_A$ for $x_A, y_A \in \mathcal{C}(A)$ with $x_A \sqsubseteq y_A$ (where \sqsubseteq is the ‘‘Scott order’’). Interaction with copycat does not induce any causal loops (we say it is **deadlock-free**), so that configurations of $\mathcal{S} \circledast \mathbb{C}_A$ correspond to pairs of configurations $x_S \in \mathcal{C}(S)$ and $y_A \parallel x_A \in \mathcal{C}(\mathbb{C}_A)$ such that $\sigma x_S = x_A \parallel x_B$ for some $x_B \in \mathcal{C}(B)$. Similarly, configurations of $\mathbb{C}_B \circledast \mathcal{S}$ are $(y_B \parallel x_B) \circledast x_S$, where $x_B \sqsubseteq y_B$ (so $y_B \parallel x_B \in \mathcal{C}(\mathbb{C}_B)$), and $\sigma x_S = x_A \parallel y_B$ for some $x_A \in \mathcal{C}(A)$.

Lemma 2.22 ([CCRW17, Cas17]). *There are isomorphisms $\rho_\sigma : \mathcal{S} \circledast \mathbb{C}_A \rightarrow \mathcal{S}$ and $\lambda_\sigma : \mathbb{C}_B \circledast \mathcal{S} \rightarrow \mathcal{S}$, making the above diagrams commute, and acting on configurations as follows.*

- To $x_S \circledast (y_A \parallel x_A) \in \mathcal{C}(\mathcal{S} \circledast \mathbb{C}_A)$, where $\sigma x_S = x_A \parallel x_B$ and $x_A \sqsubseteq y_A$, ρ_σ associates the unique $x_S^* \in \mathcal{C}(S)$ such that $x_S^* \sqsubseteq x_S$ and $\sigma x_S^* = y_A \parallel x_B$.
- To $(x_B \parallel y_B) \circledast x_S \in \mathcal{C}(\mathbb{C}_B \circledast \mathcal{S})$, where $\sigma x_S = x_A \parallel x_B$ and $y_B \sqsubseteq x_B$, λ_σ associates the unique $x_S^* \in \mathcal{C}(S)$ such that $x_S^* \sqsubseteq x_S$ and $\sigma x_S^* = x_A \parallel y_B$.

2.4.3 Associativity of composition

Consider three strategies $\sigma : \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$, $\tau : \mathcal{T} \rightarrow \mathcal{B}^\perp \parallel \mathcal{C}$ and $\rho : \mathcal{R} \rightarrow \mathcal{C}^\perp \parallel \mathcal{D}$. Their composition is not associative on the nose, that is: $\rho \circledast (\tau \circledast \sigma) \neq (\rho \circledast \tau) \circledast \sigma$. However, the strategies are *isomorphic* via a map $\alpha_{\sigma, \tau, \rho} : (\mathcal{R} \circledast \mathcal{T}) \circledast \mathcal{S} \Rightarrow \mathcal{R} \circledast (\mathcal{T} \circledast \mathcal{S})$ such that

$$\begin{array}{ccc} (\mathcal{R} \circledast \mathcal{T}) \circledast \mathcal{S} & \xrightarrow{\alpha_{\sigma, \tau, \rho}} & \mathcal{R} \circledast (\mathcal{T} \circledast \mathcal{S}) \\ & \searrow & \swarrow \\ & (\rho \circledast \tau) \circledast \sigma & \rho \circledast (\tau \circledast \sigma) \\ & \mathcal{A}^\perp \parallel \mathcal{D} & \end{array}$$

commutes. The map was given in [CCRW17] and extended to games with symmetry in [Cas17]. It is obtained by first considering a canonical isomorphism

$$\alpha_{\rho, \tau, \sigma}^{\circledast} : (\mathcal{R} \circledast \mathcal{T}) \circledast \mathcal{S} \Rightarrow \mathcal{R} \circledast (\mathcal{T} \circledast \mathcal{S})$$

between the two obvious ways to make σ, τ , and ρ interact, and carefully restricting this to an isomorphism $\alpha_{\rho, \tau, \sigma} : (\mathcal{R} \circledast \mathcal{T}) \circledast \mathcal{S} \rightarrow \mathcal{R} \circledast (\mathcal{T} \circledast \mathcal{S})$ involving only visible events.

Lemma 2.23. *There is a strong isomorphism of strategies $\alpha_{\sigma, \tau, \rho} : (\rho \circledast \tau) \circledast \sigma \Rightarrow \rho \circledast (\tau \circledast \sigma)$ of strategies, natural in its arguments σ, τ and ρ , and such that $\alpha_{\rho, \tau, \sigma}((x_R \circledast x_T) \circledast x_S) = x_R \circledast (x_T \circledast x_S)$.*

The isomorphisms α, λ , and ρ are instrumental in defining a bicategory of games and strategies; this is what is done in [CCRW17, Cas17]. We give a different presentation, noticing that games and strategies form part of a larger object (a *pseudo-double category*) from which the bicategory eventually arises.

2.5 A pseudo-double category

We first give some motivation.

The bicategorical story. Associators and unitors are examples of *2-cells* – morphisms between morphisms, of the form

$$\begin{array}{ccc} \mathcal{S} & \xrightarrow{f} & \mathcal{S}' \\ \sigma \searrow & & \swarrow \sigma' \\ & \mathcal{A}^\perp \parallel \mathcal{B} & \end{array}$$

That the diagram commutes means that for any $s \in \mathcal{S}$, s and $f(s)$ correspond to the same move of the game. Because games are equipped with symmetry, we relax this slightly, allowing s and $f(s)$ to be *different copies* of the same move. All copies of each Opponent move must appear in both \mathcal{S} and \mathcal{S}' anyway, by the receptivity property, so this relaxation is only interesting for Player moves. With this relaxed notion of 2-cell, two strategies are considered *isomorphic* if they only differ in Player's choice of copy indices.

So in the bicategory of concurrent games, a 2-cell will be a map $f : \mathcal{S} \rightarrow \mathcal{S}'$ such that

$$\begin{array}{ccc} \mathcal{S} & \xrightarrow{f} & \mathcal{S}' \\ \sigma \searrow & \sim^+ & \swarrow \sigma' \\ & \mathcal{A}^\perp \parallel \mathcal{B} & \end{array}$$

where \sim^+ is an equivalence relation on maps defined as follows:

Definition 2.24. Two maps $f, f' : \mathcal{S} \rightarrow \mathcal{S}'$ of essps are said to be **symmetric**, written $f \sim f'$, if for all $x \in \mathcal{C}(S)$ the bijection $\varphi_x = \{(f(s), f'(s)) \mid s \in x\}$ is in the isomorphism family \cong_S . If \mathcal{A} is a game with symmetry, then we say $g, g' : \mathcal{S} \rightarrow \mathcal{A}$ are **positively symmetric** ($g \sim^+ g'$) if they are symmetric with respect to the isomorphism family \cong_A^+ .

We will see that games (with symmetry), strategies and maps of strategies form a bicategory. As a reminder, and to fix the notation, a **bicategory** consists of:

- a set \mathbb{C} of objects;
- for each $A, B \in \mathbb{C}$, a category $\mathbb{C}[A, B]$, whose morphisms are called **2-cells** and denoted with a double arrow, e.g. $f : \sigma \Rightarrow \sigma'$ for $\sigma, \sigma' \in \mathbb{C}[A, B]$;
- for each $A, B, C \in \mathbb{C}$, a composition functor $\odot : \mathbb{C}[B, C] \times \mathbb{C}[A, B] \rightarrow \mathbb{C}[A, C]$;
- for each $A \in \mathbb{C}$, an identity morphism $\text{id}_A \in \mathbb{C}[A, A]$;
- for each $A, B \in \mathbb{C}$ and for each $\sigma \in \mathbb{C}[A, B]$, invertible 2-cells $\lambda_\sigma : \text{id}_B \odot \sigma \Rightarrow \sigma$ and $\rho_\sigma : \sigma \odot \text{id}_A \Rightarrow \sigma$, the left and right **unitors**; and
- for each $A, B, C, D \in \mathbb{C}$ and for each $\sigma \in \mathbb{C}[A, B]$, $\tau \in \mathbb{C}[B, C]$ and $\eta \in \mathbb{C}[C, D]$, an invertible 2-cell $\alpha_{\sigma, \tau, \eta} : (\eta \odot \tau) \odot \sigma \Rightarrow \eta \odot (\tau \odot \sigma)$ called the **associator**.

Unitors and associator are subject to naturality and coherence conditions which we omit; a standard reference for this is [Lei98]. We do not normally distinguish between the name of a category and that of its set of objects; if there is any ambiguity we write $|\mathbb{C}|$ for the latter.

It is worth saying explicitly that the 2-cells of a bicategory can be composed in two distinct ways: as morphisms in a hom-category $\mathbb{C}[A, B]$, and via the composition functor \odot . In the bicategory of games this is as follows: given

$$\begin{array}{ccc} \mathcal{S} & \xrightarrow{f} & \mathcal{S}' \\ \sigma \searrow & \sim^+ & \swarrow \sigma' \\ \mathcal{A}^\perp \parallel \mathcal{B} & & \end{array} \qquad \begin{array}{ccc} \mathcal{S}' & \xrightarrow{f'} & \mathcal{S}'' \\ \sigma' \searrow & \sim^+ & \swarrow \sigma'' \\ \mathcal{A}^\perp \parallel \mathcal{B} & & \end{array}$$

we define their **vertical composition** as the map

$$\begin{array}{ccc} \mathcal{S} & \xrightarrow{f' \circ f} & \mathcal{S}'' \\ \sigma \searrow & \sim^+ & \swarrow \sigma'' \\ \mathcal{A}^\perp \parallel \mathcal{B} & & \end{array}$$

By contrast, for maps

$$\begin{array}{ccc} \mathcal{S} & \xrightarrow{f} & \mathcal{S}' \\ \sigma \searrow & \sim^+ & \swarrow \sigma' \\ \mathcal{A}^\perp \parallel \mathcal{B} & & \end{array} \qquad \begin{array}{ccc} \mathcal{T} & \xrightarrow{g} & \mathcal{T}' \\ \tau \searrow & \sim^+ & \swarrow \tau' \\ \mathcal{B}^\perp \parallel \mathcal{C} & & \end{array}$$

there exists a map

$$\begin{array}{ccc} \mathcal{T} \odot \mathcal{S} & \xrightarrow{g \odot f} & \mathcal{T}' \odot \mathcal{S}' \\ \tau \odot \sigma \searrow & \sim^+ & \swarrow \tau' \odot \sigma' \\ \mathcal{A}^\perp \parallel \mathcal{C} & & \end{array}$$

called the **horizontal composition** of f and g . Understanding why the map $g \odot f$ exists is technical, and an important contribution of [CCW19]; we will give more details later on.

Generalised maps and lifting. The notion of maps between strategies can be generalised further. In [CCRW17, Lemma 4.4] maps of the form

$$\begin{array}{ccc} \mathcal{S} & \xrightarrow{f} & \mathcal{T} \\ \sigma \downarrow & & \downarrow \tau \\ \mathcal{A}^\perp \parallel \mathcal{B} & \xrightarrow{g^\perp \parallel h} & \mathcal{C}^\perp \parallel \mathcal{D} \end{array}$$

are briefly discussed. The authors notice that for maps

$$\begin{array}{ccc}
\mathcal{S} & \xrightarrow{f} & \mathcal{S}' \\
\sigma \downarrow & & \downarrow \sigma' \\
\mathcal{A}^\perp \parallel \mathcal{B} & \xrightarrow{h_1^\perp \parallel h_2} & \mathcal{A}'^\perp \parallel \mathcal{B}'
\end{array}
\qquad
\begin{array}{ccc}
\mathcal{T} & \xrightarrow{g} & \mathcal{T}' \\
\tau \downarrow & & \downarrow \tau' \\
\mathcal{B}^\perp \parallel \mathcal{C} & \xrightarrow{h_2^\perp \parallel h_3} & \mathcal{B}'^\perp \parallel \mathcal{C}'
\end{array}$$

they can construct a map

$$\begin{array}{ccc}
\mathcal{T} \odot \mathcal{S} & \xrightarrow{g \odot f} & \mathcal{T}' \odot \mathcal{S}' \\
\tau \odot \sigma \downarrow & & \downarrow \tau' \odot \sigma' \\
\mathcal{A}^\perp \parallel \mathcal{C} & \xrightarrow{h_1^\perp \parallel h_3} & \mathcal{A}'^\perp \parallel \mathcal{C}'
\end{array}$$

Strictly speaking, our situation is different, since in [CCRW17] games have no symmetry. We will consider a “weak” variant, with the square commuting only up to \sim^+ , and see that the construction of $g \odot f$ proposed in [CCW19] readily extends. The only additional requirement is that maps between games behave well with respect to all three isomorphism families:

Definition 2.25. For games with symmetry \mathcal{A} and \mathcal{B} , a **map of games** $f : \mathcal{A} \rightarrow \mathcal{B}$ is a map of essps which additionally preserves positive and negative symmetry: if $\theta \in \cong_A^+$ then $f\theta \in \cong_B^+$ and if $\theta \in \cong_A^-$ then $f\theta \in \cong_B^-$.

Generalised maps of strategies are not given a central place in [CCRW17], yet they are used to show a significant result about “lifting”, a method for constructing strategies from maps of essps between the games themselves. We start by recalling the definition:

Definition 2.26. Let \mathcal{A} and \mathcal{B} be games. Let $f : \mathcal{A} \rightarrow \mathcal{B}$ be a map of event structures with symmetry which is courteous, receptive, and \sim -receptive. Then the composite

$$\mathbb{C}_{\mathcal{A}} \xrightarrow{\omega_{\mathcal{A}}} \mathcal{A}^\perp \parallel \mathcal{A} \xrightarrow{\mathcal{A}^\perp \parallel f} \mathcal{A}^\perp \parallel \mathcal{B}$$

is a strategy from \mathcal{A} to \mathcal{B} , called the **lifting** of f and denoted \hat{f} . Similarly for a courteous, receptive and \sim -receptive map $g : \mathcal{B}^\perp \rightarrow \mathcal{A}^\perp$, its **co-lifting** \check{g} is the strategy

$$\mathbb{C}_{\mathcal{B}} \xrightarrow{\omega_{\mathcal{B}}} \mathcal{B}^\perp \parallel \mathcal{B} \xrightarrow{g \parallel \mathcal{B}} \mathcal{A}^\perp \parallel \mathcal{B}$$

from \mathcal{A} to \mathcal{B} .

This construction is helpful because maps of essps are often easier to describe than strategies. For instance, the symmetry of parallel composition is evident from the isomorphism of essps

$$b : \mathcal{A} \parallel \mathcal{B} \rightarrow \mathcal{B} \parallel \mathcal{A},$$

so when \mathcal{A} and \mathcal{B} are games this can be lifted to a strategy

$$\hat{b} : \mathbb{C}_{\mathcal{A} \parallel \mathcal{B}} \rightarrow (\mathcal{A} \parallel \mathcal{B})^\perp \parallel \mathcal{B} \parallel \mathcal{A}.$$

In the bicategory of games and strategies, \widehat{b} is not strictly an isomorphism, but there are invertible 2-cells

$$\alpha_{\mathcal{A}} \cong \widehat{b}^{-1} \odot \widehat{b} \quad \widehat{b} \odot \widehat{b}^{-1} \cong \alpha_{\mathcal{B}}$$

making \widehat{b} an *equivalence*. To see why those 2-cells exist, we must investigate composition with a lifted map. This is where generalised maps come in.

Lemma 2.27 ([CCW19]). *Let $f : \mathcal{B} \rightarrow \mathcal{C}$ be courteous, receptive and \sim -receptive, and let $\sigma : \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$ be a strategy. Then, there is an invertible 2-cell $\widehat{f} \odot \sigma \cong (\mathcal{A}^\perp \parallel f) \circ \sigma$, i.e.*

$$\begin{array}{ccc} \mathbb{C}_{\mathcal{B}} \odot \mathcal{S} & \xrightarrow{\cong} & \mathcal{S} \\ & \searrow \widehat{f} \odot \sigma & \downarrow \sigma \\ & & \mathcal{A}^\perp \parallel \mathcal{B} \\ & & \downarrow \mathcal{A}^\perp \parallel f \\ & & \mathcal{A}^\perp \parallel \mathcal{C} \end{array}$$

Proof. There are generalised maps

$$\begin{array}{ccc} \mathcal{S} & \xrightarrow{\text{id}_{\mathcal{S}}} & \mathcal{S} \\ \sigma \downarrow & & \downarrow \sigma \\ \mathcal{A}^\perp \parallel \mathcal{B} & \xrightarrow{\text{id}_{\mathcal{A}^\perp} \parallel \text{id}_{\mathcal{B}}} & \mathcal{A}^\perp \parallel \mathcal{B} \end{array} \quad \begin{array}{ccc} \mathbb{C}_{\mathcal{B}} & \xrightarrow{\text{id}_{\mathbb{C}_{\mathcal{B}}}} & \mathbb{C}_{\mathcal{B}} \\ \alpha_{\mathcal{B}} \downarrow & & \downarrow \widehat{f} \\ \mathcal{B}^\perp \parallel \mathcal{B} & \xrightarrow{\mathcal{B}^\perp \parallel f} & \mathcal{B}^\perp \parallel \mathcal{C} \end{array}$$

which we can compose “horizontally”:

$$\begin{array}{ccc} \mathbb{C}_{\mathcal{B}} \odot \mathcal{S} & \xrightarrow{\text{id}_{\mathbb{C}_{\mathcal{B}}} \odot \text{id}_{\mathcal{S}}} & \mathbb{C}_{\mathcal{B}} \odot \mathcal{S} \\ \alpha_{\mathcal{B}} \odot \sigma \downarrow & & \downarrow \widehat{f} \odot \sigma \\ \mathcal{A}^\perp \parallel \mathcal{B} & \xrightarrow{\mathcal{A}^\perp \parallel f} & \mathcal{A}^\perp \parallel \mathcal{C} \end{array}$$

This map is an isomorphism, which we combine with the unitor

$$\begin{array}{ccc} \mathcal{S} & \xrightarrow{\lambda^{-1}} & \mathbb{C}_{\mathcal{B}} \odot \mathcal{S} \\ & \searrow \sigma & \downarrow \alpha_{\mathcal{B}} \odot \sigma \\ & & \mathcal{A}^\perp \parallel \mathcal{B} \end{array}$$

to get the result. \square

A pseudo-double category. By considering generalised maps of strategies, we step outside of the bicategory: the maps are now between strategies on different games. As it turns out, games, strategies, maps of games, and (generalised) maps of strategies form a *pseudo-double category*. In a double category there are two kinds of morphisms (referred to as *horizontal* and *vertical*), and therefore two kinds of composition, for which the laws of identity and associativity are required to hold

on the nose. In a *pseudo*-double category, the laws for horizontal composition are relaxed and allowed to hold up to coherent invertible 2-cells, much like in a bicategory. And indeed any pseudo-double category gives rise to a bicategory (its “horizontal bicategory”), obtained by forgetting the vertical dimension. Note that pseudo-double categories have recently appeared in several places within game semantics: see for instance [EH19, Mel19].

The formal definition is as follows. (The notation reflects that our pseudo-double category of interest consists of games and strategies.)

Definition 2.28. A **pseudo-double category** \mathcal{D} consists of two categories \mathbb{D}_0 and \mathbb{D}_1 , together with functors

$$\begin{aligned} \mathfrak{c} &: \mathbb{D}_0 \rightarrow \mathbb{D}_1 \\ \text{src}, \text{tgt} &: \mathbb{D}_1 \rightarrow \mathbb{D}_0 \\ \odot &: \mathbb{D}_1 \times_{\mathbb{D}_0} \mathbb{D}_1 \rightarrow \mathbb{D}_1 \end{aligned}$$

(where the pullback is over $\mathbb{D}_1 \xrightarrow{\text{src}} \mathbb{D}_0 \xleftarrow{\text{tgt}} \mathbb{D}_1$) such that for $\mathcal{A} \in \mathbb{D}_0$, $\text{src}(\mathfrak{c}_A) = \text{tgt}(\mathfrak{c}_A) = A$, and for $(\tau, \sigma) \in \mathbb{D}_1 \times_{\mathbb{D}_0} \mathbb{D}_1$, $\text{src}(\tau \odot \sigma) = \text{src}(\sigma)$ and $\text{tgt}(\tau \odot \sigma) = \text{tgt}(\tau)$. Additionally \mathcal{D} is equipped with natural isomorphisms

$$\begin{aligned} \lambda &: \mathfrak{c}_B \odot \sigma \Rightarrow \sigma \\ \rho &: \sigma \odot \mathfrak{c}_A \Rightarrow \sigma \\ \alpha &: (\eta \odot \tau) \odot \sigma \Rightarrow \eta \odot (\tau \odot \sigma) \end{aligned}$$

subject to the same coherence axioms as in a bicategory, and such that $\text{src}(\lambda)$, $\text{tgt}(\lambda)$, $\text{src}(\rho)$, $\text{tgt}(\rho)$, $\text{src}(\alpha)$, and $\text{tgt}(\alpha)$ are identity maps in \mathbb{D}_0 .

Some standard terminology and notation: morphisms of \mathbb{D}_0 are called **vertical morphisms**, and we write them $f : \mathcal{A} \rightarrow \mathcal{B}$. Objects of \mathbb{D}_1 are called **horizontal morphisms** and written $\sigma : \mathcal{A} \dashv\vdash \mathcal{B}$ (where $\mathcal{A} = \text{src}(\sigma)$ and $\mathcal{B} = \text{tgt}(\sigma)$; this matches the usual concurrent games notation). Finally a morphism α in \mathbb{D}_1 (known as a **2-cell**) is pictured as a square

$$\begin{array}{ccc} A & \xrightarrow{\sigma} & B \\ f \downarrow & \Downarrow \alpha & \downarrow g \\ A' & \xrightarrow{\sigma'} & B' \end{array}$$

where $f = \text{src}(\alpha)$ and $g = \text{tgt}(\alpha)$. Squares can be glued either horizontally or vertically, following one of the two notions of composition for 2-cells.

We will see that in the pseudo-double category \mathcal{G} of concurrent games, \mathbb{G}_0 has games as objects and maps of games as morphisms, while \mathbb{G}_1 has strategies as objects and generalised maps as morphisms. Later on it will be important to recover the usual bicategory, and for this we will perform the following construction on \mathcal{G} :

Definition 2.29. The **horizontal bicategory** $\mathcal{H}(\mathcal{D})$ of a pseudo-double category \mathcal{D} has

- objects: objects of \mathbb{D}_0 ;

- morphisms: objects of \mathbb{D}_1 ;
- 2-cells: **globular** morphisms of \mathbb{D}_1 , *i.e.* those α such that $\text{src}(\alpha)$ and $\text{tgt}(\alpha)$ are identity maps in \mathbb{D}_0 . Instead of the square

$$\begin{array}{ccc} A & \xrightarrow{\sigma} & B \\ \parallel & \Downarrow \alpha & \parallel \\ A & \xrightarrow{\sigma'} & B \end{array}$$

we often write a globular 2-cell as $\alpha : \sigma \Rightarrow \sigma'$.

In our pseudo-double category, a globular map will be one of the form

$$\begin{array}{ccc} \mathcal{S} & \xrightarrow{f} & \mathcal{T} \\ \sigma \downarrow & \sim^+ & \downarrow \tau \\ \mathcal{A}^\perp \parallel \mathcal{B} & \xrightarrow{\text{id}_\mathcal{A}^\perp \parallel \text{id}_\mathcal{B}} & \mathcal{A}^\perp \parallel \mathcal{B} \end{array}$$

i.e. a map in the usual sense.

Outline. The construction of a pseudo-double category of concurrent games and strategies is very natural, and it seems important enough in its own right. Furthermore, we will see that the lifting and co-lifting constructions are justified by the (established) double-categorical concepts of *companions* and *conjoins*.

Still, the result has practical relevance, since it allows us to identify symmetric monoidal structure in the bicategory of concurrent games. A theorem of Shulman [Shu10] states that if \mathcal{D} is symmetric monoidal, and is *isofibrant* (in concurrent games terms, “all isomorphisms can be lifted”), then the bicategory $\mathcal{H}(\mathcal{D})$ is symmetric monoidal. It is *significantly* easier to show that a pseudo-double category is symmetric monoidal than to carry out the full proof directly in the horizontal bicategory. As we will see in Section 2.6, symmetric monoidal structure is an important step towards cartesian closure, and the latter is needed for applications in semantics.

So the rest of the section is as follows: in 2.5.1 we define the pseudo-double category \mathcal{G} of concurrent games, discussing in particular the difficulty with horizontal composition in the presence of symmetry. In 2.5.2, we explore lifting in more detail, and in 2.5.3 we describe the symmetric monoidal structure of \mathcal{G} .

2.5.1 The pseudo-double category \mathcal{G}

We start with two categories.

- \mathbb{G}_0 has games \mathcal{A}, \mathcal{B} , *etc.* as objects, and maps $f : \mathcal{A} \rightarrow \mathcal{B}$ of games as morphisms, with the usual notions of identity and composition for maps of event structures.
- The objects of \mathbb{G}_1 are strategies $\sigma : \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$, $\tau : \mathcal{T} \rightarrow \mathcal{C}^\perp \parallel \mathcal{D} \dots$, where $\mathcal{A}, \mathcal{B}, \mathcal{C}, \mathcal{D}$, *etc.* are games. Its morphisms are tuples (f, g, h) with f a map of

essps and g, h maps of games, such that

$$\begin{array}{ccc} \mathcal{S} & \xrightarrow{f} & \mathcal{T} \\ \sigma \downarrow & \sim^+ & \downarrow \tau \\ \mathcal{A}^\perp \parallel \mathcal{B} & \xrightarrow{g^\perp \parallel h} & \mathcal{C}^\perp \parallel \mathcal{D} \end{array}$$

The identity morphism on a strategy $\sigma : \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$ is $(\text{id}_\mathcal{S}, \text{id}_\mathcal{A}, \text{id}_\mathcal{B})$, and the composition operation (*vertical* composition) is the componentwise composition, which is ensured to commute up to \sim^+ :

Lemma 2.30. *Let (f, g, h) and (f', g', h') be maps of strategies given by*

$$\begin{array}{ccccc} \mathcal{S} & \xrightarrow{f} & \mathcal{S}' & \xrightarrow{f'} & \mathcal{S}'' \\ \sigma \downarrow & \sim^+ & \downarrow \sigma' & \sim^+ & \downarrow \sigma'' \\ \mathcal{A}^\perp \parallel \mathcal{B} & \xrightarrow{g^\perp \parallel h} & \mathcal{A}'^\perp \parallel \mathcal{B}' & \xrightarrow{g'^\perp \parallel h'} & \mathcal{A}''^\perp \parallel \mathcal{B}'' \end{array}$$

Then $(f' \circ f, g' \circ g, h' \circ h)$ is a map of strategies from σ to σ'' .

Proof. We check that merging the two squares preserves the commutativity up to \sim^+ . For $x \in \mathcal{C}(\mathcal{S})$, by definition of \sim^+ there are two “canonical” symmetries

$$\varphi_x : \sigma'(fx) \cong (g^\perp \parallel h)(\sigma x) \quad \phi_{fx} : \sigma''(f'(fx)) \cong (g'^\perp \parallel h')(\sigma'(fx))$$

which are in $\cong_{\mathcal{A}'^\perp \parallel \mathcal{B}'}^+$ and $\cong_{\mathcal{A}''^\perp \parallel \mathcal{B}''}^+$, respectively. Since the map $g'^\perp \parallel h'$ preserves positive symmetry (as a map of games), the map

$$\phi_{fx} \circ (g'^\perp \parallel h') \varphi_x : \sigma''((f' \circ f)x) \cong ((g' \circ g)^\perp \parallel (h' \circ h))(\sigma x)$$

is also in $\cong_{\mathcal{A}''^\perp \parallel \mathcal{B}''}^+$, and it is clear that this coincides with the canonical bijection in the definition of \sim^+ , so we are done. \square

So \mathbb{G}_0 is our category of objects, and \mathbb{G}_1 our category of morphisms. The “domain” and “codomain” functors $\text{src}, \text{tgt} : \mathbb{G}_1 \rightarrow \mathbb{G}_0$ are as expected:

$$\begin{aligned} \text{src}(\mathcal{S} \xrightarrow{\sigma} \mathcal{A}^\perp \parallel \mathcal{B}) &= \mathcal{A} \\ \text{tgt}(\mathcal{S} \xrightarrow{\sigma} \mathcal{A}^\perp \parallel \mathcal{B}) &= \mathcal{B} \\ \text{src}(f, g, h) &= g \\ \text{tgt}(f, g, h) &= h \end{aligned}$$

For the “identity” functor $\mathfrak{c} : \mathbb{G}_0 \rightarrow \mathbb{G}_1$, we use the following property:

Lemma 2.31. *If $f : \mathcal{A} \rightarrow \mathcal{B}$ is a map of essps, then the map $f^\perp \parallel f : \mathcal{A}^\perp \parallel \mathcal{A} \rightarrow \mathcal{B}^\perp \parallel \mathcal{B}$ is also a map $\mathbb{C}_\mathcal{A} \rightarrow \mathbb{C}_\mathcal{B}$, denoted \mathfrak{c}_f .*

Proof. It preserves configurations since for every $x \parallel y \in \mathcal{C}(\mathbb{C}_\mathcal{A})$, $y \sqsubseteq x$, so $fy \sqsubseteq fx$ and therefore $fx \parallel fy = f(x \parallel y)$ is a configuration of $\mathbb{C}_\mathcal{B}$. It is locally injective because $f^\perp \parallel f$ is, and preserves symmetry, since it preserves \sqsubseteq and symmetries in copycat are of the form $\psi \parallel \phi$ with $\phi \sqsubseteq \psi$. \square

So we define $\mathfrak{c} : \mathbb{G}_0 \rightarrow \mathbb{G}_1$ as $\mathfrak{c}(\mathcal{A}) = \mathfrak{c}_{\mathcal{A}}$ and $\mathfrak{c}(f : \mathcal{A} \rightarrow \mathcal{B}) = (\mathfrak{c}_f, f, f)$.

We proceed with the horizontal composition functor \odot . The composition of strategies $\sigma : \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$ and $\tau : \mathcal{T} \rightarrow \mathcal{B}^\perp \parallel \mathcal{C}$ is the strategy $\tau \odot \sigma : \mathcal{T} \odot \mathcal{S} : \mathcal{A}^\perp \parallel \mathcal{C}$ obtained via interaction and hiding, as defined in 2.3.2. Defining the action of the functor \odot on morphisms of $\mathbb{G}_1 \times_{\mathbb{G}_0} \mathbb{G}_1$ is less straightforward. Suppose given maps

$$\begin{array}{ccc} \mathcal{S} & \xrightarrow{f} & \mathcal{S}' \\ \sigma \downarrow & \sim^+ & \downarrow \sigma' \\ \mathcal{A}^\perp \parallel \mathcal{B} & \xrightarrow{h_1^\perp \parallel h_2} & \mathcal{A}'^\perp \parallel \mathcal{B}' \end{array} \quad \begin{array}{ccc} \mathcal{T} & \xrightarrow{g} & \mathcal{T}' \\ \tau \downarrow & \sim^+ & \downarrow \tau' \\ \mathcal{B}^\perp \parallel \mathcal{C} & \xrightarrow{h_2^\perp \parallel h_3} & \mathcal{B}'^\perp \parallel \mathcal{C}' \end{array}$$

The situation for interactions can be summarised in a diagram, with polarity ignored as usual:

$$\begin{array}{ccccc} & & \mathcal{T} \circledast \mathcal{S} & & \\ & \swarrow \Pi_1 & & \searrow \Pi_2 & \\ & & \mathcal{T}' \circledast \mathcal{S}' & & \\ & \swarrow \Pi'_1 & & \searrow \Pi'_2 & \\ \mathcal{S} \parallel \mathcal{C} & \xrightarrow{f \parallel h_3} & \mathcal{S}' \parallel \mathcal{C}' & & \mathcal{A}' \parallel \mathcal{T}' \xleftarrow{h_1 \parallel g} \mathcal{A} \parallel \mathcal{T} \\ & \searrow \sigma \parallel \mathcal{C} & \sim & \swarrow \sigma' \parallel \mathcal{C}' & \swarrow \mathcal{A}' \parallel \tau' \\ & & \mathcal{A} \parallel \mathcal{B} \parallel \mathcal{C} & \xrightarrow{h_1 \parallel h_2 \parallel h_3} & \mathcal{A}' \parallel \mathcal{B}' \parallel \mathcal{C}' \xleftarrow{h_1 \parallel h_2 \parallel h_3} \mathcal{A} \parallel \mathcal{B} \parallel \mathcal{C} \\ & & & & \swarrow \mathcal{A} \parallel \tau \end{array}$$

The outer diagram, as well as the inner-most square, commute on the nose, because $\mathcal{T} \circledast \mathcal{S}$ and $\mathcal{T}' \circledast \mathcal{S}'$ are defined as pullbacks. But because the two squares at the bottom only commute up to symmetry, we are *not* in a position to apply the universal property of $\mathcal{T}' \circledast \mathcal{S}'$ in order to obtain a map $g \circledast f : \mathcal{T} \circledast \mathcal{S} \rightarrow \mathcal{T}' \circledast \mathcal{S}'$.

But it turns out that $\mathcal{T}' \circledast \mathcal{S}'$ enjoys a stronger universal property:

Lemma 2.32 (Bi-pullback property, [CCW19]). *If \mathcal{X} is an ess and there are maps $\varphi : \mathcal{X} \rightarrow \mathcal{S}' \parallel \mathcal{C}'$ and $\psi : \mathcal{X} \rightarrow \mathcal{A}' \parallel \mathcal{T}'$ such that $(\sigma' \parallel \mathcal{C}') \circ \varphi \sim (\mathcal{A}' \parallel \tau') \circ \psi$, then there exists a map $\omega : \mathcal{X} \rightarrow \mathcal{T} \circledast \mathcal{S}$, unique up to \sim , such that $\varphi \sim \Pi'_1 \circ \omega$ and $\psi \sim \Pi'_2 \circ \omega$.*

We can apply this to the above diagram. This gives a map $\omega : \mathcal{T} \circledast \mathcal{S} \rightarrow \mathcal{T}' \circledast \mathcal{S}'$, defined up to symmetry such that (simplifying the diagram)

$$\begin{array}{ccc} \mathcal{T} \circledast \mathcal{S} & \xrightarrow{\omega} & \mathcal{T}' \circledast \mathcal{S}' \\ \tau \circledast \sigma \downarrow & \sim & \downarrow \tau' \circledast \sigma' \\ \mathcal{A} \parallel \mathcal{B} \parallel \mathcal{C} & \xrightarrow{h_1 \parallel h_2 \parallel h_3} & \mathcal{A}' \parallel \mathcal{B}' \parallel \mathcal{C}' \end{array}$$

By restricting to visible events, we obtain a map $\omega|_{T \odot S}$ such that

$$\begin{array}{ccc} \mathcal{T} \odot \mathcal{S} & \xrightarrow{\omega|_{T \odot S}} & \mathcal{T}' \odot \mathcal{S}' \\ \tau \odot \sigma \downarrow & \sim & \downarrow \tau' \odot \sigma' \\ \mathcal{A} \parallel \mathcal{C} & \xrightarrow{h_1 \parallel h_3} & \mathcal{A}' \parallel \mathcal{C}' \end{array}$$

But the map $\omega|_{T \odot S}$ is only defined up to \sim , and it is not necessarily a map of strategies: for this the diagram above must commute up to \sim^+ with the projection to the game (and $\sim^+ \subseteq \sim$). The next lemma, found in [CCW19], implies that there is a unique 2-cell in the \sim -equivalence class of $\omega|_{T \odot S}$.

Lemma 2.33 ([CCW19]). *Let $\sigma : \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$ and $\sigma' : \mathcal{S}' \rightarrow \mathcal{A}'^\perp \parallel \mathcal{B}'$ be strategies, and let $g : \mathcal{A} \rightarrow \mathcal{A}'$ and $h : \mathcal{B} \rightarrow \mathcal{B}'$ be maps of games. Then, for every $f : \mathcal{S} \rightarrow \mathcal{S}'$ such that $\sigma' \circ f \sim (g^\perp \parallel h) \circ \sigma$, there is a unique $f' : \mathcal{S} \rightarrow \mathcal{S}'$ such that $f \sim f'$ and $\sigma' \circ f' \sim (g^\perp \parallel h) \circ \sigma$.*

(The statement of the lemma in [CCW19] mentions only globular maps of strategies. But in fact the proof is valid also for generalised maps.)

So we define $g \odot f$ to be the unique map $\mathcal{T} \odot \mathcal{S} \rightarrow \mathcal{T}' \odot \mathcal{S}'$ in the \sim -equivalence class of $\omega|_{T \odot S}$.

It is often convenient to reason directly at the level of interactions, so we note that $g \odot f$ has a unique *witness* $g \circledast f : \mathcal{T} \circledast \mathcal{S} \rightarrow \mathcal{T}' \circledast \mathcal{S}'$. From the results in [CCW19] we can derive:

Lemma 2.34. *Suppose f and g are maps of strategies as above. Then there is a unique map $g \circledast f : \mathcal{T} \circledast \mathcal{S} \rightarrow \mathcal{T}' \circledast \mathcal{S}'$ such that the diagram*

$$\begin{array}{ccc} \mathcal{T} \circledast \mathcal{S} & \xrightarrow{g \circledast f} & \mathcal{T}' \circledast \mathcal{S}' \\ \tau \circledast \sigma \downarrow & & \downarrow \tau' \circledast \sigma' \\ \mathcal{A} \parallel \mathcal{B} \parallel \mathcal{C} & \xrightarrow{h_1 \parallel h_2 \parallel h_3} & \mathcal{A}' \parallel \mathcal{B}' \parallel \mathcal{C}' \end{array}$$

commutes up to the isomorphism family $\cong_{A^\perp}^+ \parallel \cong_B \parallel \cong_C^+$ on $\mathcal{A}' \parallel \mathcal{B}' \parallel \mathcal{C}'$.

Proof. Suppose there are two such maps ω, ω' . By Lemma 2.32, $\omega \sim \omega'$, and so by Lemma 2.33, ω and ω' must agree on A -moves and C -moves. We show by induction on x that $\omega x = \omega' x$ for every $x \in \mathcal{C}(T \circledast S)$. For $x = \emptyset$ this is clear, and by the previous remark we only need to consider extensions $x \dashv^e y$ for e a B -move. Without loss of generality, suppose e is a σ -action. As $\omega \sim \omega'$, we have a bijection $\varphi : \omega y \cong_{T' \circledast S'} \omega' y$. Its projection to $\cong_{S'}$ is a positive extension of the identity bijection on the projection to S of $\omega x (= \omega' x)$, therefore by thinness of \mathcal{S}' it must itself be the identity. \square

So in particular $g \circledast f$ restricts to $g \odot f$.

Theorem 2.35. *There is a pseudo-double category \mathcal{G} with*

- *objects: games with symmetry;*
- *vertical morphisms: maps of games;*
- *horizontal morphisms: strategies; and*
- *2-cells: “generalised” maps of strategies.*

Proof. We have given all the data. It remains to check (1) naturality of $\lambda_\sigma, \rho_\sigma$, and α_σ , (2) that the latter satisfy the necessary coherence axioms, and (3) that horizontal composition is functorial (also known as the “interchange” law). \square

2.5.2 Lifting

We consider the lifting construction (Definition 2.26) in the context of double category theory.

For \mathcal{D} a (pseudo-)double category, the following are well-known concepts [GP04]:

Definition 2.36. Let $f : \mathcal{A} \rightarrow \mathcal{B}$ be a vertical morphism, *i.e.* a morphism in \mathbb{D}_0 . A **companion** to f is a horizontal morphism $\hat{f} : \mathcal{A} \rightarrow \mathcal{B}$ together with 2-cells

$$\begin{array}{ccc} \mathcal{A} & \xrightarrow{\hat{f}} & \mathcal{B} \\ f \downarrow & \Downarrow & \parallel \\ \mathcal{B} & \xrightarrow[\alpha_{\mathcal{B}}]{} & \mathcal{B} \end{array} \quad \begin{array}{ccc} \mathcal{A} & \xrightarrow[\alpha_{\mathcal{A}}]{} & \mathcal{A} \\ \parallel & \Downarrow & \downarrow f \\ \mathcal{A} & \xrightarrow[\hat{f}]{} & \mathcal{B} \end{array}$$

subject to coherence axioms [Shu10]. A **conjoint** of f is a horizontal morphism $\check{f} : \mathcal{B} \rightarrow \mathcal{A}$ which is a companion of f in the dual double category (the explicit data is easily recovered).

The lifting construction for maps of essps is an instance of the above:

Lemma 2.37. *If $f : \mathcal{A} \rightarrow \mathcal{B}$ is a courteous, receptive and \sim -receptive map of essps, then the strategy \hat{f} of Definition 2.26 is a companion to f in \mathcal{G} .*

If $f^\perp : \mathcal{A}^\perp \rightarrow \mathcal{B}^\perp$ is a courteous, receptive and \sim -receptive map of essps, then the strategy \check{f} is a conjoint to f in \mathcal{G} .

Proof. From the definition of \hat{f} it follows that the diagrams

$$\begin{array}{ccc} \mathbb{C}_{\mathcal{A}} & \xrightarrow{\alpha_f} & \mathbb{C}_{\mathcal{B}} \\ \hat{f} \downarrow & & \downarrow \alpha_{\mathcal{B}} \\ \mathcal{A}^\perp \parallel \mathcal{B} & \xrightarrow{f^\perp \parallel \mathcal{B}} & \mathcal{B}^\perp \parallel \mathcal{B} \end{array} \quad \begin{array}{ccc} \mathbb{C}_{\mathcal{A}} & \xrightarrow{\text{id}_{\mathbb{C}_{\mathcal{A}}}} & \mathbb{C}_{\mathcal{A}} \\ \alpha_{\mathcal{A}} \downarrow & & \downarrow \hat{f} \\ \mathcal{A}^\perp \parallel \mathcal{A} & \xrightarrow{\mathcal{A}^\perp \parallel f} & \mathcal{A}^\perp \parallel \mathcal{B} \end{array}$$

commute, and this provides the required data. The axioms can be verified, and the second part of the statement has a symmetric proof. \square

Say \mathcal{D} is **isofibrant** [Shu10] if every isomorphism has both a companion and a conjoint. This is easily verified in \mathcal{G} :

Lemma 2.38. \mathcal{G} is isofibrant.

Proof. A map $f : \mathcal{A} \rightarrow \mathcal{B}$ can be lifted provided it is courteous, receptive and \sim -receptive. It can be co-lifted if the map f^\perp is courteous, receptive and \sim -receptive. When f is an isomorphism, both of the above are immediately satisfied. \square

2.5.3 Symmetric monoidal structure

A **symmetric monoidal** pseudo-double category \mathcal{D} is one where:

- \mathbb{D}_0 and \mathbb{D}_1 are symmetric monoidal categories,
- the functor \mathfrak{c} preserves the monoidal unit,
- the functors src and tgt are *strict* symmetric monoidal,

and additionally \mathcal{D} is equipped with globular, invertible 2-cells

$$\begin{aligned} \phi_{\mathcal{A},\mathcal{B}} : \mathfrak{c}_{\mathcal{A} \otimes \mathcal{B}} &\Rightarrow \mathfrak{c}_{\mathcal{A}} \otimes \mathfrak{c}_{\mathcal{B}} \\ \phi_{(\sigma_1, \sigma_2), (\tau_1, \tau_2)} : (\tau_1 \parallel \tau_2) \odot (\sigma_1 \parallel \sigma_2) &\Rightarrow (\tau_1 \odot \sigma_1) \parallel (\tau_2 \odot \sigma_2) \end{aligned}$$

subject to coherence axioms which we omit [Shu10]. Those axioms are considerably less intimidating than those for a symmetric monoidal *bicategory* [Sta16], which makes the following result attractive:

Theorem 2.39 (Shulman [Shu10]). *If \mathcal{D} is an isofibrant symmetric monoidal pseudo-double category, then its horizontal bicategory $\mathcal{H}(\mathcal{D})$ is symmetric monoidal.*

The symmetric monoidal structure in \mathcal{G} is given by the parallel composition operation and the associated structural maps. Unsurprisingly, the next result will play a fundamental role:

Lemma 2.40. *The category **Essp** of event structures with symmetry and polarity is symmetric monoidal, with monoidal product \parallel and unit \emptyset .*

Proof. The parallel composition operation \parallel extends to a functor in a natural way, and there are isomorphisms

$$\begin{aligned} \underline{l}_{\mathcal{A}} : \emptyset \parallel \mathcal{A} &\rightarrow \mathcal{A} & \underline{a}_{\mathcal{A},\mathcal{B},\mathcal{C}} : (\mathcal{A} \parallel \mathcal{B}) \parallel \mathcal{C} &\rightarrow \mathcal{A} \parallel (\mathcal{B} \parallel \mathcal{C}) \\ \underline{r}_{\mathcal{A}} : \mathcal{A} \parallel \emptyset &\rightarrow \mathcal{A} & \underline{b}_{\mathcal{A},\mathcal{B}} : \mathcal{A} \parallel \mathcal{B} &\rightarrow \mathcal{B} \parallel \mathcal{A} \end{aligned}$$

There is a faithful functor $\mathbf{Essp} \rightarrow \mathbf{Set}$ sending the above to the structural data associated with the symmetric monoidal category $(\mathbf{Set}, +, \emptyset)$. From this we obtain that all the necessary properties are satisfied. \square

The category \mathbb{G}_0 has games as objects and maps of games as morphisms; the category inherits the monoidal structure of **Essp** in an obvious way.

The category \mathbb{G}_1 is also symmetric monoidal. For strategies $\sigma : \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$ and $\tau : \mathcal{T} \rightarrow \mathcal{C}^\perp \parallel \mathcal{D}$, we define $\sigma \otimes \tau$ to be the composite

$$\mathcal{S} \parallel \mathcal{T} \xrightarrow{\sigma \parallel \tau} (\mathcal{A}^\perp \parallel \mathcal{B}) \parallel (\mathcal{C}^\perp \parallel \mathcal{D}) \xrightarrow{\cong} (\mathcal{A} \parallel \mathcal{C})^\perp \parallel (\mathcal{B} \parallel \mathcal{D})$$

where the second arrow is the appropriate reordering. This is clearly a strategy. This construction extends to 2-cells: given

$$\begin{array}{ccc}
\mathcal{S} & \xrightarrow{f_1} & \mathcal{S}' \\
\sigma \downarrow & \sim^+ & \downarrow \sigma' \\
\mathcal{A}^\perp \parallel \mathcal{B} & \xrightarrow{g_1^\perp \parallel h_1} & \mathcal{A}'^\perp \parallel \mathcal{B}'
\end{array}
\qquad
\begin{array}{ccc}
\mathcal{T} & \xrightarrow{f_2} & \mathcal{T}' \\
\tau \downarrow & \sim^+ & \downarrow \tau' \\
\mathcal{C}^\perp \parallel \mathcal{D} & \xrightarrow{g_2^\perp \parallel h_2} & \mathcal{C}'^\perp \parallel \mathcal{D}'
\end{array}$$

we construct their monoidal product as

$$\begin{array}{ccc}
\mathcal{S} \parallel \mathcal{T} & \xrightarrow{f_1 \parallel f_2} & \mathcal{S}' \parallel \mathcal{T}' \\
\sigma \parallel \tau \downarrow & \sim^+ & \downarrow \sigma' \parallel \tau' \\
(\mathcal{A}^\perp \parallel \mathcal{B}) \parallel (\mathcal{C}^\perp \parallel \mathcal{D}) & \xrightarrow{(g_1^\perp \parallel h_1) \parallel (g_2^\perp \parallel h_2)} & (\mathcal{A}'^\perp \parallel \mathcal{B}') \parallel (\mathcal{C}'^\perp \parallel \mathcal{D}') \\
\cong \downarrow & & \downarrow \cong \\
(\mathcal{A}^\perp \parallel \mathcal{C}^\perp) \parallel (\mathcal{B} \parallel \mathcal{D}) & \xrightarrow{(g_1^\perp \parallel g_2^\perp) \parallel (h_1 \parallel h_2)} & (\mathcal{A}'^\perp \parallel \mathcal{C}'^\perp) \parallel (\mathcal{B}' \parallel \mathcal{D}')
\end{array}$$

where the top square commutes up to \sim^+ by definition of the positive symmetry in $(\mathcal{A}'^\perp \parallel \mathcal{B}') \parallel (\mathcal{C}'^\perp \parallel \mathcal{D}')$, and the bottom square commutes by naturality of the structural maps in **Essp**.

Two final pieces of data are required: for every \mathcal{A} and \mathcal{B} , an isomorphism

$$\phi_{\mathcal{A}, \mathcal{B}} : \mathbb{C}_{\mathcal{A} \parallel \mathcal{B}} \Rightarrow \mathbb{C}_{\mathcal{A}} \parallel \mathbb{C}_{\mathcal{B}}$$

and for every $\sigma_1 : \mathcal{A}_1 \rightarrow \mathcal{B}_1$, $\tau_1 : \mathcal{B}_1 \rightarrow \mathcal{C}_1$ and $\sigma_2 : \mathcal{A}_2 \rightarrow \mathcal{B}_2$, $\tau_2 : \mathcal{B}_2 \rightarrow \mathcal{C}_2$, an isomorphism

$$\phi_{(\sigma_1, \sigma_2), (\tau_1, \tau_2)} : (\tau_1 \parallel \tau_2) \odot (\sigma_1 \parallel \sigma_2) \Rightarrow (\tau_1 \odot \sigma_1) \parallel (\tau_2 \odot \sigma_2).$$

We give both maps by defining their action on configurations. For $(x_A \parallel x_B) \parallel (y_A \parallel y_B) \in \mathcal{C}(\mathbb{C}_{\mathcal{A} \parallel \mathcal{B}})$, $\phi_{\mathcal{A}, \mathcal{B}}((x_A \parallel x_B) \parallel (y_A \parallel y_B)) = (x_A \parallel y_A) \parallel (x_B \parallel y_B)$, and for $(y_1 \parallel y_2) \odot (x_1 \parallel x_2) \in \mathcal{C}((\mathcal{T}_1 \parallel \mathcal{T}_2) \odot (\mathcal{S}_1 \parallel \mathcal{S}_2))$, let $\phi_{(\sigma_1, \sigma_2), (\tau_1, \tau_2)}((y_1 \parallel y_2) \odot (x_1 \parallel x_2)) = y_1 \odot x_1 \parallel y_2 \odot x_2$. This is easily shown to preserve symmetry, polarised inclusions, cardinality, and unions, so that by Lemma 2.6, $\phi_{(\sigma_1, \sigma_2), (\tau_1, \tau_2)}$ and $\phi_{\mathcal{A}, \mathcal{B}}$ are generated by maps of essps.

Theorem 2.41. *The pseudo-double category \mathcal{G} is symmetric monoidal.*

Proof. The axioms can be verified directly without any difficulty. \square

Using Shulman's theorem, we immediately deduce:

Corollary 2.42. *The bicategory $\mathcal{H}(\mathcal{G})$ is symmetric monoidal.*

Rest of the chapter. Symmetric monoidal structure is not sufficient for our purposes. So in the next section, we will impose some restrictions on the bicategory $\mathcal{H}(\mathcal{G})$, to get a sub-bicategory **G** of negative arenas and so-called *negative, well-threaded* strategies. We will show that:

- \mathbf{G} is symmetric monoidal closed;
- \mathbf{G} has finite products; and
- there is a pseudo-comonad $! : \mathbf{G} \rightarrow \mathbf{G}$ such that the Kleisli bicategory $\mathbf{G}_!$ is cartesian closed.

We will give the necessary definitions along the way.

2.6 A cartesian closed bicategory

To interpret the higher-order languages we are concerned with in this thesis, we construct a bicategory with more structure.

Our first step is to restrict the objects of the model to the *negative arenas* of 2.2.2. Then we restrict the morphisms:

Definition 2.43. A strategy $\sigma : \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$ is **negative** if all initial moves of S are negative. It is **well-threaded** if for every $s \in S$, $[s]$ contains a unique initial move, denoted $\text{init}(s)$.

This is a valid restriction:

Lemma 2.44. *Negative and well-threaded strategies between negative arenas are stable under composition.*

Proof. Let $\sigma : \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$ and $\tau : \mathcal{T} \rightarrow \mathcal{B}^\perp \parallel \mathcal{C}$ be negative and well-threaded, with $\mathcal{A}, \mathcal{B}, \mathcal{C}$ negative arenas.

Maps of event structures preserve initial moves, so if $e \in S$ is initial, then it σe is initial in $\mathcal{A}^\perp \parallel \mathcal{B}$, and therefore because e is negative it is necessarily a B -move. Similarly, initial moves of T are all C -moves.

If $e \in T \otimes S$ is initial, then $\Pi_1 e$ is initial in $S \parallel C$ and $\Pi_2 e$ is initial in $A \parallel T$, so by the above remarks e must be a C -move; in particular it is visible and negative. This shows that $\tau \odot \sigma$ is negative.

Let $d \in T \otimes S$ be visible and suppose $e, e' \leq d$ with e and e' initial. If $\Pi_2 d$ is in the T component of $A \parallel T$, then we must have $\Pi_2 e, \Pi_2 e' \leq \Pi_2 d$, which implies $e = e'$ because τ is well-threaded.

If d is an A -move, then there must be B -moves $c, c' \in T \otimes S$ such that $\Pi_1 c, \Pi_1 c' \leq \Pi_1 d$, $\Pi_2 e \leq \Pi_2 c$ and $\Pi_2 e' \leq \Pi_2 c'$. The moves c and c' can be chosen so that $\Pi_1 c$ and $\Pi_1 c'$ are minimal in S , in which case we must have $c = c'$ by well-threadedness of S . Therefore $e = e'$ as τ is well-threaded. \square

It is not difficult to see why negative and well-threaded strategies are closed under \otimes , and the copycat strategy on a negative arena is negative and well-threaded. In particular, any strategy obtained by lifting an isomorphism of negative arenas is negative and well-threaded.

Hence, by restricting the strategies and arenas of $\mathcal{H}(\mathcal{G})$ in this way, we obtain a symmetric monoidal bicategory, denoted \mathbf{G} , with

- objects: negative arenas,

- morphisms: well-threaded, negative strategies,
- 2-cells: (globular) maps of strategies,

and with symmetric monoidal structure inherited from $\mathcal{H}(\mathcal{G})$.

At this stage it is helpful to recall some notions of bicategory theory. First, for bicategories \mathbb{C} and \mathbb{D} , a **pseudo-functor** from \mathbb{C} to \mathbb{D} consists of:

- a map $F : |\mathbb{C}| \rightarrow |\mathbb{D}|$;
- for each $A, B \in \mathbb{C}$, functors $F_{A,B} : \mathbb{C}[A, B] \rightarrow \mathbb{D}[FA, FB]$, often written F ;
- for each $A \in \mathbb{C}$, an invertible 2-cell $\Phi_A : \text{id}_{FA} \Rightarrow F\text{id}_A$;
- for each $A, B, C \in \mathbb{C}$, $\sigma \in \mathbb{C}[A, B]$ and $\tau \in \mathbb{C}[B, C]$, an invertible 2-cell $\Phi_{\sigma, \tau} : F\tau \odot F\sigma \Rightarrow F(\tau \odot \sigma)$,

subject to coherence axioms. We will write (F, Φ) , or simply F for a pseudo-functor with data as above.

For pseudo-functors $F, G : \mathbb{C} \rightarrow \mathbb{D}$, a **pseudo-natural transformation** $a : F \rightarrow G$ consists of maps $a_A \in \mathbb{C}[FA, GA]$ for each $A \in \mathbb{C}$, together with for each $\sigma \in \mathbb{C}[A, B]$, an invertible 2-cell a_σ filling the usual “naturality square”, *i.e.* $a_\sigma : a_B \odot F\sigma \Rightarrow G\sigma \odot a_A$, subject to more coherence axioms. Given two pseudo-natural transformations a, b between pseudo-functors F and G , a **modification** $\mu : a \Rightarrow b$ is a family of 2-cells $\mu_A : a_A \Rightarrow b_A$ which commutes (in the appropriate sense) with the 2-cells a_σ and b_σ .

An **adjunction** in a bicategory \mathbb{C} consists of two objects $A, B \in \mathbb{C}$, morphisms $\sigma \in \mathbb{C}[A, B]$ and $\sigma^\bullet \in \mathbb{C}[B, A]$ and 2-cells $\eta : \text{id}_A \Rightarrow \sigma^\bullet \odot \sigma$ and $\varepsilon : \sigma \odot \sigma^\bullet \Rightarrow \text{id}_B$ satisfying the usual “triangle identities” [ML13]. (In this situation σ is *left adjoint* to σ^\bullet ; we write $\sigma \dashv \sigma^\bullet$.) If η and ε are invertible, we call $(A, B, \sigma, \sigma^\bullet, \eta, \varepsilon)$ an **adjoint equivalence**.

An **equivalence** is weaker than an adjoint equivalence in the sense that η and ε need not satisfy the triangle identities. But conveniently, every equivalence gives rise to an adjoint equivalence between the same objects [ML13], so that whenever an adjoint equivalence is required it is sufficient to exhibit an equivalence. We will make much use of this.

A few more concepts will be introduced along the way as necessary. In doing certain proofs, the **coherence theorem** for bicategories is an important tool: it says that any diagram of 2-cells made up of instances of α , λ and ρ must commute.

2.6.1 Closed structure

We now show that \mathbf{G} is symmetric monoidal closed. The formal definition is as follows:

Definition 2.45. A symmetric monoidal bicategory \mathbb{C} is **closed** if for every object B , the pseudo-functor $- \otimes B$ has a right biadjoint, denoted $B \multimap -$.

Unfolding the definition, to prove \mathbb{C} is closed it suffices to give, for all objects A and B ,

- an object $A \multimap B$;
- a morphism $\text{ev}_{A,B} : (A \multimap B) \otimes A \rightarrow B$
- an adjoint equivalence

$$\begin{array}{ccc}
& \xrightarrow{\text{ev}_{B,C} \circ (- \otimes B)} & \\
\mathbb{C}[A, B \multimap C] & \perp & \mathbb{C}[A \otimes B, C] \\
& \xleftarrow{\Lambda} &
\end{array}$$

for every $C \in \mathbb{C}$.

We show \mathbf{G} is closed.

Given (negative and well-threaded) games A and B (without symmetry for now), the game $A \multimap B$ is defined to be the event structure with events $\bigsqcup_{b \in \min(B)} A^\perp \uplus B$ (where $\min(B)$ is the set of initial events of B), polarity induced from A and B , and causality relation

$$\leq (\|_{b \in \min(B)} A) \| B \cup \{((2, b), (1, (a, b))) \mid b \in \min(B), a \in A^\perp\}.$$

Observe that there is a canonical function $\chi : A \multimap B \rightarrow A^\perp \| B$, sending B to itself and mapping every copy of A^\perp in $A \multimap B$ to the A^\perp component of $A^\perp \| B$. Thus the function χ reflects order and preserves polarity, and the consistency relation on $A \multimap B$ is defined so as to turn χ into a valid map of event structures. Given $X \subseteq A \multimap B$, $X \in \text{Con}_{A \multimap B}$ iff the restriction $\chi|_{[X]}$ is injective and $\chi X \in \text{Con}_{A^\perp \| B}$.

When \mathcal{A} and \mathcal{B} are games with symmetry, we define three families of bijections on $A \multimap B$:

- The set $\cong_{A \multimap B}$ contains order-isomorphisms $\theta : x \cong y$ such that $\chi\theta : \chi x \cong \chi y$ is in $\cong_A \| \cong_B$.
- The set $\cong_{A \multimap B}^-$ contains order-isomorphisms $\theta : x \cong y$ such that $\chi\theta : \chi x \cong \chi y$ is in $\cong_A^+ \| \cong_B^-$.
- The set $\cong_{A \multimap B}^+$ contains order-isomorphisms $\theta : x \cong y$ such that $\chi\theta : \chi x \cong \chi y$ is in $\cong_A^- \| \cong_B^+$.

The following technical lemma will be useful as we instantiate the lifting construction with χ later on:

Lemma 2.46. *The map $\chi : \mathcal{A} \multimap \mathcal{B} \rightarrow \mathcal{A}^\perp \| \mathcal{B}$ is a courteous, receptive and \sim -receptive map of essps.*

Proof. Additional causal links in $\mathcal{A} \multimap \mathcal{B}$ are from initial moves of B to initial moves of A . Because \mathcal{A} and \mathcal{B} are negative arenas, this respects courtesy.

We check receptivity. Let $x \in \mathcal{C}(A \multimap B)$ and $\chi x \multimap^e y$ for some negative move e . If e is a B -move, then $x \cup \{e\} \in \mathcal{C}(A \multimap B)$ and we are done. If e is an A -move, then because e is negative it cannot be initial in A . Let a be its predecessor in A ; a

must be in χx , so $a = \chi e'$ for some (necessarily unique) $e' \in x$. There is a unique $e'' \in A \multimap B$ such that $e' \rightarrow e''$ and $\chi e'' = e$, so $x \cup \{e''\}$ is the witness to receptivity.

Finally let $\theta : x \cong_{A \multimap B} y$ and suppose $\chi \theta \multimap^{(e, e')} \phi : z \cong_{A^\perp \parallel B} w$ is a negative extension. By receptivity there are unique x' and y' in $\mathcal{C}(A \multimap B)$ extending x and y , and the (clearly unique) bijection $\theta' : x' \cong y'$ extending θ such that $\chi \theta' = \phi$ is an order isomorphism, by courtesy of χ . So $\theta' \in \cong_{A \multimap B}$ which concludes the proof. \square

Lemma 2.47. *The tuple $\mathcal{A} \multimap \mathcal{B} = (A \multimap B, \cong_{A \multimap B}, \cong_{A \multimap B}^-, \cong_{A \multimap B}^+)$ is a negative arena.*

Proof. We check that $\cong_{A \multimap B}$ is an isomorphism family. The map χ preserves identities, composition and inverses of bijections, so the (Groupoid) axiom is satisfied since $\cong_A \parallel \cong_B$ is an isomorphism family.

If $\theta : x \cong_{A \multimap B} y$, and $x \subseteq x'$, then $\chi x \subseteq \chi x'$ and therefore by the (Extension) axiom for $\cong_{A^\perp \parallel B}$ there exists $\varphi : \chi x' \cong_{\cong_{A^\perp \parallel B}} z$ such that $\chi y \subseteq z$ and $\chi \theta \subseteq \varphi$. From the definition of $A \multimap B$ we see that there is at most one $b \in B$ such that y contains elements of the form $(1, (b, a))$. If it exists, we can define $y' = y \cup \{(1, (b, a)) : a \in z \setminus \chi y\}$. Therefore $\chi y' = z$, and in particular ϕ is the image under χ of some bijection $\theta' : x' \cong_{A \multimap B} y'$ such that $\theta \subseteq \theta'$. So the (Extension) axiom holds for $\cong_{A \multimap B}$.

Suppose now $\theta : x \cong_{A \multimap B} y$, and $x' \subseteq x$. Then, there is a bijection $\phi : \chi x' \cong_{\cong_{A^\perp \parallel B}} z$ such that $\phi \subseteq \chi \theta$. But since $z \subseteq \chi y$ and χ is a map of event structures there exists $y' \subseteq y$ such that $\chi y' = z$, and in particular ϕ is the image under χ of some $\theta' : x' \cong_{A \multimap B} y'$, with $\theta' \subseteq \theta$. So the (Restriction) property is satisfied and $\cong_{A \multimap B}$ is an isomorphism family. Reproducing this argument for $\cong_{A \multimap B}^-$ and $\cong_{A \multimap B}^+$ is straightforward.

We now check the axioms from Definition 2.12. That $\cong_{A \multimap B}^-$ and $\cong_{A \multimap B}^+$ are subsets of $\cong_{A \multimap B}$ follows from the fact that $\cong_A^- \parallel \cong_B^+$ and $\cong_A^+ \parallel \cong_B^-$ are subsets of $\cong_A \parallel \cong_B$. If $\theta : x \cong y$ is in $\cong_{A \multimap B}^+ \cap \cong_{A \multimap B}^-$ then $\chi \theta \in \cong_A^- \parallel \cong_B^+ \cap \cong_A^+ \parallel \cong_B^-$, so $\chi \theta$ must be the identity and in particular $\chi x = \chi y$. Observe that this implies that θ is the identity bijection on the restriction of x and y to the B component of $A \multimap B$. Suppose there exists $e \in x$ such that $\theta e \neq e$. Then e and θe must be in different copies of A^\perp , so by definition $\text{init}(e) \neq \text{init}(\theta e)$. But because θ is an order-isomorphism, we have $\text{init}(\theta e) = \theta(\text{init}(e))$ and so this implies $\text{init}(e) \neq \theta(\text{init}(e))$. This is a contradiction, as $\text{init}(e)$ and $\theta(\text{init}(e))$ are both in the B component of $A \multimap B$ on which θ is the identity.

Finally, suppose θ is in $\cong_{A \multimap B}^-$ and $\theta \subseteq^- \theta'$ for some $\theta' \in \cong_{A \multimap B}$. Then we have that $\chi \theta \subseteq^- \chi \theta'$, where $\chi \theta \in \cong_{A^\perp \parallel B}^-$ and $\chi \theta' \in \cong_{A^\perp \parallel B}$. So $\chi \theta' \in \cong_{A^\perp \parallel B}^-$ and so $\theta' \in \cong_{A \multimap B}^-$. The last axiom is checked in the same way. \square

With this definition of $\mathcal{A} \multimap \mathcal{B}$, the evaluation strategy $\text{ev}_{\mathcal{A}, \mathcal{B}}$ is obtained by lifting the map $\chi : (\mathcal{A} \multimap \mathcal{B})^\perp \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$ and rearranging the bracketing of the game: the composite

$$\text{ev}_{\mathcal{A}, \mathcal{B}} : \mathbb{C}_{\mathcal{A} \multimap \mathcal{B}} \xrightarrow{\hat{\chi}} (\mathcal{A} \multimap \mathcal{B})^\perp \parallel (\mathcal{A}^\perp \parallel \mathcal{B}) \xrightarrow{\cong} ((\mathcal{A} \multimap \mathcal{B})^\perp \parallel \mathcal{A}^\perp) \parallel \mathcal{B}$$

is a strategy $(\mathcal{A} \multimap \mathcal{B}) \otimes \mathcal{A} \rightarrow \mathcal{B}$. To see that ev induces a family of adjoint equivalences between hom-categories, we will use the following result.

Lemma 2.48. *For arenas $\mathcal{A}, \mathcal{B}, \mathcal{C}$, there is an isomorphism of categories*

$$\Phi : \mathbf{G}[\mathcal{A}, \mathcal{B} \multimap \mathcal{C}] \rightarrow \mathbf{G}[\mathcal{A} \otimes \mathcal{B}, \mathcal{C}].$$

Proof. For $\sigma : \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B} \multimap \mathcal{C}$, define $\Phi(\sigma)$ to be the composite

$$\mathcal{S} \xrightarrow{\sigma} \mathcal{A}^\perp \parallel \mathcal{B} \multimap \mathcal{C} \xrightarrow{\mathcal{A}^\perp \parallel \chi} \mathcal{A}^\perp \parallel (\mathcal{B}^\perp \parallel \mathcal{C}) \xrightarrow{\cong} (\mathcal{A} \parallel \mathcal{B})^\perp \parallel \mathcal{C},$$

a strategy $\mathcal{A} \otimes \mathcal{B} \rightarrow \mathcal{C}$.

We show that Φ has an inverse. Given a strategy $\tau : \mathcal{T} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}^\perp \parallel \mathcal{C}$, we define a strategy $\tau' : \mathcal{T} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B} \multimap \mathcal{C}$ such that $\Phi(\tau') = \tau$. For $t \in T$, there are two cases: either $\tau(t)$ is an A -move a , or $\tau(t) = e$ for some $e \in \mathcal{B}^\perp \parallel \mathcal{C}$. If $\tau(t)$ is an A -move a , we set $\tau'(t) = \tau(t)$. Otherwise, if $\tau(t)$ is a $\mathcal{B}^\perp \parallel \mathcal{C}$ -move e , we need to set $\tau'(t) = e'$ for some $e' \in \mathcal{B} \multimap \mathcal{C}$ with $(\chi_{\mathcal{B}, \mathcal{C}})e' = e$. If $e = (2, c)$ for some $c \in \mathcal{C}$, then $e' = e$ will do. If $e = (1, b)$ for $b \in \mathcal{B}$, then we can set $e' = (1, (\text{init}(t), b))$.

If $x \in \mathcal{C}(T)$ and $\tau'x \multimap y \in \mathcal{C}(\mathcal{A}^\perp \parallel \mathcal{B} \multimap \mathcal{C})$, then $(\mathcal{A}^\perp \parallel \chi)(\tau'x) \multimap (\mathcal{A}^\perp \parallel \chi)y$ and so by receptivity of $\tau = (\mathcal{A}^\perp \parallel \chi) \circ \tau'$, there is a unique $x' \in \mathcal{C}(T)$ with $x \multimap x'$ and $\tau x' = (\mathcal{A}^\perp \parallel \chi)y$. Then $\tau'x' = y$, since they are both negative extensions of $\tau'x$ with the same image under $(\mathcal{A}^\perp \parallel \chi)$. So τ' is receptive, and the argument for \sim -receptivity is the same. Thinness is a property of \mathcal{T} , so τ' is thin because τ is.

We investigate the action of Φ on 2-cells. Given strategies $\sigma : \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B} \multimap \mathcal{C}$ and $\sigma' : \mathcal{S}' \rightarrow \mathcal{A}^\perp \parallel \mathcal{B} \multimap \mathcal{C}$, and a map $f : \sigma \rightarrow \sigma'$, define $\Phi(f)$ to be the same $f : \mathcal{S} \rightarrow \mathcal{S}'$: this is also a map $\Phi(\sigma) \rightarrow \Phi(\sigma')$ because χ preserves positive symmetry. Conversely if $f : \Phi(\sigma) \Rightarrow \Phi(\sigma')$, we show $f : \sigma \Rightarrow \sigma'$. Let $x \in \mathcal{S}$. By definition of maps of strategies, the bijection $\{((\mathcal{A}^\perp \parallel \chi)(\sigma(s)), (\mathcal{A}^\perp \parallel \chi)((\sigma' \circ f)(s))) \mid s \in x\}$ is in $\cong_{\mathcal{A}^\perp}^+ \parallel \cong_{\mathcal{B}^\perp \parallel \mathcal{C}}^+$, so $\{\sigma s, (\sigma' \circ f)s \mid s \in x\}$ meets the conditions for being in $\cong_{\mathcal{A}^\perp}^+ \parallel \cong_{\mathcal{B} \multimap \mathcal{C}}^+$, and we are done. \square

Lemma 2.49. *The bicategory \mathbf{G} is symmetric monoidal closed.*

Proof. For arenas \mathcal{A} and \mathcal{B} , we have defined $\mathcal{A} \multimap \mathcal{B}$; the strategy $\text{ev}_{\mathcal{A}, \mathcal{B}}$ is the image of $\mathbb{C}_{\mathcal{A} \multimap \mathcal{B}}$ under the functor Φ of Lemma 2.48, so

$$\text{ev}_{\mathcal{A}, \mathcal{B}} : \mathbb{C}_{\mathcal{A} \multimap \mathcal{B}} \longrightarrow (\mathcal{A} \multimap \mathcal{B} \parallel \mathcal{A})^\perp \parallel \mathcal{B}.$$

We show that the functor Φ^{-1} together with

$$\text{ev}_{\mathcal{B}, \mathcal{C}} \odot (- \otimes \mathbb{C}_{\mathcal{B}}) : \mathbf{G}[\mathcal{A}, \mathcal{B} \multimap \mathcal{C}] \rightarrow \mathbf{G}[\mathcal{A} \otimes \mathcal{B}, \mathcal{C}],$$

forms an equivalence of categories, from which one can always obtain an adjoint equivalence. To do this we show that for any $\sigma : \mathcal{A} \multimap \mathcal{B} \multimap \mathcal{C}$ there is an isomorphism $\text{ev}_{\mathcal{B}, \mathcal{C}} \odot (\sigma \otimes \mathbb{C}_{\mathcal{B}}) \cong \Phi(\sigma)$, and that this is natural in σ ; from this the equivalence can be immediately derived, because Φ is an isomorphism.

So fix $\sigma : \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B} \multimap \mathcal{C}$, and consider the strategy on the LHS: $\mathbb{C}_{\mathcal{B} \multimap \mathcal{C}} \odot (\mathcal{S} \parallel \mathbb{C}_{\mathcal{B}})$. Its configurations are of the form $(x_{\mathcal{B} \multimap \mathcal{C}}, y_{\mathcal{B} \multimap \mathcal{C}}) \odot (x_{\mathcal{S}} \parallel (z_{\mathcal{B}}, y_{\mathcal{B}}))$ where configurations of copycat are written as pairs; and the variable names emphasise the matching conditions: $\chi y_{\mathcal{B} \multimap \mathcal{C}} = y_{\mathcal{B}} \parallel y_{\mathcal{C}}$, $\sigma x_{\mathcal{S}} = x_{\mathcal{A}} \parallel x_{\mathcal{B} \multimap \mathcal{C}}$, and so on. The isomorphism sends $(x_{\mathcal{B} \multimap \mathcal{C}}, y_{\mathcal{B} \multimap \mathcal{C}}) \odot (x_{\mathcal{S}} \parallel (z_{\mathcal{B}}, y_{\mathcal{B}}))$ to the unique $x_{\mathcal{S}}^\dagger \in \mathcal{C}(\mathcal{S})$ such that $x_{\mathcal{S}}^\dagger \sqsubseteq_{\mathcal{S}} x_{\mathcal{S}}$ and $\Phi(\sigma)(x_{\mathcal{S}}) = x_{\mathcal{A}} \parallel z_{\mathcal{B}} \parallel y_{\mathcal{C}}$. \square

2.6.2 Finite products

Finite products exist in a bicategory provided it has binary products and a terminal object – or more accurately, a *pseudo-terminal* object, but in \mathbf{G} this happens to be strict and given by the empty game \emptyset . Indeed for any negative arena \mathcal{A} , the only negative strategy on $\mathcal{A}^\perp \parallel \emptyset$ is the empty one.

We recall the definition of binary products in a bicategory.

Definition 2.50. (e.g. [FS19]) A bicategory **with binary products** is a bicategory \mathbb{C} equipped with the following data for every $A, B \in \mathbb{C}$:

- an object $A \& B$;
- projection morphisms $\varpi_1 : A \& B \rightarrow A$ and $\varpi_2 : A \& B \rightarrow B$;
- for every $C \in \mathbb{C}$, an adjoint equivalence

$$\begin{array}{ccc} & \xrightarrow{(\varpi_1 \odot -, \varpi_2 \odot -)} & \\ \mathbb{C}[C, A \& B] & \perp & \mathbb{C}[C, A] \times \mathbb{C}[C, B] \\ & \xleftarrow{\langle -, - \rangle} & \end{array}$$

We proceed to define this data for \mathbf{G} . Given arenas \mathcal{A} and \mathcal{B} , the arena $\mathcal{A} \& \mathcal{B}$ is defined to have the same events, causality and polarity as $A \parallel B$, with consistent sets restricted to those of the form $X_A \parallel \emptyset$ for $X_A \in \text{Con}_A$ and $\emptyset \parallel X_B$ for $X_B \in \text{Con}_B$. The isomorphism families $\cong_{A \& B}$, $\cong_{A \& B}^-$ and $\cong_{A \& B}^+$ are restrictions of $\cong_{A \parallel B}$, $\cong_{A \parallel B}^-$ and $\cong_{A \parallel B}^+$ to $\mathcal{C}(A \& B)$.

The projections $\varpi_1 : \mathcal{A} \& \mathcal{B} \rightarrow \mathcal{A}$ and $\varpi_2 : \mathcal{A} \& \mathcal{B} \rightarrow \mathcal{B}$ are obtained by colifting the injections $\mathcal{A}^\perp \rightarrow (\mathcal{A} \& \mathcal{B})^\perp$ and $\mathcal{B}^\perp \rightarrow (\mathcal{A} \& \mathcal{B})^\perp$, respectively. Explicitly, this gives strategies

$$\varpi_1 : \mathbb{C}_A \rightarrow (\mathcal{A} \& \mathcal{B})^\perp \parallel \mathcal{A} \quad \varpi_2 : \mathbb{C}_B \rightarrow (\mathcal{A} \& \mathcal{B})^\perp \parallel \mathcal{B}$$

for every \mathcal{A} and \mathcal{B} . Finally, for every $\mathcal{A}, \mathcal{B}, \mathcal{C}$ we define a pairing functor

$$\langle -, - \rangle : \mathbf{G}[\mathcal{C}, \mathcal{A}] \times \mathbf{G}[\mathcal{C}, \mathcal{B}] \rightarrow \mathbf{G}[\mathcal{C}, \mathcal{A} \& \mathcal{B}]$$

assigning to each pair of strategies $\sigma : \mathcal{S} \rightarrow \mathcal{C}^\perp \parallel \mathcal{A}$ and $\tau : \mathcal{T} \rightarrow \mathcal{C}^\perp \parallel \mathcal{B}$ the strategy

$$\langle \sigma, \tau \rangle : \mathcal{S} \& \mathcal{T} \rightarrow \mathcal{C}^\perp \parallel \mathcal{A} \& \mathcal{B}$$

defined as the composition of $\sigma \& \tau : \mathcal{S} \& \mathcal{T} \rightarrow (\mathcal{C}^\perp \parallel \mathcal{A}) \& (\mathcal{C}^\perp \parallel \mathcal{B})$ with the map $(\mathcal{C}^\perp \parallel \mathcal{A}) \& (\mathcal{C}^\perp \parallel \mathcal{B}) \rightarrow \mathcal{C}^\perp \parallel \mathcal{A} \& \mathcal{B}$ identifying the two copies of \mathcal{C}^\perp . The action on 2-cells is straightforward: if $f : \mathcal{S} \rightarrow \mathcal{S}'$ and $g : \mathcal{T} \rightarrow \mathcal{T}'$ are 2-cells $\sigma \Rightarrow \sigma'$ and $\tau \Rightarrow \tau'$ respectively, then the map $\langle f, g \rangle := f \& g : \mathcal{S} \& \mathcal{T} \rightarrow \mathcal{S}' \& \mathcal{T}'$ is easily seen to be a valid 2-cell $\langle \sigma, \tau \rangle \Rightarrow \langle \sigma', \tau' \rangle$.

Lemma 2.51. *For every $\mathcal{A}, \mathcal{B}, \mathcal{C}$, the following is an adjoint equivalence:*

$$\begin{array}{ccc} & \xrightarrow{(\varpi_1 \odot -, \varpi_2 \odot -)} & \\ \mathbf{G}[\mathcal{C}, \mathcal{A} \& \mathcal{B}] & \perp & \mathbf{G}[\mathcal{C}, \mathcal{A}] \times \mathbf{G}[\mathcal{C}, \mathcal{B}] \\ & \xleftarrow{\langle -, - \rangle} & \end{array}$$

Proof. It suffices to exhibit unit and co-unit natural isomorphisms.

(unit) We show that for every strategy $\rho : \mathcal{R} \rightarrow \mathcal{C}^\perp \parallel \mathcal{A} \& \mathcal{B}$, there is a strong isomorphism of strategies

$$\eta_\rho : \rho \Rightarrow \langle \varpi_1 \odot \rho, \varpi_2 \odot \rho \rangle$$

We construct a map $\eta_\rho : \mathcal{R} \rightarrow (\mathbb{C}_A \odot \mathcal{R}) \& (\mathbb{C}_B \odot \mathcal{R})$.

Note first that any initial move of R must be mapped to an initial move of $\mathcal{C}^\perp \parallel A \& B$, and because R is negative, it must be in the $A \& B$ component (initial moves of \mathcal{C}^\perp are positive). So, because ρ is well-threaded, events of R can be partitioned as $R_A \uplus R_B$ where R_A (resp. R_B) contains those events depending on an initial move e with $\rho(e)$ in the A (resp. B) component. An event of R_A and one of R_B cannot be consistent, since ρ reflects conflict, so any configuration $x \in \mathcal{C}(R)$ satisfies $x \subseteq R_A$ or $x \subseteq R_B$. By restricting the structure of \mathcal{R} we get essps appropriately we get strategies $\rho_A : \mathcal{R}_A \rightarrow \mathcal{C}^\perp \parallel \mathcal{A}$ and $\rho_B : \mathcal{R}_B \rightarrow \mathcal{C}^\perp \parallel \mathcal{B}$, and an iso $\mathcal{R} \rightarrow \mathcal{R}_A \& \mathcal{R}_B$. This is natural in \mathcal{R} : by the condition on maps of strategies any $f \in \mathbf{f} : \rho \Rightarrow \rho'$ sends events of \mathcal{R}_A to those of \mathcal{R}'_A , and events of \mathcal{R}_B to those of \mathcal{R}'_B .

Consider the composition $\varpi_1 \odot \rho$. A configuration of $\mathbb{C}_A \odot \mathcal{R}$ is $(x_A \parallel y_A) \odot x_R$ where $x_R \in \mathcal{C}(R)$ with $\rho x_R = x_C \parallel \iota_A y_A$, where $\iota_A : A \rightarrow A \& B$ is the usual injection. So $x_R \subseteq R_A$; and we observe that (a minor variant of) the map λ_{ρ_A} gives a natural isomorphism $\rho_A \Rightarrow \varpi_1 \odot \rho$. A similar reasoning gives an iso $\rho_B \Rightarrow \varpi_2 \odot \rho$.

So we take η_ρ to be the composite isomorphism

$$\mathcal{R} \xrightarrow{\cong} \mathcal{R}_A \& \mathcal{R}_B \xrightarrow{\lambda_{\rho_A}^{-1} \& \lambda_{\rho_B}^{-1}} (\mathbb{C}_A \odot \mathcal{R}) \& (\mathbb{C}_B \odot \mathcal{R}),$$

natural in ρ .

(co-unit) For every $\sigma : \mathcal{C} \rightarrow \mathcal{A}$ and $\tau : \mathcal{C} \rightarrow \mathcal{B}$ there is a strong invertible 2-cell

$$\varepsilon_{(\sigma, \tau)} : (\varpi_1 \odot \langle \sigma, \tau \rangle, \varpi_2 \odot \langle \sigma, \tau \rangle) \Longrightarrow (\sigma, \tau)$$

which by the remarks of the previous paragraph is equal to $(\lambda_\sigma, \lambda_\tau)$. This is clearly natural, since λ_σ and λ_τ are. □

2.6.3 A linear exponential pseudo-comonad

We define a pseudo-comonad $!$ on \mathbf{G} . It is its *Kleisli bicategory* that we are ultimately interested in: there, a strategy from \mathcal{A} to \mathcal{B} can play the moves of \mathcal{A} as many times as necessary.

Definition 2.52. A pseudo-comonad on a bicategory \mathbb{C} consists of:

- a pseudo-functor $! : \mathbb{C} \rightarrow \mathbb{C}$;
- pseudo-natural transformations $\delta_A : !A \rightarrow !!A$ and $\varepsilon_A : !A \rightarrow A$;
- invertible modifications \mathbf{m} , \mathbf{l} and \mathbf{r} as follows:

$$\begin{array}{ccc} !A & \xrightarrow{\delta_A} & !!A \\ \delta_A \downarrow & \uparrow \mathbf{m} & \downarrow !\delta_A \\ !!A & \xrightarrow{\delta_{!A}} & !!!A \end{array} \quad \begin{array}{ccc} & & !A \\ & \curvearrowright & \downarrow \delta_A \\ !A & \xleftarrow[\varepsilon_A]{} & !!A \xrightarrow[\varepsilon_{!A}]{} !A \\ & \curvearrowleft & \uparrow \mathbf{r} \end{array}$$

subject to coherence axioms [Lac00].

The action of $!$ on objects of \mathbf{G} was defined in 2.2.2: recall in particular that for any $\text{essp } \mathcal{E}$ the $\text{essp } !\mathcal{E}$ has underlying $\text{esp } \parallel_{i \in \omega} E$, and isomorphism family \cong_E making permutable the different copies of E : bijections $\theta : \parallel_{i \in \omega} x_i \cong_{!E} \parallel_{i \in \omega} y_i$ between configurations of $!E$ consist of a permutation $\pi : \omega \rightarrow \omega$ and bijections $\theta_i : x_i \cong_E y_{\pi(i)}$ for each i .

The action of $!$ on strategies is as follows: if $\sigma : \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$ then there is a map $!\mathcal{S} \rightarrow !(\mathcal{A}^\perp \parallel \mathcal{B})$ sending $(i, s) \mapsto (i, \sigma(s))$. The strategy $!\sigma$ is obtained as the composite $!\mathcal{S} \rightarrow !(\mathcal{A}^\perp \parallel \mathcal{B}) \rightarrow !\mathcal{A}^\perp \parallel !\mathcal{B}$ where the second map is the obvious reindexing. Similarly, for every 2-cell $f : \sigma \Rightarrow \sigma'$ we get a 2-cell $!f : !\sigma \Rightarrow !\sigma'$ (and $!f$ is strong whenever f is).

Lemma 2.53. *There is a pseudo-functor $! : \mathbf{G} \rightarrow \mathbf{G}$ whose action on objects, morphisms and 2-cell is as above.*

Proof. That $!$ has a well-defined action on strategies and 2-cells is a routine verification. It is clear that $!$ also preserves identity 2-cells and vertical composition.

To turn $!$ into a pseudo-functor we need 2-cells $\Phi_{\sigma, \tau} : !\tau \odot !\sigma \Rightarrow !(\tau \odot \sigma)$ and $\Phi_{\mathcal{A}} : \mathbb{C}_{!A} \Rightarrow !\mathbb{C}_{\mathcal{A}}$. The former is defined by $\Phi_{\sigma, \tau}(\parallel_{i \in \omega} y_i \odot \parallel_{i \in \omega} x_i) = \parallel_{i \in \omega} (y_i \odot x_i)$ (the causal compatibility conditions hold directly given the argument), while the latter is $\Phi_{\mathcal{A}}(\parallel_{i \in \omega} x_i, \parallel_{i \in \omega} y_i) = \parallel_{i \in \omega} (x_i, y_i)$, writing configurations of $\mathcal{C}(\mathbb{C}_{!A})$ as pairs of configurations of $!A$.

The axioms follow from this definition without difficulty. \square

We continue with the rest of the pseudo-comonad data. Some technical aspects are subtle because of the extra symmetry induced by $!$: in particular it is the first time the structural data requires 2-cells which do not commute strictly (but only up to \sim^+) with the projection to the game.

We start by fixing a bijection $\alpha : \omega^2 \rightarrow \omega$; the particular choice does not matter. Then note that for an $\text{essp } \mathcal{E}$ the map $\underline{\delta} : !!\mathcal{E} \rightarrow !\mathcal{E}$ defined as $\underline{\delta}(i, (j, e)) = (\alpha(i, j), e)$ preserves symmetry (it is a valid map of event structures with symmetry) but does not reflect it in general: suppose \mathcal{E} has a unique event $*$, so that the sets $!E$ and $!!E$ can be identified with ω and ω^2 , respectively. Then, configurations $\{(1, 1), (1, 2)\}$ and $\{(1, 3), (2, 4)\}$ are mapped to symmetric configurations (in $!\mathcal{E}$ any two configurations of same cardinality are symmetric) whilst not being symmetric in $!!\mathcal{E}$ (it is easy to see that no bijection $\pi : \omega \rightarrow \omega$ could be appropriate). The desire to turn $!$ into a monad on event structures was part of Winskel's motivation for introducing symmetry [Win07], as the monadic laws only hold up to \sim . We are interested in positive symmetry, so Winskel's result has to be adapted, but the reasoning is exactly the same:

Lemma 2.54 ([Win07]). *The triple $(!, \underline{\delta}, \underline{\varepsilon})$ is a monad up to \sim^+ on the category of positive arenas and maps between them.*

We see how this lifts to a pseudo-comonad on \mathbf{G} . To construct the necessary structural 2-cells it seems unavoidable to make use of “generalised” maps of strategies, so we reason in the pseudo-double category \mathcal{G} defined previously. It is the case [GP99] that every pseudo-double category is equivalent to a strict double category,

so it is harmless to reason as if composition of strategies was strictly associative and unital; this is what we will do.

For each *negative* arena \mathcal{A} , define a strategy $\delta_{\mathcal{A}}$ as the co-lifting of the map $\underline{\delta}_{\mathcal{A}^\perp}$ above. Similarly let $\varepsilon_{\mathcal{A}}$ be the co-lifting of the injection $\underline{\varepsilon}_{\mathcal{A}^\perp} : \mathcal{A}^\perp \rightarrow !\mathcal{A}^\perp$ defined as $\underline{\varepsilon}_{\mathcal{A}}(a) = (0, a)$. Explicitly, we have strategies

$$\delta_{\mathcal{A}} : \mathbb{C}_{!!\mathcal{A}} \rightarrow !\mathcal{A}^\perp \parallel !!\mathcal{A} \quad \text{and} \quad \varepsilon_{\mathcal{A}} : \mathbb{C}_{\mathcal{A}} \rightarrow !\mathcal{A}^\perp \parallel \mathcal{A}$$

respectively the comultiplication and counit of the pseudo-comonad. That these are pseudo-natural in \mathcal{A} is the statement of the next lemma.

Lemma 2.55. *For every strategy $\sigma : \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$, there are invertible 2-cells $\tilde{\delta}_\sigma : !!\sigma \odot \delta_{\mathcal{A}} \Rightarrow \delta_{\mathcal{B}} \odot !\sigma$ and $\tilde{\varepsilon}_\sigma : \sigma \odot \varepsilon_{\mathcal{A}} \Rightarrow \varepsilon_{\mathcal{B}} \odot !\sigma$, natural in σ , making δ and ε pseudo-natural transformations.*

Proof. Define $\tilde{\varepsilon}_\sigma$ to be the map

$$\begin{array}{ccccccc} !\mathcal{A} & \xrightarrow{\varepsilon_{\mathcal{A}}} & \mathcal{A} & \xrightarrow{\sigma} & \mathcal{B} & \xrightarrow{\varepsilon_{\mathcal{B}}} & \mathcal{B} \\ \parallel & \Downarrow & \downarrow \underline{\varepsilon}_{\mathcal{A}} & \Downarrow \underline{\varepsilon}_{\mathcal{S}} & \downarrow \underline{\varepsilon}_{\mathcal{B}} & \Downarrow & \parallel \\ !\mathcal{A} & \xrightarrow{\varepsilon_{!\mathcal{A}}} & !\mathcal{A} & \xrightarrow{!\sigma} & !\mathcal{B} & \xrightarrow{\varepsilon_{\mathcal{B}}} & \mathcal{B} \end{array}$$

It is not obviously invertible, since $\underline{\varepsilon}_{\mathcal{S}} : \mathcal{S} \rightarrow !\mathcal{S}$ has no inverse. But configurations of $\mathbb{C}_{\mathcal{B}} \odot !\mathcal{S}$ (the composition of $\varepsilon_{\mathcal{B}}$ and $!\sigma$) are those $(x_B \parallel y_B) \odot \parallel_{i \in \omega} x_S^i$ where the matching condition requires that $(!\sigma)_{\mathcal{B}}(\parallel_{i \in \omega} x_S^i) = \underline{\varepsilon}_{\mathcal{B}} x_B$. This implies that $\sigma_{\mathcal{B}} x_0^{\mathcal{S}} = x_B$, and all other $x_i^{\mathcal{S}}$ are empty. So configurations of $\mathbb{C}_{\mathcal{B}} \odot !\mathcal{S}$ correspond to those of $\mathbb{C}_{\mathcal{B}} \odot \mathcal{S}$, making $\tilde{\varepsilon}_\sigma$ an isomorphism.

It is also not immediate that the map $\tilde{\delta}_\sigma$ defined as

$$\begin{array}{ccccccc} !\mathcal{A} & \xrightarrow{\delta_{\mathcal{A}}} & !!\mathcal{A} & \xrightarrow{!!\sigma} & !!\mathcal{B} & \xrightarrow{\varepsilon_{!!\mathcal{B}}} & !!\mathcal{B} \\ \parallel & \Downarrow & \downarrow \underline{\delta}_{\mathcal{A}} & \Downarrow \underline{\delta}_{\mathcal{S}} & \downarrow \underline{\delta}_{\mathcal{B}} & \Downarrow & \parallel \\ !\mathcal{A} & \xrightarrow{\varepsilon_{!\mathcal{A}}} & !\mathcal{A} & \xrightarrow{!\sigma} & !\mathcal{B} & \xrightarrow{\delta_{\mathcal{B}}} & !!\mathcal{B} \end{array}$$

has an inverse: as per the discussion above, the map $\underline{\delta}_{\mathcal{A}} : !!\mathcal{A} \rightarrow !\mathcal{A}$ is an isomorphism between the underlying event structures but does not reflect symmetry (and consequently has no inverse in the category of essps).

But in fact the bijection $\text{id} \odot \underline{\delta}_{\mathcal{S}} : \mathbb{C}_{!!\mathcal{B}} \odot !!\mathcal{S} \rightarrow \mathbb{C}_{!!\mathcal{B}} \odot !\mathcal{S}$ does reflect symmetry. To see this, let

$$(\theta, \phi) \odot \psi : \left(\parallel_{i,j} x_{i,j}, \parallel_{i,j} y_{i,j} \right) \odot \left(\parallel_k z_k \right) \cong_{\mathbb{C}_{!!\mathcal{B}} \odot !\mathcal{S}} \left(\parallel_{i,j} x'_{i,j}, \parallel_{i,j} y'_{i,j} \right) \odot \left(\parallel_k z'_k \right)$$

be a symmetry between configurations of $\mathbb{C}_{!!\mathcal{B}} \odot !\mathcal{S}$. By Lemma 2.19, and by definition of the symmetry on copycat, this is determined by a family of bijections: $\theta_{i,j} : x_{i,j} \cong_{!!\mathcal{B}} x'_{\pi(i), \pi_i(j)}$, $\phi_{i,j} : y_{i,j} \cong_{!!\mathcal{B}} y'_{\pi(i), \pi_i(j)}$ and $\psi_k : z_k \cong_{!S} z'_{\xi(k)}$ for $i, j, k \in \omega$ (where π, π_i and ξ are bijections $\omega \rightarrow \omega$).

The difficulty is to show that the image of $\psi : \parallel_k z_k \cong_{!S} \parallel_k z'_{\xi(k)}$ under $\underline{\delta}_S^{-1}$, the bijection $\underline{\delta}_S^{-1}\psi : \parallel_i \parallel_j z_{\alpha(i,j)} \cong \parallel_i \parallel_j z_{\xi(\alpha(i,j))}$ (α is the pairing function), is an element of $\cong_{!!S}$. For this it suffices to show that there are bijections $\kappa, \kappa_i : \omega \rightarrow \omega$, for $i \in \omega$, such that $\xi(\alpha(i, j)) = \alpha(\kappa(i), \kappa_i(j))$ for each i, j .

But by the matching condition of Lemma 2.19, the projection of $\underline{\delta}_S\psi$ to $!B$ coincides with the bijection $\underline{\delta}_B\theta : \parallel_i \parallel_j x_{\alpha(i,j)} \cong_{!!B} \parallel_i \parallel_j x'_{\alpha\pi(i), \pi_i(j)}$, and since all initial moves of S are B -moves, the result holds, we must have $\xi(\alpha(i, j)) = \alpha(\pi(i), \pi_i(j))$ as required. \square

Finally we must exhibit coherence modifications as part of the pseudo-comonad structure. These are obtained by combining the various 2-cells associated with co-lifted strategies, and the monadic laws for $!$ on positive arenas. For instance $\mathbf{l}_{\mathcal{A}}$ is given by a 2-cell of the form

$$\begin{array}{ccccc}
!A & \xrightarrow{\delta_{\mathcal{A}}} & !!A & \xrightarrow{!\varepsilon_{\mathcal{A}}} & !A & \xrightarrow{\alpha_{!A}} & !A \\
\parallel & \Downarrow & \parallel & \Downarrow & \downarrow \text{!}\varepsilon_{\mathcal{A}} & \parallel & \\
!A & \xrightarrow{\delta_{\mathcal{A}}} & !!A & \xrightarrow{\alpha_{!A}} & !!A & (\star) & \\
\parallel & \Downarrow & \parallel & \Downarrow & \downarrow \delta_{\mathcal{A}} & \parallel & \\
!A & \xrightarrow{\alpha_{!A}} & !A & \xrightarrow{\alpha_{!A}} & !A & &
\end{array}$$

where the 2-cell marked (\star) is constructed using one of the monad laws for maps of essps (Lemma 2.54).

Lemma 2.56. *The data $(!, \delta, \varepsilon, \mathbf{m}, \mathbf{r}, \mathbf{l})$ forms a pseudo-comonad on \mathbf{G} .*

2.6.4 The cartesian closed Kleisli bicategory

Going towards applications in semantics, we move to a cartesian closed setting by considering the *Kleisli bicategory* for the pseudo-comonad $!$. Familiarity with categorical models of linear logic should make the results of this section unsurprising; the path we take is reminiscent of the Seely categories described (for 1-categories) in *e.g.* [Mel09].

Definition 2.57 (*e.g.* [FGHW08]). Let \mathbb{C} be a bicategory and $!$ be a pseudo-comonad. The **Kleisli bicategory** $\mathbb{C}_!$ has objects those of \mathbb{C} , and morphisms and 2-cells defined by $\mathbb{C}_![A, B] = \mathbb{C}[!A, B]$. The composition in $\mathbb{C}_!$ of $\sigma : !A \rightarrow B$ and $\tau : !B \rightarrow C$, written $g \circ_! f$ is defined as $g \circ !f \circ \delta_A$, and the identity morphism on A is the component at A of the co-unit for $!$: $\varepsilon_A \in \mathbb{C}_![A, A]$.

Associator and unitors are defined using the pseudo-comonad data and shown to satisfy the necessary coherence axioms; we omit the details as the thesis does not require them.

Instantiating Definition 2.45 with the “cartesian” symmetric monoidal structure yields the following: a **cartesian closed bicategory** is a bicategory \mathbb{D} with finite

products equipped with, for objects A and B , an object $A \rightarrow B$, a map $\text{Ev}_{A,B} : (A \rightarrow B) \& A \rightarrow B$, and an adjoint equivalence

$$\begin{array}{ccc} & \xrightarrow{\text{Ev}_{B,C} \circ (-\& B)} & \\ \mathbb{D}[A, B \rightarrow C] & \perp & \mathbb{D}[A \& B, C] \\ & \xleftarrow{\text{Cur}} & \end{array}$$

Theorem 2.58. *Let \mathbb{C} be a symmetric monoidal closed bicategory with finite products. Let $!$ be a pseudo-comonad on \mathbb{C} and suppose that there is a natural family of adjoint equivalences $m_{A,B} : !(A \& B) \rightarrow !A \otimes !B$. Suppose also that there is an invertible modification*

$$\begin{array}{ccc} !(A \& B) & \xrightarrow{\delta} & !!(A \& B) \\ \downarrow m & \Rightarrow \bar{m} & \downarrow \langle !\varpi_1, !\varpi_2 \rangle \\ !A \otimes !B & \xrightarrow{\delta \otimes \delta} & !!A \otimes !!B \\ & & \downarrow m \\ & & !(A \& B) \end{array}$$

Then, the Kleisli bicategory $\mathbb{C}_!$ is cartesian closed.

Proof. It is a standard result [Mel09] that $\mathbb{C}_!$ has binary products induced from those in \mathbb{C} . Explicitly, $A \& B$ is the product in \mathbb{C} , projections are defined as $\varpi_1 : !(A \& B) \xrightarrow{\varepsilon} A \& B \xrightarrow{\varpi_1} A$ and $\varpi_2 : !(A \& B) \xrightarrow{\varepsilon} A \& B \xrightarrow{\varpi_2} B$. The terminal object in $\mathbb{C}_!$ is the same as in \mathbb{C} .

For the closed structure, define $A \rightarrow B = !A \multimap B$ and define the map $\text{Ev}_{A,B} \in \mathbb{C}_![A \rightarrow B \& A, B]$ to be the composite

$$!(A \rightarrow B \& A) \xrightarrow{m} !(A \rightarrow B) \otimes !A \xrightarrow{\varepsilon \otimes !A} A \rightarrow B \otimes !A \xrightarrow{\text{ev}_{!A,B}} B.$$

Given $\sigma \in \mathbb{C}_![A \& B, C]$ we define its currying $\text{Cur}(\sigma) \in \mathbb{C}_![A, B \Rightarrow C]$ as $\Lambda(\sigma \circ m_{A,B}^\bullet)$. Because $m_{A,B}$ is an adjoint equivalence, by a straightforward argument there is an adjoint equivalence

$$\begin{array}{ccc} & \xrightarrow{\Lambda^{-1}(-) \circ m_{A,B}} & \\ \mathbb{C}_![A, B \rightarrow C] & \perp & \mathbb{C}_![A \& B, C] \\ & \xleftarrow{\Lambda(- \circ m_{A,B}^\bullet)} & \end{array}$$

and because adjoints are preserved by natural isomorphism, to obtain the required adjoint equivalence

$$\begin{array}{ccc} & \xrightarrow{\text{Ev}_{B,C} \circ (-\& \varepsilon_B)} & \\ \mathbb{C}_![A, B \rightarrow C] & \perp & \mathbb{C}_![A \& B, C] \\ & \xleftarrow{\text{Cur}} & \end{array}$$

it suffices to exhibit an isomorphism $\text{Ev}_{B,C} \circ !(\sigma \& \varepsilon_B) \cong \Lambda^{-1}(\sigma) \circ m$, natural in σ . This is given by the following pasting diagram:

$$\begin{array}{ccccc}
!(A \& B) & \xrightarrow{\delta} & !!(A \& B) & \xrightarrow{! \langle \sigma \circ !\pi_1, \varepsilon_B \circ !\pi_2 \rangle} & !((B \rightarrow C) \& B) \\
\downarrow m & & \searrow ! \langle !\pi_1, !\pi_2 \rangle & & \swarrow !(\sigma \& \varepsilon_B) \\
& & & & & !(A \& !B) \\
& & & & & \downarrow m \\
!A \otimes !B & \xrightarrow{\delta \otimes \delta} & !!A \otimes !!B & \xrightarrow{!\sigma \otimes !\varepsilon_B} & !(B \rightarrow C) \otimes !B \\
\downarrow \Lambda^{-1}(\sigma) & \searrow \sigma \otimes !B & \downarrow \varepsilon_{!A} \otimes !\varepsilon_B & \swarrow \varepsilon_{B \rightarrow C} \otimes !B & \\
& & !A \otimes !B & & \\
& & \downarrow \sigma \otimes !B & & \\
& & (B \rightarrow C) \otimes !B & & \\
& & \downarrow \text{ev} & & \\
& & C & &
\end{array}$$

□

We show that the bicategory \mathbf{G} can be equipped with this structure. Fix a bijection $\beta : \omega \uplus \omega \rightarrow \omega$, and for negative arenas \mathcal{A}, \mathcal{B} , define $\underline{m}_{\mathcal{A}, \mathcal{B}} : !(A \& B) \rightarrow !A \parallel !B$ and $\underline{m}_{\mathcal{A}, \mathcal{B}}^\bullet : !A \parallel !B \rightarrow !(A \& B)$ as follows, for every $i \in \omega$, $a \in A$ and $b \in B$:

$$\begin{array}{ll}
\underline{m}_{\mathcal{A}, \mathcal{B}}(i, a) = (1, (i, a)) & \underline{m}_{\mathcal{A}, \mathcal{B}}^\bullet(1, (i, a)) = (\beta(1, i), a) \\
\underline{m}_{\mathcal{A}, \mathcal{B}}(i, b) = (2, (i, b)) & \underline{m}_{\mathcal{A}, \mathcal{B}}^\bullet(2, (i, b)) = (\beta(2, i), b)
\end{array}$$

We get two maps of essps, inverses up to positive symmetry in the case where \mathcal{A} and \mathcal{B} are *positive* games.

Lemma 2.59. *For positive games \mathcal{A} and \mathcal{B} , $\underline{m}_{\mathcal{A}, \mathcal{B}} \circ \underline{m}_{\mathcal{A}, \mathcal{B}}^\bullet \sim^+ \text{id}_{\mathcal{B}}$ and $\underline{m}_{\mathcal{A}, \mathcal{B}}^\bullet \circ \underline{m}_{\mathcal{A}, \mathcal{B}} \sim^+ \text{id}_{\mathcal{A}}$.*

Proof. Direct verification. □

So in particular, for \mathcal{A}, \mathcal{B} *negative* games, the lemma applies to maps $\underline{m}_{\mathcal{A}^\perp, \mathcal{B}^\perp}$ and $\underline{m}_{\mathcal{A}^\perp, \mathcal{B}^\perp}^\bullet$ which can also be checked to be courteous, receptive and \sim -receptive.

It is an easy consequence of Lemma 2.27 that any isomorphism of essps lifts to an equivalence. This can be generalised to isomorphisms “up to \sim^+ ”:

Lemma 2.60. *Let \mathcal{A} and \mathcal{B} be games with symmetry, and let $f : \mathcal{A} \rightarrow \mathcal{B}$ and $f^\bullet : \mathcal{B} \rightarrow \mathcal{A}$ be courteous, receptive and \sim -receptive maps of essps such that $f \circ f^\bullet \sim^+ \text{id}_{\mathcal{B}}$ and $f^\bullet \circ f \sim^+ \text{id}_{\mathcal{A}}$. Then there is an adjoint equivalence*

$$\begin{array}{ccc}
& \xrightarrow{\widehat{f}} & \\
\mathcal{A} & & \mathcal{B} \\
& \xleftarrow{\widehat{f}^\bullet} &
\end{array}$$

in the bicategory \mathbf{G} . Similarly, if $g : \mathcal{B}^\perp \rightarrow \mathcal{A}^\perp$ and $g^\bullet : \mathcal{A}^\perp \rightarrow \mathcal{B}^\perp$ are courteous, receptive and \sim -receptive maps of essps such that $g \circ g^\bullet \sim^+ \text{id}_{\mathcal{B}^\perp}$ and $g^\bullet \circ g \sim^+ \text{id}_{\mathcal{A}^\perp}$. Then

$$\mathcal{A} \begin{array}{c} \xrightarrow{\check{g}} \\ \xleftarrow{g^\bullet} \end{array} \mathcal{B}$$

forms an adjoint equivalence in \mathbf{G} .

Proof. There is an isomorphism $\alpha_{\mathcal{A}} \Rightarrow \hat{f}^\bullet \odot \hat{f}$ obtained as follows:

$$\begin{array}{ccccc} \mathcal{A} & \xrightarrow{\alpha_{\mathcal{A}}} & \mathcal{A} & \xrightarrow{\alpha_{\mathcal{A}}} & \mathcal{A} \\ \parallel & & \downarrow & & \downarrow f \\ \mathcal{A} & \xrightarrow{\quad} & \mathcal{B} & & \mathcal{B} \\ \parallel & & \hat{f} & \cong & \parallel \\ \mathcal{A} & \xrightarrow{\hat{f}} & \mathcal{B} & \xrightarrow{\alpha_{\mathcal{A}}} & \mathcal{B} \\ \parallel & & \parallel & & \downarrow f^\bullet \\ \mathcal{A} & \xrightarrow{\hat{f}} & \mathcal{B} & \xrightarrow{\hat{f}^\bullet} & \mathcal{A} & \xrightarrow{\alpha_{\mathcal{A}}} & \mathcal{A} \end{array} \quad (*)$$

In the diagram the 2-cell labelled $(*)$ is constructed using the assumption that $f^\bullet \circ f \sim^+ \text{id}_{\mathcal{A}}$. There is also an iso $\hat{f} \odot \hat{f}^\bullet \Rightarrow \alpha_{\mathcal{B}}$ described using a symmetric diagram. This proves the first part of the statement; the second part is similar. \square

By the second part of Lemma 2.60, the co-lifting construction gives strategies $m_{\mathcal{A},\mathcal{B}} : \mathbb{C}_{!_{\mathcal{A}} \parallel !_{\mathcal{B}}} \rightarrow !(\mathcal{A} \& \mathcal{B})^\perp \parallel !_{\mathcal{A}} \parallel !_{\mathcal{B}}$ and $m_{\mathcal{A},\mathcal{B}}^\bullet : \mathbb{C}_{!(\mathcal{A} \& \mathcal{B})} \rightarrow !(\mathcal{A} \parallel !_{\mathcal{B}})^\perp \parallel !(\mathcal{A} \& \mathcal{B})$ which together form an adjoint equivalence. This forms a natural transformation, since for each $\sigma \in \mathbf{G}[\mathcal{A}, \mathcal{A}']$ and $\tau \in \mathbf{G}[\mathcal{B}, \mathcal{B}']$ there is an invertible 2-cell $\tilde{m}_{\sigma,\tau} : m_{\mathcal{A}',\mathcal{B}'} \odot !(\sigma \& \tau) \Rightarrow !\sigma \otimes !\tau \odot m_{\mathcal{A},\mathcal{B}}$.

Theorem 2.61. *The Kleisli bicategory $\mathbf{G}_!$ is cartesian closed.*

Proof. By Theorem 2.58 it is enough to give a family of invertible 2-cells

$$w_{\mathcal{A},\mathcal{B}} : (\delta_{\mathcal{A}} \otimes \delta_{\mathcal{B}}) \odot m_{\mathcal{A},\mathcal{B}} \Rightarrow m_{!_{\mathcal{A}},!_{\mathcal{B}}} \odot \langle !\pi_1, !\pi_2 \rangle \odot \delta_{\mathcal{A} \& \mathcal{B}}$$

satisfying the modification axiom. For this we observe that the strategy $\langle !\pi_1, !\pi_2 \rangle$ is isomorphic to the co-lifted strategy \check{g} where $g : !_{\mathcal{A}}^\perp \& !_{\mathcal{B}}^\perp \rightarrow !(\mathcal{A} \& \mathcal{B})^\perp$ has the obvious action on events. (The isomorphism is a map $\mathbb{C}_{!(!_{\mathcal{A}} \& !_{\mathcal{B}})} \rightarrow !(\mathbb{C}_{\mathcal{A}} \& \mathbb{C}_{\mathcal{B}})$.) Using this, the map $w_{\mathcal{A},\mathcal{B}}$ is derived by repeated applications of Lemma 2.27. \square

2.6.5 Colimits of strategies

Finally, we show the existence of certain ω -colimits in hom-categories $\mathbf{G}[\mathcal{A}, \mathcal{B}]$. These will be used to interpret recursion operators, and more generally to describe strategies via ω -chains of finite approximations. Fix an arbitrary arena \mathcal{A} . We study colimits in $\mathbf{G}[\emptyset, \mathcal{A}]$, and the results apply to $\mathbf{G}[\mathcal{A}, \mathcal{B}]$ for any \mathcal{B} since there is an isomorphism $\mathbf{G}[\mathcal{A}, \mathcal{B}] \cong \mathbf{G}[\emptyset, \mathcal{A} \multimap \mathcal{B}]$.

For each $i \in \omega$, let $\sigma_i : \mathcal{S}_i \rightarrow \mathcal{A}$ be a strategy, and let there be a chain

$$\sigma_0 \xrightarrow{f_0} \sigma_1 \xrightarrow{f_1} \dots$$

in $\mathbf{G}[\emptyset, \mathcal{A}]$.

Strong inclusions. The simplest case is that in which $\mathcal{S}_i \subseteq \mathcal{S}_{i+1}$ for every i , the f_i are inclusion maps, and $\sigma_i(s) = \sigma_{i+1}(s)$ for every $s \in \mathcal{S}_i$. Note that this is inclusion of event structures with symmetry, *i.e.* componentwise inclusion $(\mathcal{S}_i, \leq_{\mathcal{S}_i}, \text{Con}_{\mathcal{S}_i}, \cong_{\mathcal{S}_i}) \subseteq (\mathcal{S}_{i+1}, \leq_{\mathcal{S}_{i+1}}, \text{Con}_{\mathcal{S}_{i+1}}, \cong_{\mathcal{S}_{i+1}})$. In this case, the colimit $\bigvee_{i \in \omega} \sigma_i$ exists and can be obtained by taking the componentwise union:

$$\bigvee_{i \in \omega} \mathcal{S}_i \triangleq \bigcup_{i \in \omega} (\mathcal{S}_i, \leq_{\mathcal{S}_i}, \text{Con}_{\mathcal{S}_i}, \cong_{\mathcal{S}_i}).$$

The map $\bigvee_{i \in \omega} \sigma_i : \bigvee_{i \in \omega} \mathcal{S}_i \rightarrow \mathcal{A}$ is determined by the σ_i , and the 2-cells $\sigma_i \Rightarrow \bigvee_{i \in \omega} \sigma_i$ are the obvious inclusions.

Strong embeddings. Suppose more generally that the f_i are *embeddings*, *i.e.* injective, order-preserving maps of essps. The above construction still applies because every embedding can be turned into an inclusion.

Indeed, given strategies $\sigma : \mathcal{S} \rightarrow \mathcal{A}$ and $\tau : \mathcal{T} \rightarrow \mathcal{A}$ and an embedding $f : \sigma \Rightarrow \tau$, we can relabel events in \mathcal{T} to obtain an essp \mathcal{T}^\dagger with $\mathcal{T} \cong \mathcal{T}^\dagger$ and $\mathcal{S} \subseteq \mathcal{T}^\dagger$. We get a strategy $\tau^\dagger : \mathcal{T}^\dagger \rightarrow \mathcal{A}$, a strong isomorphism $\tau \cong \tau^\dagger$ and a strong inclusion $\iota : \sigma \Rightarrow \tau^\dagger$, such that

$$\begin{array}{ccc} \sigma & \xrightarrow{f} & \tau \\ & \searrow \iota & \parallel \cong \\ & & \tau^\dagger \end{array}$$

commutes.

So we construct colimit for the σ_i by considering the chain of strong inclusions

$$\sigma_0^\dagger \Rightarrow \sigma_1^\dagger \Rightarrow \sigma_2^\dagger \Rightarrow \dots$$

where $\sigma_0^\dagger = \sigma_0$ and each subsequent σ_i^\dagger is obtained inductively by applying the above to the composite $\sigma_{i-1} \cong \sigma_{i-1}^\dagger \xrightarrow{f_i} \sigma_i$. Standard reasoning then shows that the strategy $\bigvee_{i \in \omega} \sigma_i^\dagger$ is a colimit for the σ_i , with each $\sigma_j \Rightarrow \bigvee_{i \in \omega} \sigma_i^\dagger$ factoring through σ_j^\dagger .

Preservation of colimits. Importantly, ω -colimits of strong embeddings commute with the various constructions of this chapter.

Proposition 2.62. *Let $\sigma_i : \mathcal{S}_i \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$ be a strategy for $i \in \omega$, and let there be an ω -chain $\sigma_0 \xrightarrow{f_0} \sigma_1 \xrightarrow{f_1} \dots$ of strong embeddings. If $\tau \in \mathbf{G}[\mathcal{B}, \mathcal{C}]$, then there is a chain $\tau \odot \sigma_0 \xrightarrow{\text{id} \odot f_0} \tau \odot \sigma_1 \xrightarrow{\text{id} \odot f_1} \dots$ of strong embeddings, and we have $\tau \odot (\bigvee_{i \in \omega} \sigma_i) \cong \bigvee_{i \in \omega} (\tau \odot \sigma_i)$. Similarly, the colimit commutes with precomposition, tensor, currying, pairing, and !.*

Proof. Routine. □

Chapter 3

Probability in concurrent games

In this chapter, we enrich the strategies of Chapter 2 with probability.

By way of illustration, consider a game \mathcal{A} with three events a^- , b^+ and c^+ , and trivial consistency, polarity, and symmetry. (Events are directly annotated with their polarity; we will use this notation throughout the thesis.) Consider the following strategy $\sigma : \mathcal{S} \rightarrow \mathcal{A}$:

$$a^- \begin{array}{l} \triangleright b^+ \\ \triangleleft \begin{array}{l} \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \end{array} \\ \triangleright c^+ \end{array}$$

Here Player waits until Opponent plays move a ; and then plays either b or c , nondeterministically. A *probabilistic* version of this strategy must additionally give a probability coefficient to every possible outcome of the nondeterministic branching:

$$a^- \begin{array}{l} \triangleright b^+ \quad p_b \\ \triangleleft \begin{array}{l} \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \end{array} \\ \triangleright c^+ \quad p_c \end{array}$$

where p_b and p_c are positive reals with $p_b + p_c \leq 1$. (The inequality is because Player may choose not to play either.)

Though this may seem straightforward, there are difficulties in adding probability to strategies in general. This is discussed in 3.1. Later on we see how probabilistic strategies can be composed and organised into a bicategory (this involves defining identities and composition in this setting), as in Chapter 2. This bicategory has good structural properties, many of which can be lifted directly from the structure of \mathbf{G} .

The work presented here is based on Winskel's model of basic concurrent games and probabilistic strategies [Win13a]. The addition of symmetry, as well as the study of the categorical structure, are contributions of this thesis.

3.1 Probabilistic strategies

We have seen in the example above how probability can be added to an event structure with polarity representing a simple nondeterministic branching. When

event structures are not *tree-shaped*, however, the addition of probability is more subtle.

Consider for instance removing the conflict between b and c in \mathcal{S} above: define a strategy $\tau : \mathcal{T} \rightarrow \mathcal{A}$ to be

$$a^- \begin{array}{c} \triangleright b^+ \\ \triangleleft c^+ \end{array}$$

The situation for Player is no longer a simple nondeterministic choice; indeed b and c may both be played in the same execution. We are still interested in the respective probability of b and c occurring, but we may also want to express probabilistic dependency between the two events. So probability coefficients are assigned to configurations rather than individual events. As we will see, this approach also allows for a smoother integration with the rest of the structure: causality, consistency, and (later on) symmetry.

More precisely, for each $x \in \mathcal{C}(E)$ we assign coefficients to **positive extensions** of x , *i.e.* configurations $y \in \mathcal{C}(E)$ such that $x \sqsubseteq^+ y$. We write $v(y | x)$ for this coefficient, representing the conditional probability that y will occur given that x has. If $v(- | -)$ is to make sense as a form of conditional probability, we must have $v(x | x) = 1$, and a *chain rule*: $v(z | x) = v(y | x)v(z | y)$, when $x \sqsubseteq^+ y \sqsubseteq^+ z$.

We must also ensure that $v(- | x)$ is a probability distribution on the positive extensions of x . If those extensions are pairwise incompatible, then indeed the sum $\sum_{x \sqsubseteq^+ y} v(y | x)$ must be ≤ 1 ; if instead extensions y_1, \dots, y_n are not pairwise mutually exclusive then we must account for any overlap, using the inclusion-exclusion principle. This is condition (3) in the definition below, called **drop condition** in [Win13a]. Condition (4) formalises the requirement that Player and Opponent, whenever they are causally independent, are also probabilistically independent; and finally condition (5) forces Player to play *symmetric* configurations with equal probability.

Definition 3.1. A **conditional valuation** on an esp S is a family $(v(y | x))_{x \sqsubseteq^+ y}$ of coefficients in $[0, 1]$ satisfying

- (1) $v(x | x) = 1$ for all $x \in \mathcal{C}(S)$;
- (2) if $x \sqsubseteq^+ y \sqsubseteq^+ z$ then $v(z | x) = v(y | x)v(z | y)$;
- (3) if $x \sqsubseteq^+ y_1, \dots, y_n$, then

$$\sum_I (-1)^{|I|+1} v(\cup_{i \in I} y_i | x) \leq 1$$

where I ranges over nonempty subsets of $\{1, \dots, n\}$ such that $\cup_{i \in I} y_i \in \mathcal{C}(S)$;

- (4) if $x \sqsubseteq^+ y$, $x \sqsubseteq^- z$ and $y \cup z \in \mathcal{C}(S)$, then $v(y | x) = v(y \cup z | z)$;
- (5) if $\theta : x \cong_S y$ and $\theta \sqsubseteq^+ \theta' : x' \cong_S y'$, then $v(x' | x) = v(y' | y)$.

Conditional valuations are an intuitive way of making Player's behaviour probabilistic. Given a conditional valuation on an esp \mathcal{S} , one can recover a form of *absolute* (*i.e. unconditional*) probability $v(x)$ for each configuration $x \in \mathcal{C}(S)$. Due

to the presence of Opponent events this will not strictly speaking be an unconditional probability distribution; more accurately, $v(x)$ is the probability that an execution will *reach* configuration x , provided Opponent chooses to play all negative moves in x .

To define $v : \mathcal{C}(S) \rightarrow [0, 1]$ from the conditional family $\{v(y \mid x)\}_{x \subseteq^+ y}$ we compute the product of the conditional probabilities along a covering chain. A key observation is that the particular choice of covering chain does not affect the resulting value:

Lemma 3.2. *If $y \subseteq x \in \mathcal{C}(S)$ and there are chains*

$$\begin{aligned} y \subseteq^+ x_1 \subseteq^- x_2 \subseteq^+ x_3 \subseteq^- \cdots \subseteq^- x_{n-1} \subseteq^+ x_n = x \\ y \subseteq^+ x'_1 \subseteq^- x'_2 \subseteq^+ x'_3 \subseteq^- \cdots \subseteq^- x'_{m-1} \subseteq^+ x'_m = x \end{aligned}$$

we have

$$\begin{aligned} v(x_1 \mid y) \times v(x_3 \mid x_2) \times \cdots \times v(x_n \mid x_{n-1}) \\ = v(x'_1 \mid y) \times v(x'_3 \mid x'_2) \times \cdots \times v(x'_m \mid x'_{m-1}). \end{aligned}$$

Proof. By induction on $|x \setminus y|$. The equality holds trivially when $x = y$.

If $x_1 = x'_1 = y$ then we can assume *w.l.o.g.* that $\emptyset \subsetneq^- x_2$ and $\emptyset \subsetneq^- x'_2$. Then, fix a chain of the form $x_2 \subseteq^- x_2 \cup x'_2 \subseteq^+ z_1 \subseteq \dots \subseteq z_k \subseteq^+ x$; this has length strictly less than $|x \setminus y|$, so by the induction hypothesis we have

$$\begin{aligned} v(x_3 \mid x_2) \times \cdots \times v(x \mid x_{n-1}) \\ = v(z_1 \mid x_2 \cup x'_2) \times \cdots \times v(x \mid z_k). \end{aligned}$$

and by a symmetric argument the RHS is also equal to $v(x'_3 \mid x'_2) \times \cdots \times v(x \mid x'_{m-1})$, from which we conclude using that $v(x_1 \mid y) = v(x'_1 \mid y)$.

If $x_1 = y$ but $y \subsetneq x_1$, then by axiom (4) for conditional valuations we have $v(x_1 \mid y) = v(x_2 \cup x'_1 \mid x_2)$. Then, by fixing a chain $x_2 \cup x'_1 \subseteq^+ z_1 \subseteq \cdots \subseteq z_k \subseteq^+ x$, we get

$$\begin{aligned} v(x'_1 \mid y) \times v(x'_3 \mid x'_2) \times \cdots \times v(x \mid x_{n-1}) \\ &= v(x_2 \cup x'_1 \mid x_2) \times v(x'_3 \mid x'_2) \times \cdots \times v(x \mid x'_{m-1}) \\ &= v(x_2 \cup x'_1 \mid x_2) \times v(z_1 \mid x_2 \cup x'_1) \cdots \times v(x \mid z_k) && \text{by IH for } x'_1 \subseteq x \\ &= v(z_1 \mid x_2) \cdots \times v(x \mid z_k) && \text{by axiom (2)} \\ &= v(x_3 \mid x_2) \times \cdots \times v(x \mid x_{n-1}) && \text{by IH for } x_2 \subseteq x \end{aligned}$$

which is the desired result as $y = x_1$.

Finally if $y \subsetneq x_1$ and $y \subsetneq x'_1$, then $v(x_1 \mid y) \times v(x_1 \cup x'_1 \mid x_1) = v(x'_1 \mid y) \times v(x_1 \cup x'_1 \mid x'_1)$, by axiom (2) for conditional valuations, so that the result follows using the induction hypothesis for $x_1 \subseteq x$ and $x'_1 \subseteq x$. \square

The resulting function $v : \mathcal{C}(S) \rightarrow [0, 1]$ is an instance of the following:

Definition 3.3. A **valuation** on an event structure with polarity S is a function $v : \mathcal{C}(S) \rightarrow [0, 1]$ satisfying

- (1) $v(\emptyset) = 1$;
- (2) $v(x) = v(y)$ when $x \subseteq^- y$;
- (3) for every $x \subseteq^+ y_1, \dots, y_n$,

$$v(x) \geq \sum_I (-1)^{|I|+1} v(\cup_{i \in I} y_i)$$

where I ranges over nonempty subsets of $\{1, \dots, n\}$ such that $\cup_{i \in I} y_i \in \mathcal{C}(S)$;

- (4) if $\theta : x \cong_S y$, then $v(x) = v(y)$;

We check:

Lemma 3.4. *If $\{v(y | x)\}_{x \subseteq y}$ is a conditional valuation on \mathcal{S} , then defining $v(x) = v(x_1 | \emptyset) \times v(x_3 | x_2) \times \dots \times v(x | x_{n-1})$ with respect to any covering chain $\emptyset \subseteq^+ x_1 \subseteq^- x_2 \subseteq^+ x_3 \subseteq^- \dots \subseteq^- x_{n-1} \subseteq^+ x_n$ yields a valuation on \mathcal{S} (the choice of covering chain does not matter by Lemma 3.2).*

Proof. Axioms (1)-(3) follow directly from the definition. For (4), it suffices to observe that if $\theta : x \cong y$ and $\emptyset \subset x_1 \subset \dots \subset x_{n-1} \subset x$ is a covering chain for x , then there is a covering chain $\emptyset \subset y_1 \subset \dots \subset y_{n-1} \subset y$ for y such that θ restricts to $\theta_i : x_i \cong_S y_i$ for each i . The property then follows from axiom (5) for conditional valuations. \square

Note that valuations are necessarily anti-monotone:

Lemma 3.5. *If v is a valuation on \mathcal{S} and $x \subseteq y \in \mathcal{C}(S)$ then $v(x) \geq v(y)$.*

Proof. If $x \subseteq^- y$, then $v(y) = v(x)$ by the second condition in Definition 3.3, and if $x \subseteq^+ y$ then the third condition simplifies to $v(x) \geq v(y)$. Since for every inclusion $x \subseteq y$ there exists a chain $x \subseteq^+ x_1 \subseteq^- x_2 \subseteq^+ \dots \subseteq^+ x_n = y$, the statement follows by transitivity of \geq . \square

Lemma 3.6. *If v is a valuation on \mathcal{S} , then the following satisfies the axioms for a conditional valuation:*

$$v(y | x) = \begin{cases} \frac{v(y)}{v(x)} & \text{if } v(x) \neq 0 \\ 1 & \text{if } v(x) = 0 \text{ and } x = y \\ 0 & \text{if } v(x) = 0 \text{ and } x \neq y \end{cases}$$

for every $x \subseteq^+ y \in \mathcal{C}(S)$.

Proof. We check the conditions of Definition 3.1 for the family $\{v(y | x)\}$. Condition (1) follows directly from the definition. Suppose $x \subseteq^+ y \subseteq^+ z$; we show that $v(z | x) = v(y | x)v(z | y)$. If any of the two inclusions is an equality, the statement follows from (1). If both inclusions are strict and $v(x) = 0$, both LHS and RHS are zero by definition. If $v(x) \geq 0$ and $v(y) = 0$, then clearly RHS = 0; but $v(z) = 0$ by Lemma 3.5, and so LHS = 0 as required. Finally if $v(y) \geq 0$, $v(y | x)v(z | y) = \frac{v(y)}{v(x)} \frac{v(z)}{v(y)} = \frac{v(z)}{v(x)} = v(z | x)$. So condition (2) holds.

For (3), suppose $x \subseteq^+ y_1, \dots, y_n$ in $\mathcal{C}(S)$. Because v is a valuation, we know that

$$v(x) \geq \sum_{\substack{\emptyset \neq I \subseteq \{1, \dots, n\} \\ \cup_i y_i \in \mathcal{C}(S)}} (-1)^{|I|+1} v(\cup_{i \in I} y_i). \quad (3.1)$$

We check that $\sum_{\substack{\emptyset \neq I \subseteq \{1, \dots, n\} \\ \cup_i y_i \in \mathcal{C}(S)}} (-1)^{|I|+1} v(\cup_{i \in I} y_i \mid x) \leq 1$. If $v(x) = 0$ then $v(\cup_{i \in I} y_i \mid x) = 0$ for every I so the inequality holds directly, and if $v(x) > 0$ then it is obtained by dividing both sides of (3.1) by $v(x)$.

Suppose $x \subseteq^- y$ and $x \subseteq^+ z$, with $y \cup z \in \mathcal{C}(S)$. If $v(x) = 0$, then $v(y) = 0$, and therefore $v(y \cup z \mid y) = v(z \mid x) = 0$. Otherwise, because $z \subseteq^- y \cup z$, we have $v(z) = v(y \cup z)$, so that $v(y \cup z \mid y) = \frac{v(y \cup z)}{v(y)} = \frac{v(z)}{v(x)} = v(z \mid x)$. This proves (4). Condition (5) is straightforward, since $\theta : x \cong_S y$ implies $v(x) = v(y)$. \square

Note that the above defines a one-to-one correspondence between valuations and those conditional valuations satisfying the following property: if $v(y \mid x) = 0$ and $y \subseteq z \subseteq^+ w$ then $v(w \mid z) = 0$.

Conditional valuations are arguably more intuitive than valuations, since they provide a more explicit representation for the probability coefficients assigned to Player's behaviour. They are also marginally more general, since nothing prevents having, say, $v(y \mid x) = 0$ but $v(z \mid y) \geq 0$ for some $x \subseteq^+ y \subseteq^+ z$; such a situation cannot arise from a valuation. (The mismatch is not surprising: it is well-known in probability theory that conditioning on a zero-probability event is not well-defined.)

Nevertheless, we shall see that valuations provide a convenient way of making strategies probabilistic; in particular they allow for a cleaner definition of composition for probabilistic strategies than is possible with conditional valuations. So in the first part of this thesis where only *discrete* probability is considered, we stick to valuations, albeit making informal use of the correspondence with conditional valuations, notably when drawing pictures of strategies. ¹

Definition 3.7. A **probabilistic strategy** on a game \mathcal{A} is a strategy $\sigma : \mathcal{S} \rightarrow \mathcal{A}$ together with a valuation v_S on \mathcal{S} .

In the following sections, we study two classes of probabilistic strategies:

- the *Markov* ones, in which any two Player actions are probabilistically independent, conditionally on their causal history.
- the *deterministic* ones, in which the behaviour of Player is fully determined by that of Opponent.

¹Another reason to introduce conditional valuations is the essential role they will play when we move to continuous probability distributions in strategies, in Chapter 7.

3.2 Markov strategies

Recall the example strategy $\tau : \mathcal{T} \rightarrow \mathcal{A}$ at the beginning of 3.1,

$$\begin{array}{ccc} & & b^+ \\ & \nearrow & \\ a^- & & \\ & \searrow & \\ & & c^+ \end{array}$$

A valuation v on T must respect $v(\{a, b\}) + v(\{a, c\}) - v(\{a, b, c\}) \leq 1$, by condition (3) of Definition 3.3. A suitable candidate for $v(\{a, b, c\})$ is the product $v(\{a, b\})v(\{a, c\})$, which always satisfies the above and indicates that events b and c are *probabilistically independent* (conditionally on a). In a Markov strategy, this will be the case for any two Player moves, and indeed for any set of moves:

Definition 3.8. A probabilistic strategy $\sigma : \mathcal{S} \rightarrow \mathcal{A}$ is **Markov** (we may alternatively say the *valuation* v is Markov) if for any $y, z \in \mathcal{C}(S)$ such that $y \cup z \in \mathcal{C}(S)$, $v(y \cup z)v(y \cap z) = v(y)v(z)$.

This condition can be equivalently written as $v(y \cup z \mid y \cap z) = v(y \mid y \cap z) \cdot v(z \mid y \cap z)$, which says that y and z are probabilistically independent, conditionally on their common history ($y \cap z$).

The following equivalent condition is convenient:

Lemma 3.9. A strategy $\sigma : \mathcal{S} \rightarrow \mathcal{A}$ is Markov if and only if for any $x, x' \in \mathcal{C}(S)$, if $x \text{--}c^s x'$ then $v(x') \cdot v([s]) = v(x) \cdot v([s])$.

In conditional form, the above says $v(x' \mid x) = v([s] \mid [s])$. So informally, a strategy is Markov if the probability of playing an event s only depends on its causal history. (This justifies the name ‘‘Markov’’, as it is reminiscent of the Markov condition for probabilistic graphical models. See [Pea00, KF09] for details.)

Proof. The ‘‘only if’’ direction is clear, taking $y = x$ and $z = [s]$. We show the converse. For $y, z \in \mathcal{C}(S)$ with $y \cup z \in \mathcal{C}(S)$, we show that $v(y \cup z)v(y \cap z) = v(y)v(z)$ by induction on $|y \cup z|$. Assume the equality holds, and suppose $y \text{--}c^s y'$ with $s \notin z$.

Assume $v(y) \neq 0$; otherwise the statement holds directly. By assumption, since we have both $y \text{--}c^s y'$ and $y \cup z \text{--}c^s y \cup z$, we have $v(y' \mid y) = v([s] \mid [s]) = v(y' \cup z \mid y \cup z)$, and thus $v(y' \cup z) = v(y \cup z) \cdot \frac{v(y')}{v(y)}$. Note also that $y' \cap z = y \cap z$, so we derive the equation as follows:

$$\begin{aligned} v(y' \cup z)v(y \cap z) &= \frac{v(y')}{v(y)}v(y \cup z)v(y' \cap z) \\ &= \frac{v(y')}{v(y)}v(y \cup z)v(y \cap z) \\ &= \frac{v(y')}{v(y)}v(y)v(z) = v(y')v(z). \end{aligned}$$

□

Finally, we prove a factorisation result:

Lemma 3.10. *A strategy $\sigma : \mathcal{S} \rightarrow \mathcal{A}$ is Markov if and only if for all $x \in \mathcal{C}(S)$,*

$$v(x) = \prod_{\substack{e \in x \\ \text{pol}(e)=+}} v([e] \mid [e]).$$

Proof. The (if) direction is an easy verification. For the (only if) direction, suppose σ is Markov. The argument is by induction on x . The property holds for $x = \emptyset$ since $v(\emptyset) = 1$ by definition.

Suppose the property holds for some $x \in \mathcal{C}(S)$ and suppose $x \prec^s x'$. If s is negative then the property holds since $v(x) = v(x')$ and the positive events of x are exactly those of x' .

Suppose s is positive. Notice that $v(x') = v(x' \mid x) \cdot v(x)$. By Lemma 3.9 $v(x' \mid x) = v([s] \mid [s])$, so using the induction hypothesis for x we get

$$v(x') = v([s] \mid [s]) \times \prod_{\substack{e \in x \\ \text{pol}(e)=+}} v([e] \mid [e]) = \prod_{\substack{e \in x' \\ \text{pol}(e)=+}} v([e] \mid [e]).$$

□

Markov strategies are *not* stable under composition, because the probabilistic independence condition is not closed under hiding: two causally independent events may have a hidden “common cause”, may not be probabilistic independent.

However, the interaction $\mathcal{T} \otimes \mathcal{S}$ satisfies the Markov property:

Lemma 3.11. *Let $x \in \mathcal{C}(T \otimes S)$ and suppose $x \prec^s x'$. Then,*

$$v_{T \otimes S}(x') \cdot v_{T \otimes S}([s]) = v_{T \otimes S}(x) \cdot v_{T \otimes S}([s]).$$

Proof. If s is negative, then $v_{T \otimes S}(x) = v_{T \otimes S}(y)$ and $v_{T \otimes S}([s]) = v_{T \otimes S}([s])$ by Lemma 3.18 and we are done.

Suppose s is a σ -action, and write $x = x_T \otimes x_S$ and $x' = x'_T \otimes x'_S$. Write $[s] = z_T \otimes z_S$ and $[s] = z'_T \otimes z'_S$. As $x_T \subseteq^- x'_T$, $v_T(x_T) = v_T(x'_T)$, and similarly $z_T \subseteq^- z'_T$ so $v_T(z_T) = v_T(z'_T)$. Since S is Markov, $v_S(x'_S) \cdot v_S(z_S) = v_S(x_S) \cdot v_S(z'_S)$, and therefore

$$\begin{aligned} v_{T \otimes S}(x') \cdot v_{T \otimes S}([s]) &= v_T(x'_T) \cdot v_S(x'_S) \cdot v_T(z_T) \times v_S(z_S) \\ &= v_T(x_T) \cdot v_S(x_S) \cdot v_T(z'_T) \times v_S(z'_S) \\ &= v_{T \otimes S}(x) \cdot v_{T \otimes S}([s]). \end{aligned}$$

The proof for s a τ -action is symmetric. □

The Markov property is a sub-condition of the *sequential innocence* condition which we define in the next chapter. The lemma above is a necessary step in the proof that sequential innocent strategies are stable under composition. We will see that those strategies are sufficiently constrained for the hiding problem to disappear.

We continue with a discussion of *deterministic* strategies.

3.3 Determinism and race-freeness

3.3.1 Deterministic strategies

In a deterministic strategy, Player behaves well-enough that simply assigning probability 1 to every configuration yields a valid probabilistic strategy. Note first that this is not always the case. Recall the example strategy $\sigma : \mathcal{S} \rightarrow \mathcal{A}$ above

$$a^- \begin{array}{c} \nearrow b^+ \\ \searrow \\ \triangleleft \\ \triangleright c^+ \end{array}$$

There, assigning 1 to both $\{a, b\}$ and $\{a, c\}$ would violate condition (3) of valuations.

The issue here is that \mathcal{S} contains a minimal conflict between two Player moves, b and c . So we define *deterministic* strategies to be those in which this situation does not happen, and the behaviour of Player is in some sense completely determined by that of Opponent.

Definition 3.12. A strategy $\sigma : \mathcal{S} \rightarrow \mathcal{A}$ is **deterministic** when for every $x \in \mathcal{C}(S)$, if $x \dashv\vdash^s$ and $x \dashv\vdash^{s'}$, with $\text{pol}(s) = +$, then $x \cup \{s, s'\} \in \mathcal{C}(S)$.

Lemma 3.13. *If $\sigma : \mathcal{S} \rightarrow \mathcal{A}$ is deterministic, then the map $v : \mathcal{C}(S) \rightarrow [0, 1]$ as $v(x) = 1$ for every $x \in \mathcal{C}(S)$ satisfies the axioms for a valuation.*

Proof. We check that v satisfies all three axioms. The first two are straightforward, since by definition $v(\emptyset) = 1$ and $v(x) = v(y)$ whenever $x \subseteq^- y$. If $x \subseteq^+ y_1, \dots, y_n$, then because σ is deterministic, $\cup_{i \in I} y_i \in \mathcal{C}(S)$ for every $I \subseteq \{1, \dots, n\}$. So

$$\begin{aligned} v(x) - \sum_{\substack{\emptyset \neq I \subseteq \{1, \dots, n\} \\ \cup_{i \in I} y_i \in \mathcal{C}(S)}} (-1)^{|I|+1} v(\cup_{i \in I} y_i) &= 1 - \sum_{\emptyset \neq I \subseteq \{1, \dots, n\}} (-1)^{|I|+1} \\ &= \sum_{I \subseteq \{1, \dots, n\}} (-1)^{|I|} = 0, \end{aligned}$$

where the last equality is a well-known combinatorial result. \square

The converse to Lemma 3.13 is false due to the possible presence of *races* in games and strategies. By *race* we mean a minimal conflict between a Player move and an Opponent move, as in the following game \mathcal{B} :

$$a^+ \rightsquigarrow b^-$$

Races do not mix well with probability: take the strategy $\sigma : \mathcal{S} \rightarrow \mathcal{B}$ where $\mathcal{S} = \mathcal{B}$ and σ is the identity. Although $v(\{a\})$ can be set to any $p \in [0, 1]$, the operational behaviour is ambiguous: Player does not have control over the move b and may not be able to play a at all.

Thus we introduce the class of “race-free” games [Win].

3.3.2 Race-free games

Definition 3.14. An event structure with polarity E is **race-free** if for all $x \in \mathcal{C}(E)$, whenever $x \sqsubseteq^+ y$ and $x \sqsubseteq^- z$ then y and z are compatible: $y \cup z \in \mathcal{C}(E)$.

By extension we call race-free the games and strategies whose underlying esp is race-free.

Lemma 3.15. *If \mathcal{A} is a race-free game, then any strategy $\sigma : \mathcal{S} \rightarrow \mathcal{A}$ is race-free.*

Proof. Suppose $x \in \mathcal{C}(S)$ with $x \text{--} \text{C}^+ y$ and $x \text{--} \text{C}^- z$. The map σ is defined on events, so $\sigma x \text{--} \text{C}^+ \sigma y$ and $\sigma x \text{--} \text{C}^- \sigma z$. By assumption, \mathcal{A} is race-free so $\sigma y \cup \sigma z = \sigma(y \cup z) \in \mathcal{C}(A)$. Note that $\sigma y \text{--} \text{C}^- \sigma(y \cup z)$ so by receptivity of σ there is a unique $w \in \mathcal{C}(S)$ such that $y \text{--} \text{C}^- w$ and $\sigma w = \sigma(y \cup z)$.

We show that $w = y \cup z$, so $y \cup z \in \mathcal{C}(S)$ as required. We write $y = x \cup \{s^+\}$, $z = x \cup \{s'^-\}$, and $w = c \cup \{s^+, t^-\}$; we must show $t = s'$. By courtesy, $s \vdash t$, and therefore $x \cup \{t\} \in \mathcal{C}(S)$. But $\sigma t = \sigma s'$ since $\sigma w = \sigma(y \cup z)$, and by receptivity there is a unique x' with $x \subseteq x'$ such that $\sigma x' = \sigma z$. Thus $x \cup \{t\} = z$, *i.e.* $s' = t$, and in particular $w = y \cup z$. \square

It is stated in [Win12, Lemma 2] that a game \mathcal{A} is race-free if and only if the copycat strategy $\alpha_{\mathcal{A}} : \mathbb{C}_{\mathcal{A}} \rightarrow \mathcal{A}^\perp \parallel \mathcal{A}$ is deterministic. This will be required for copycat to be the identity morphism in the bicategory we construct below.

Lemma 3.16. *For a race-free game \mathcal{A} , a strategy $\sigma : \mathcal{S} \rightarrow \mathcal{A}$ is deterministic if and only if the function $v : \mathcal{C}(S) \rightarrow [0, 1] : x \mapsto 1$ is a valuation.*

Proof. The (only if) direction is Lemma 3.13, so we check the converse. Suppose $\sigma : \mathcal{S} \rightarrow \mathcal{A}$ is not deterministic, so there exist configurations $z \text{--} \text{C}^+ y_1$ and $z \text{--} \text{C}^+ y_2$ with $y_1 \cup y_2 \notin \text{Con}$. But \mathcal{S} is race-free by Lemma 3.15, so $z \text{--} \text{C}^+ y_2$. By the valuation axioms, $v(z) \geq v(y_1) + v(y_2)$. It is clear that one cannot have $v(x) = 1$ for all $x \in \mathcal{C}(S)$. \square

The probabilistic copycat strategy is defined accordingly.

Definition 3.17. The **probabilistic copycat strategy** on a (race-free) game \mathcal{A} is the copycat strategy $\alpha_{\mathcal{A}} : \mathbb{C}_{\mathcal{A}} \rightarrow \mathcal{A}^\perp \parallel \mathcal{A}$ equipped with the constantly 1 valuation $v : \mathcal{C}(\mathbb{C}_{\mathcal{A}}) \rightarrow [0, 1]$.

So, from now on, we assume all games are race-free, and the rest of the chapter is devoted to the construction of a cartesian closed bicategory of probabilistic strategies on race-free games.

3.4 The bicategory PG

3.4.1 Composition of probabilistic strategies

Suppose $\sigma : \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$ and $\tau : \mathcal{T} \rightarrow \mathcal{B}^\perp \parallel \mathcal{C}$ are probabilistic strategies with valuations $v_{\mathcal{S}}$ and $v_{\mathcal{T}}$ respectively. Their **interaction** is $\tau \circledast \sigma : \mathcal{T} \circledast \mathcal{S} \rightarrow \mathcal{A} \parallel \mathcal{B} \parallel \mathcal{C}$, and we can define a function $v_{\mathcal{T} \circledast \mathcal{S}}$ on configurations of $\mathcal{T} \circledast \mathcal{S}$ as

$$v_{\mathcal{T} \circledast \mathcal{S}}(x_{\mathcal{T}} \circledast x_{\mathcal{S}}) = v_{\mathcal{S}}(x_{\mathcal{S}}) \times v_{\mathcal{T}}(x_{\mathcal{T}})$$

for every $x_T \circledast x_S \in \mathcal{C}(T \circledast S)$. This satisfies the axioms for a valuation, treating neutral events as positive. For $x, y \in \mathcal{C}(T \circledast S)$, we write $x \subseteq^{+,0} y$ to mean that the extension contains only positive and neutral events.

Lemma 3.18 ([Win13a]). *The map $v_{T \circledast S} : \mathcal{C}(T \circledast S) \rightarrow [0, 1]$ satisfies the following properties:*

- $v_{T \circledast S}(\emptyset) = 1$;
- $v_{T \circledast S}(x) = v_{T \circledast S}(y)$ if $x \subseteq^- y$; and
- for $x \subseteq^{+,0} y_1, \dots, y_n$, $v_{T \circledast S}(x) - \sum_I (-1)^{|I|+1} v_{T \circledast S}(\bigcup_{i \in I} y_i) \geq 0$, where I ranges over nonempty subsets of $\{1, \dots, n\}$ such that $\bigcup_{i \in I} y_i$ is a configuration.

Then, the **composition** of σ and τ is $\tau \odot \sigma : \mathcal{T} \odot \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{C}$, equipped with the valuation $v_{T \odot S}$ defined as

$$v_{T \odot S}(x) = v_{T \circledast S}([x])$$

for every $x \in \mathcal{C}(T \odot S)$, or equivalently,

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

for every $x_T \odot x_S \in \mathcal{C}(T \odot S)$. Showing that this satisfies the axioms of a valuation is easily deduced from Lemma 3.18 [Win13a].

Remark 3.19. Chapters 6 and 7 of this thesis contain the development of a games model for *continuous* probability distributions in a measure-theoretical setting. The following points are worth clarifying.

- We sometimes refer to the probabilistic model presented here as “discrete”, since coefficients are assigned directly to elements of $\mathcal{C}(S)$, in a way that is reminiscent of probabilistic reasoning in discrete spaces (such as countable sets). But this terminology is short-sighted, and somewhat misleading, since event structures can be infinite, and real numbers can be encoded as infinite sequences of events; then, by assigning coefficients to finite sequences one is able to model all probability distributions on the reals. This is briefly discussed in [Win]. We leave as further work the comparison of this approach with the measure-theoretic framework of Chapters 6 and 7.
- The “discrete” model arises as a sub-bicategory of the probabilistic games model of Chapters 6 and 7. Therefore, to avoid redundancy, the structural proofs are omitted in this chapter. In particular we make no mention of the ambient pseudo-double category, and build a bicategory directly. We do state all definitions and results; they will provide some intuition for the general case, and suffice for the time being. (In the next two chapters, on the untyped λ -calculus and PCF, only discrete distributions are considered.)

3.4.2 2-Cells

The 2-cells in the bicategory of arenas and probabilistic strategies are maps of strategies with an additional property relating valuations.

Definition 3.20. A map of probabilistic strategies from $\sigma : \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$ to $\sigma' : \mathcal{S}' \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$ is a map $f : \sigma \Rightarrow \sigma'$ satisfying

$$v_{\mathcal{S}}(x) \leq v_{\mathcal{S}'}(fx)$$

for every $x \in \mathcal{C}(\mathcal{S})$.

Lemma 3.21. *Maps of probabilistic strategies are stable under vertical and horizontal composition.*

Proof. For vertical composition, notice that if $f : \sigma \Rightarrow \sigma'$ and $f' : \sigma' \Rightarrow \sigma''$, then $v_{\mathcal{S}}(x) \leq v_{\mathcal{S}'}(fx) \leq v_{\mathcal{S}''}((f' \circ f)x)$.

For horizontal composition, let $f : \sigma \Rightarrow \sigma'$ and $g : \tau \Rightarrow \tau'$. If $x_T \odot x_S \in \mathcal{C}(T \odot S)$, then $(g \odot f)(x_T \odot x_S)$ is a configuration of $T' \odot S'$, equal to $y_{T'} \odot y_{S'}$ for some causally compatible pair $(y_{S'}, y_{T'}) \in \mathcal{C}(S') \times \mathcal{C}(T')$, where there are bijections $\varphi_{S'} : fx_S \cong_{S'} y_{S'}$ and $\varphi_{T'} : gx_T \cong_{T'} y_{T'}$. Thus we have $v_{T \odot S}(x_T \odot x_S) = v_S(x_S)v_T(x_T) \leq v_{S'}(fx_S)v_{T'}(gx_T) = v_{S'}(y_{S'})v_{T'}(y_{T'}) = v_{T' \odot S'}(y_{T'} \odot y_{S'}) = v_{T' \odot S'}((g \odot f)(x_T \odot x_S))$. \square

This gives a bicategory.

Theorem 3.22. *There is a bicategory **PG** having:*

- *objects: race-free arenas;*
- *morphisms $\mathcal{A} \rightarrow \mathcal{B}$: probabilistic strategies;*
- *2-cells: maps of probabilistic strategies.*

3.5 Rigid maps and push-forward valuations

This section is an aside. We describe how *rigid* maps of event structures can be used to transport a valuation from one strategy to another.

Definition 3.23. A map of event structures $f : S \rightarrow T$ is **rigid** if it preserves causality, *i.e.* if $s \leq_S s'$ then $f(s) \leq_T f(s')$.

Rigid maps will play a fundamental role in the theory of measurable event structures to be introduced in the second part of this thesis, because they correspond precisely to those maps for which the induced map on configurations is a *discrete fibration*.

For now we simply observe that rigid maps of strategies have the convenient property that one can *push-forward* valuations across them. The concept is standard in probability theory, where given a function $f : X \rightarrow Y$ one can turn a probability distribution on X into one on Y (provided f is “measurable”).

Here, given a map of strategies $f : \sigma \Rightarrow \tau$ and a valuation v on S , define a map $f_*v : \mathcal{C}(T) \rightarrow [0, 1]$ as

$$(f_*v)(y) = \sum_{x \in f^{-1}\{y\}} v(x),$$

for $y \in \mathcal{C}(T)$, where $f^{-1}\{y\}$ is the pre-image of y under f as a function of configurations, *i.e.* the set $\{x \in \mathcal{C}(S) \mid fx = y\}$. We show that when the underlying $f : S \rightarrow T$ is rigid, f_*v satisfies the axioms for a valuation. A special case of this result is known, for rigid maps of basic strategies without symmetry [Win15].

To handle the addition of symmetry the following technical lemma is key:

Lemma 3.24. *Let $f : \sigma \Rightarrow \tau$ be a map of strategies, and let $x, y \in \mathcal{C}(T)$ with $\theta : x \cong_T y$. Then there exists a bijective function $\alpha : f^{-1}\{x\} \rightarrow f^{-1}\{y\}$ equipped with a symmetry $\phi_z : z \cong_S \alpha(z)$ for every $z \in f^{-1}\{x\}$.*

Proof. The proof is by induction on the size of θ , where the case $\theta = \text{id}_\emptyset$ is straightforward.

Suppose $\theta \text{---} \text{c}^+ \theta' : x' \cong_T y'$ and we are given α with symmetries ϕ_z as above. If $z \in f^{-1}\{x\}$ and $z \subseteq z' \in f^{-1}\{x'\}$, since $\phi_z : z \cong_S \alpha(z)$, there is $w \in \mathcal{C}(S)$ with a bijection $\phi_{z'} : z' \cong_S w$ extending ϕ_z . It must be the case that $fw = y'$, because from $f\phi_{z'} : x' \cong_T fw$ and $\theta' : x' \cong_T y'$ we get a symmetry $f\phi_{z'} \circ \theta'^{-1} : y' \cong_T fw$ which extends id_y positively, and \mathcal{T} is thin. So we define $\alpha'(z') = w$. Since $f\phi_z = \theta$ and f is defined on events, we have $f\phi_{z'} = \theta'$.

The resulting $\alpha' : f^{-1}\{x'\} \rightarrow f^{-1}\{y'\}$ is injective: assume $\alpha'(z') = \alpha'(z'')$ for some $z', z'' \in f^{-1}\{x'\}$. Assume z' and z'' are extensions of $z_0, z_1 \in f^{-1}\{x\}$, respectively; then the bijections $\phi_{z_0} : z_0 \cong_S \alpha(z_0)$ and $\phi_{z_1} : z_1 \cong_S \alpha(z_1)$ extend to $\phi_{z'} : z' \cong_S \alpha'(z')$ and $\phi_{z''} : z'' \cong_S \alpha'(z'')$, so in particular $\alpha(z_0) = \alpha(z_1)$ and therefore $z_0 = z_1$ (α is bijective). So, writing z for z_0 (and z_1), we have $\text{id}_z \text{---} \text{c}^+ \phi_{z''}^{-1} \circ \phi_{z'}$, which by thinness of \mathcal{S} implies $z' = z''$. That α is surjective uses a similar argument.

Now suppose $\theta \text{---} \text{c}^- \theta' : x' \cong_T y'$. Given that f is a map of strategies, for every $z \in \mathcal{C}(S)$ there are positive symmetries $\psi_z : \sigma z \cong_A^+ \tau x$ and $\psi_{\alpha(z)} : \sigma(\alpha(z)) \cong_A^+ \tau y$. By assumption, $\tau x \subseteq^- \tau x'$ and $\tau y \subseteq^- \tau y'$, which (by Lemma 2.13) determines unique extensions $\sigma z \subseteq^- u$ and $\sigma(\alpha(z)) \subseteq w$ with a symmetry $u \cong_T w$ extending $\sigma\phi_z$. By \sim -receptivity of σ , this can be lifted uniquely to an extension of ϕ_z . \square

From this we can prove the push-forward result:

Lemma 3.25. *Let $f : \sigma \Rightarrow \sigma'$ be a map of strategies such that the underlying map is rigid, and let v be a valuation on \mathcal{S} . Then, the map f_*v is a valuation on \mathcal{T} .*

Proof. We check the four axioms of Definition 3.3. It is shown in [Win] that axioms (1), (2), and (3) hold for f_*v provided f is rigid and *receptive*. To see why the latter holds, consider $z \in \mathcal{C}(S)$ and $fz \subseteq^- y \in \mathcal{C}(T)$. Applying τ we get $\tau(fz) \subseteq^- \tau y$. Recall f is a map of strategies, so by definition there is $\varphi_z : \sigma z \cong_A^+ \tau(fz)$, which by Lemma 2.13 extends uniquely to some $\psi : w \cong_A^+ \tau y$. As $\sigma z \subseteq^- w$, by receptivity of σ we get a unique z' extending z such that $\sigma z' = w$. Because ψ is unique we must have $fz' = y$.

Finally, axiom (4) follows directly from Lemma 3.24: for any $\theta : x \cong_T y$, the equality

$$\sum_{z \in f^{-1}\{x\}} v(z) = \sum_{w \in f^{-1}\{y\}} v(w)$$

is established using that axiom (4) holds for v . \square

Whenever the conditions are met we call f_*v the **push-forward valuation** of v across f .

3.6 Bicategorical structure

The structure of \mathbf{G} extends naturally to \mathbf{PG} . All proofs will be given in Chapter 7 (see 7.4.4).

Lifting. The lifting and co-lifting constructions only produce instances of copycat, and thus deterministic strategies. Observing this, it is not difficult to show that Lemmas 2.27 and 2.60 still hold after the addition of probability. (Of course this requires all games to be race-free, but objects of \mathbf{PG} are race-free games.)

Products. Finite products are the same as in \mathbf{G} ; $\mathcal{A} \& \mathcal{B}$ is race-free whenever \mathcal{A} and \mathcal{B} are. Recall that the pairing of strategies $\sigma : \mathcal{S} \rightarrow \mathcal{C}^\perp \parallel \mathcal{A}$ and $\tau : \mathcal{T} \rightarrow \mathcal{C}^\perp \parallel \mathcal{B}$ is a strategy $\langle \sigma, \tau \rangle : \mathcal{S} \& \mathcal{T} \rightarrow \mathcal{C}^\perp \parallel \mathcal{A} \& \mathcal{B}$. Configurations of $\mathcal{S} \& \mathcal{T}$ are either empty, or fully included in one of the two components. So for $x \in \mathcal{C}(\mathcal{S} \& \mathcal{T})$, we define $v_{\mathcal{S} \& \mathcal{T}}(x)$ to be 1 if $x = \emptyset$, and (abusing notation) $v_{\mathcal{S}}(x)$ or $v_{\mathcal{T}}(x)$ accordingly.

As strategies are negative, the incompatibility between moves of \mathcal{S} and \mathcal{T} is induced by Opponent, so that the above is a well-defined valuation.

Symmetric monoidal closed structure. Similarly, in setting up the monoidal structure, all that needs re-defining is the action of \otimes on strategies: for $\sigma : \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{A}'$ and $\tau : \mathcal{T} \rightarrow \mathcal{B}^\perp \parallel \mathcal{B}'$, their tensor product $\sigma \otimes \tau : \mathcal{S} \parallel \mathcal{T} \rightarrow (\mathcal{A} \parallel \mathcal{B})^\perp \parallel (\mathcal{A}' \parallel \mathcal{B}')$ can be equipped with the valuation $v_{\mathcal{S} \parallel \mathcal{T}} : \mathcal{C}(\mathcal{S} \parallel \mathcal{T}) \rightarrow [0, 1]$ defined by $v_{\mathcal{S} \parallel \mathcal{T}}(x_{\mathcal{S}} \parallel x_{\mathcal{T}}) = v_{\mathcal{S}}(x_{\mathcal{S}}) \times v_{\mathcal{T}}(x_{\mathcal{T}})$. (Showing that this indeed defines a valuation is slightly technical.)

All associated data is obtained in the same way as for \mathbf{G} , using the remarks on lifting above. It must also be checked that the 2-cells involved (the same as in \mathbf{G}) are valid maps of probabilistic strategies. We obtain that \mathbf{PG} is a symmetric monoidal bicategory. It is also closed: since σ and $\Lambda(\sigma)$ have the same internal event structure \mathcal{S} , the valuation remains the same and the adjunction proof goes through with no difficulty.

A linear exponential pseudo-comonad. When $\sigma : \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$ is a probabilistic strategy, we equip $!\sigma : !\mathcal{S} \rightarrow !\mathcal{A}^\perp \parallel !\mathcal{B}$ with a valuation $v_{!S}$, treating $!$ as an “infinitary tensor”:

$$v_{!S}(\parallel_{i \in \omega} x_i) = \prod_{\substack{i \in \omega \\ x_i \neq \emptyset}} v_S(x_i).$$

The rest of the proof is done as in **G**. This gives a cartesian closed bicategory:

Theorem 3.26. *The Kleisli bicategory $\mathbf{PG}_!$ is cartesian closed.*

In the next two chapters, we use this bicategory to give semantics to probabilistic extensions of well-known languages: the untyped λ -calculus (Chapter 4), and PCF (Chapter 5). To interpret recursion in the latter, we will need to consider ω -chains of strategies.

Recursion. As in Chapter 2 we give an explicit description of ω -colimits in the situation where the diagram is made up of *strong inclusions*. In this setting, we additionally assume that the valuation is preserved by the inclusion map: consider a chain $\sigma_0 \xrightarrow{f_0} \sigma_1 \xrightarrow{f_1} \dots$ of strong inclusions and suppose that for every $x \in \mathcal{C}(S_i)$ we have $v_{S_i}(x) = v_{S_{i+1}}(S_{i+1})$. Then the colimit construction given in the previous chapter extends directly: the strategy

$$\bigvee_{i \in \omega} \sigma_i : \bigcup_{i \in \omega} S_i \rightarrow \mathcal{A}$$

is equipped with the valuation $x \mapsto v_{S_i}(x)$, where i is any index such that $x \in \mathcal{C}(S_i)$.

More generally, given a chain consisting of *strong embeddings*, we construct the colimit in the same way as in Chapter 2, using that a strong embedding can always be factored as a strong inclusion composed with an isomorphism.

Chapter 4

Innocence and the untyped λ -calculus

The work presented in this chapter was conducted in collaboration with Pierre Clairambault, and shows how probabilistic concurrent strategies can be applied to a probabilistic extension of the untyped λ -calculus. We begin with a brief account of the scientific context for this work; in particular we mention the various results from which it draws inspiration.

A denotational model for the pure λ -calculus can be obtained by finding a *reflexive object* in a cartesian closed category \mathbb{C} , *i.e.* an object U with morphisms $\text{app} \in \mathbb{C}[U, U \rightarrow U]$ and $\lambda \in \mathbb{C}[U \rightarrow U, U]$ such that $\text{app} \circ \lambda = \text{id}_{U \rightarrow U}$. A λ -term with n free variables can then be interpreted as a morphism $U^n \rightarrow U$, and the induced equational theory is a well-defined λ -theory — a congruence on λ -terms closed under β -reduction. (It is also the case that any λ -theory can be obtained in this way [Sco80].) If app and λ are inverses of each other, the reflexive object is said to be *extensional* and the induced λ -theory also validates the η rule.

One may then consider reflexive objects in cartesian closed *bicategories*. There, the condition on the pair (app, λ) can be relaxed. In [See87, Zei15] it is argued that any adjunction $(U, U \rightarrow U, \text{app}, \lambda)$ leads to a model, but it remains unclear that this definition is completely appropriate: why, in particular, should the adjunction be taken in one direction and not the other? To avoid any ambiguity we consider only the extensional case, in which the adjunction must be an adjoint equivalence, and hence the direction does not matter; see Definition 4.1 below. An example of this is found in the cartesian closed bicategory of generalised species of structure [FGHW08].

Reflexive objects can also be found in categories of games, as described for example in the work of Ker et al [KNO02], carried out in a model of Hyland-Ong games. The reflexive object they consider is an arena \mathcal{U} , the *universal arena*; it additionally satisfies $\mathcal{U} \cong \mathcal{U} \rightarrow \mathcal{U}$, which implies (for abstract reasons [Sco80]) that the induced λ -theory is also closed under η -reduction (the model is said to be *extensional*). But the main result of [KNO02] is much stronger: the λ -theory of \mathcal{U} coincides with \mathcal{H}^* , a well-studied λ -theory with a central place in the model theory of the λ -calculus.

To obtain this result the authors show that there is a precise correspondence between, on one hand, *innocent* strategies, and on the other, *Nakajima trees* for

λ -terms. Innocence is a condition on strategies which plays a significant role in the game semantics of PCF [HO00], while Nakajima trees are syntactic structures used as a representation for λ -terms; they are known to capture precisely the λ -theory \mathcal{H}^* .

The work presented in this chapter follows a similar approach to [KNO02], but focusses on a λ -calculus enriched with an operator $+_p$ for probabilistic choice. We consider a probabilistic head-reduction strategy for the calculus, in the spirit of [EPT11]. The recent PhD thesis of Leventis [Lev16] provides an in-depth analysis of the calculus and its operational semantics. In particular, Leventis puts forward a probabilistic extension of Nakajima trees, and shows they characterise the probabilistic analogue of \mathcal{H}^* .

In this chapter we define game semantics for the probabilistic λ -calculus. Our contributions are organised as follows:

- In Section 4.1, we show that the bicategory $\mathbf{PG}_!$ has a reflexive object \mathcal{U} which additionally provides an interpretation for $+_p$.
- In Section 4.2, we define and investigate “sequential innocence” for probabilistic concurrent strategies. The condition is new, but largely based on innocence for non-probabilistic concurrent strategies given in [CCW15].
- In Section 4.3, we show an adequacy result for the model of sequential innocent strategies, saying that the *probability of convergence* of a λ -term can be obtained from its interpretation as a strategy.
- Finally, in Sections 4.4 and 4.5, we show how the probabilistic Nakajima trees of [Lev16] can be recovered from probabilistic innocent strategies. (Unlike [KNO02], this is not a one-to-one correspondence, due to the presence of additional branching information in probabilistic strategies.)

4.1 Syntax and semantics of the probabilistic λ -calculus

We start with some syntactic background.

4.1.1 Syntax and operational semantics

The set λ^+ of terms of the probabilistic λ -calculus is defined by the following grammar, where p ranges over the interval $[0, 1]$ and x over an infinite set Var :

$$M, N ::= x \mid \lambda x.M \mid MN \mid M +_p N.$$

Write λ^+_0 for the set of **closed terms**, *i.e.* those with no free variables, and λ^+_Γ for the set of terms with free variables in Γ .

The operator $+_p$ represents probabilistic choice, so that a term of the form $M +_p N$ has two possible reduction steps: to M , with probability p , and to N , with

probability $1 - p$. Accordingly, the reduction relation we consider is a Markov process over the set λ^+ .

It corresponds to a probabilistic variant of the standard **head-reduction** [Bar84]. It is defined inductively:

$$\frac{}{(\lambda x.M)N \xrightarrow{1} M[N/x]} \quad \frac{M \neq N}{M +_p N \xrightarrow{p} M} \quad \frac{M \neq N}{M +_p N \xrightarrow{1-p} N}$$

$$\frac{}{M +_p M \xrightarrow{1} M} \quad \frac{M \xrightarrow{p} M'}{\lambda x.M \xrightarrow{p} \lambda x.M'} \quad \frac{M \xrightarrow{p} M' \quad M \neq \lambda x.P}{MN \xrightarrow{p} M'N}$$

For $M, N \in \lambda^+$, there may be many reduction paths from M to N . The **weight** of a path $\pi : M \xrightarrow{p_1} \dots \xrightarrow{p_n} N$ is the product of the transition probabilities: $w(\pi) = \prod_{i=1}^n p_i$. The **probability of M reducing to N** is then defined as $\text{Pr}(M \rightarrow N) = \sum_{\pi: M \rightarrow^* N} w(\pi)$.

Our goal in this section is to give a denotational semantics to λ^+ -terms. We identify a reflexive object in the bicategory $\mathbf{PG}_!$, which gives a canonical interpretation for standard λ -calculus constructions; from this we will get for free a soundness result for the semantics with respect to the usual $\beta\eta$ -equality.

We will also give an explicit semantics for the operator $+_p$, but for this the meaning of this interpretation is not clear. It is not until Section 4.4 that we will get a satisfactory answer to this question, through a comparison of strategies and Böhm trees, only possible after the introduction of innocence in Section 4.2.

4.1.2 A reflexive object in $\mathbf{PG}_!$

As discussed in the introduction of this chapter, we define:

Definition 4.1. An **extensional reflexive object** in a cartesian closed bicategory \mathbb{C} is an object U equipped with an adjoint equivalence

$$\begin{array}{ccc} & \text{app} & \\ & \curvearrowright & \\ U & & U \rightarrow U \\ & \curvearrowleft & \\ & \lambda & \end{array}$$

We define an arena \mathcal{U} , starting with the underlying esp U . It has

- *events*: $(\mathbb{N} \times \mathbb{N})^*$, finite sequences of ordered pairs;
- *causality*: $s \leq t$ if s is a prefix of t ;
- *consistency*: no conflicts, $\text{Con}_U = \mathcal{P}_f(U)$;
- *polarity*: $\text{pol}_U(s) = -$ if $|s|$ is even, $+$ if it is odd.

So U is an infinite, alternating tree whose nodes are labelled with pairs of natural numbers (i, n) . In (i, n) , the integer i is the *copy index* of the node: the symmetry on U will reflect this. (The computational meaning of n is more subtle and will only become clear through the comparison with Nakajima trees.)

Accordingly, the isomorphism families on U are defined via equivalence relations on events. We define \sim , \sim_+ and \sim_- to be the smallest equivalence relations satisfying:

$$\begin{aligned} s \sim s' &\implies s \cdot (i, n) \sim s' \cdot (j, n) \\ s \sim^p s' &\implies s \cdot (i, n) \sim s' \cdot (i, n) \quad (\text{for } p \in \{+, -\}) \\ s \sim_+ s' \text{ and } |s| \text{ even} &\implies s \cdot (i, n) \sim_+ s' \cdot (j, n) \\ s \sim_- s' \text{ and } |s| \text{ odd} &\implies s \cdot (i, n) \sim_- s' \cdot (j, n) \end{aligned}$$

Then, for $x, y \in \mathcal{C}(U)$, a bijection $\theta : x \cong y$ is in \cong_U if for every $e \in x$, $e \sim \theta(e)$. The families \cong_U^- and \cong_U^+ are defined similarly using \sim_- and \sim_+ . The condition on $|s|$ in the definition of \sim_+ ensures that only Player copy indices are altered by a bijection in \cong_U^+ ; and the same holds for Opponent and \cong_U^- .

This data forms a game with symmetry:

Lemma 4.2. $\mathcal{U} = (U, \cong_U, \cong_U^-, \cong_U^+)$ is an arena.

Proof. That \mathcal{U} is forest-shaped and alternating holds by construction. We check that \cong_U , \cong_U^- and \cong_U^+ are isomorphism families on U and satisfy the axioms of Definition 2.12.

(\cong_U is an iso family). Because \sim is an equivalence relation, it is easy to check that: the identity bijection on x is in \cong_U ; if $\theta : x \cong_U y$ then $\theta^{-1} : y \cong_U x$; and if $\theta : x \cong_U y$ and $\psi : y \cong_U z$ then $\psi \circ \theta : x \cong_U z$. Suppose that $\theta : x \cong_U y$ and let $x' \subseteq x$. In particular we have that $e \sim \theta(e)$ for any $e \in x'$ and therefore the restriction of θ to x' is in \cong_U . Suppose now that $x \dashv^s x'$ for some s . If $x = \emptyset$ then $s = \varepsilon$ (the empty sequence) and θ extends trivially. If x is nonempty, let t be the unique predecessor of s in x ; write $s = t \cdot (i, n)$. By assumption, $t \sim \theta(t)$, so for any j we have $s \sim \theta(t) \cdot (j, n)$ and since y is finite, j can be chosen so that $s' \notin y$. So writing $y' = y \cup \{s'\}$, there is an extension $\theta \subseteq \theta' : x' \cong_U y'$.

The proofs for \cong_U^+ and \cong_U^- are similar, with an added subtlety when checking the extension axiom: suppose $\theta x \cong_U^+ y$ and $x \dashv^s x'$, where $s = t \cdot (i, n)$ for some $t \in x$ with $|t|$ odd. By definition, $s \sim_+ s'$ where $s' = \theta(t) \cdot (i, n)$. For the axiom to hold with $y' = y \cup \{s'\}$, it remains to check $s' \notin y$. But if $s' \in y$, $\theta^{-1}(s') \in x$ and because θ is an order-isomorphism, $\theta^{-1}(s')$ is a successor of t . By definition of \sim_+ , $\theta^{-1}(s') = t \cdot (i, n) = s$ which was assumed not to be in x .

($\cong_U^+ \subseteq \cong_U$ and $\cong_U^- \subseteq \cong_U$). This follows from the fact that the equivalence relations \sim_+ and \sim_- are subsets of \sim .

(If $\theta \in \cong_U^+ \cap \cong_U^-$ then θ is an identity bijection). If $\theta : x \cong y$ is in $\cong_U^+ \cap \cong_U^-$, then any $e \in x$ satisfies $e \sim_+ \theta(e)$ and $e \sim_- \theta(e)$. An easy inductive argument shows that $e = \theta(e)$.

(If $\theta \in \cong_U^+$ and $\theta \subseteq^+ \theta' \in \cong_U$ then $\theta' \in \cong_U^+$). It is enough to show $\theta' \in \cong_U^+$ in the case $\theta' = \theta \cup \{(s, s')\}$. We know that $|s| = |s'|$ is odd, since s and s' are positive events. We need to show that $s \sim^+ s'$. Let t and t' be the immediate predecessors of

s and s' respectively, so that $t \sim^+ t t'$ and $s = t \cdot (i, n)$ and $s = t' \cdot (j, m)$. Since $s \sim s'$, we have $n = m$, and because $|t| = |t'|$ is even, $s \sim^+ s'$. Checking the last axiom uses a similar argument. \square

The object \mathcal{U} , called the **universal arena**, is an extensional reflexive object in $\mathbf{PG}_!$: the adjoint equivalence of Definition 4.1 is obtained by lifting an isomorphism $\mathcal{U} \cong \mathcal{U} \rightarrow \mathcal{U}$ of event structures with symmetry and polarity. By Lemma 2.60 this gives an adjoint equivalence in \mathbf{PG} , which is sent to one in $\mathbf{PG}_!$ by the (canonical) pseudo-functor $\mathbf{PG} \rightarrow \mathbf{PG}_!$.

Lemma 4.3. *The essps \mathcal{U} and $\mathcal{U} \rightarrow \mathcal{U}$ are isomorphic.*

Proof. First we show that U and $U \rightarrow U$ are isomorphic as event structures. Recall the function space construction on games: $U \rightarrow U$ has events those of $!U \parallel U$, which for clarity we write as $(\{\alpha\} \times \mathbb{N} \times U) \cup (\{\beta\} \times U)$. The causality relation is defined as the transitive closure of $\leq_{!U \parallel U} \cup \{((\beta, \varepsilon), (\alpha, i, \varepsilon)) \mid i \in \mathbb{N}\}$ (where $\varepsilon \in U$ is the empty sequence).

The map

$$\begin{aligned} \Psi : U \rightarrow U &\longrightarrow U \\ (\alpha, i, s) &\longmapsto (i, 0) \cdot s \\ (\beta, \varepsilon) &\longmapsto \varepsilon \\ (\beta, (i, n) \cdot s) &\longmapsto (i, n + 1) \cdot s \end{aligned}$$

is a bijection on events, with inverse

$$\begin{aligned} \Psi^{-1} : U \rightarrow U \rightarrow U \\ \varepsilon &\longmapsto (\beta, \varepsilon) \\ (i, n) \cdot s &\longmapsto \begin{cases} (\alpha, i, s) & \text{if } n = 0 \\ (\beta, (i, n - 1) \cdot s) & \text{otherwise.} \end{cases} \end{aligned}$$

Ψ is an order-isomorphism: if $e < e'$ in $U \rightarrow U$, then either $e = (\beta, \varepsilon)$, or $e = (\alpha, i, s)$ and $e' = (\alpha, i, s')$ with $s <_U s'$, or $e = (\beta, s)$ and $e' = (\beta, s')$ with $s <_U s'$. In all cases $\Psi(e) < \Psi(e')$. If $s < s'$ in U , then s is a prefix of s' , and it is easy to check that $\Psi^{-1}(e) < \Psi^{-1}(e')$. Checking that Ψ preserves polarity is straightforward. Thus U and $U \rightarrow U$ are isomorphic as essps.

The bijection Ψ preserves symmetry. Suppose that $\theta : x \cong_{U \rightarrow U} y$, and write $x = \{\alpha\} \times (\bigcup_{i \in \mathbb{N}} \{i\} \times x_{(\alpha, i)}) \cup \{\beta\} \times x_\beta$ and $y = \{\alpha\} \times (\bigcup_{i \in \mathbb{N}} \{i\} \times y_{(\alpha, i)}) \cup \{\beta\} \times y_\beta$. By definition of $\cong_{U \rightarrow U}$, there is a reindexing bijection $\pi : \mathbb{N} \rightarrow \mathbb{N}$ and bijections $\theta_{(\alpha, i)} \in \cong_U$, $i \in \mathbb{N}$ and $\theta_\beta \in \cong_U$ such that $\theta(\alpha, i, s) = (\alpha, \pi(i), \theta_{(\alpha, i)}(s))$ and $\theta(\beta, s) = (\beta, \theta_\beta(s))$ for every s . Now for any $e \in x$, we must show that $\Psi(\theta(e)) \sim \Psi(e)$. If $e = (\alpha, i, s)$, then $\theta(e) = (\alpha, \pi(i), \theta_{(\alpha, i)}(s))$, and $\Psi(\theta(e)) = (0, \pi(i)) \cdot \theta_{(\alpha, i)}(s)$. But since $\theta_{(\alpha, i)} \in \cong_U$, we have that $s \sim \theta_{(\alpha, i)}(s)$ and therefore $\Psi(\theta(e)) \sim (i, 0) \cdot s = \Psi(e)$. Similarly if $e = (\beta, (i, n) \cdot s)$ (the case $e = (\beta, \varepsilon)$ is trivial), then $\theta(e) = (\beta, \theta_\beta((i, n) \cdot s))$. We know that $(i, n) \cdot s \sim \theta_\beta((i, n) \cdot s)$ and therefore $\theta_\beta((i, n) \cdot s) = (n, m') \cdot s'$ for some m', s' such that $s \sim s'$. We have $\Psi(e) = (n + 1, m) \cdot s \sim (n + 1, m') \cdot s' = \Psi(\theta(e))$. This concludes the proof that $\Psi\theta \in \cong_U$. We omit the proof the Ψ^{-1} also preserves symmetry, which can be verified similarly. \square

The isomorphism can be lifted to an adjoint equivalence in $\mathbf{PG}_!$:

Corollary 4.4. *The arena \mathcal{U} is an extensional reflexive object in $\mathbf{PG}_!$.*

Proof. By Lemma 4.3, there is an iso $\Psi : \mathcal{U} \rightarrow \mathcal{U} \rightarrow \mathcal{U}$. Applying Lemma 2.60, we get an adjoint equivalence

$$\begin{array}{ccc} & \xrightarrow{\widehat{\Psi^{-1}}} & \\ U & \perp & U \rightarrow U \\ & \xleftarrow{\widehat{\Psi}} & \end{array}$$

in \mathbf{PG} . The canonical pseudo-functor $\mathbf{PG} \rightarrow \mathbf{PG}_!$ is identity-on-objects and preserves adjoint equivalences, which implies the result. (Unfolding this, the adjoint equivalence

$$\begin{array}{ccc} & \xrightarrow{\text{app}} & \\ \mathcal{U} & \perp & \mathcal{U} \rightarrow \mathcal{U} \\ & \xleftarrow{\lambda} & \end{array}$$

in $\mathbf{PG}_!$ is given by $\text{app} = \widehat{\Psi^{-1}} \odot \varepsilon_{\mathcal{U}}$ and $\lambda = \widehat{\Psi} \odot \varepsilon_{\mathcal{U} \rightarrow \mathcal{U}}$. □

4.1.3 Interpretation of λ^+

Closed terms of the probabilistic λ -calculus are interpreted as probabilistic strategies on \mathcal{U} . Open terms M with free variables in Γ are interpreted as strategies $\llbracket M \rrbracket^\Gamma \in \mathbf{PG}_! [\mathcal{U}^\Gamma, \mathcal{U}]$, where $\mathcal{U}^\Gamma = \&_{x \in \Gamma} \mathcal{U}$. The interpretation of the λ -calculus constructions is standard:

$$\begin{aligned} \llbracket x \rrbracket^\Gamma &= \varpi_x, \text{ the } x^{\text{th}} \text{ projection} \\ \llbracket \lambda x.M \rrbracket^\Gamma &= \lambda \odot \Lambda(\llbracket M \rrbracket^{\Gamma, x}) \\ \llbracket MN \rrbracket^\Gamma &= \text{Ev}_{\mathcal{U}, \mathcal{U}} \odot \langle \text{app} \odot \llbracket M \rrbracket^\Gamma, \llbracket N \rrbracket^\Gamma \rangle \end{aligned}$$

For the probabilistic choice operator, we define the sum of two strategies. Let $\sigma : \mathcal{S} \rightarrow (\mathcal{U}^\Gamma)^\perp \parallel \mathcal{U}$ and $\tau : \mathcal{T} \rightarrow (\mathcal{U}^\Gamma)^\perp \parallel \mathcal{U}$ be strategies, and let $p \in [0, 1]$. The $\text{essp } \mathcal{S} +_p \mathcal{T}$ has a unique initial Opponent move (as do \mathcal{S} and \mathcal{T} — *wlog* call this move ε), and continues as either \mathcal{S} or \mathcal{T} non-deterministically. That is, it has events $\{\varepsilon\} \uplus (S \setminus \{\varepsilon\}) \uplus (T \setminus \{\varepsilon\})$, and all structure induced from \mathcal{S} and \mathcal{T} , with $X \in \text{Con}_{\mathcal{S} +_p \mathcal{T}}$ iff $X \in \text{Con}_{\mathcal{S}}$ or $X \in \text{Con}_{\mathcal{T}}$. We define $v_{\mathcal{S} +_p \mathcal{T}}(x)$ to be 1 if $x = \emptyset, \{\varepsilon\}$, $pv_{\mathcal{S}}(x)$ if $x \in \mathcal{C}(\mathcal{S})$, and $(1-p)v_{\mathcal{T}}(x)$ if $x \in \mathcal{C}(\mathcal{T})$. The obvious map $\sigma +_p \tau : \mathcal{S} +_p \mathcal{T} \rightarrow (\mathcal{U}^\Gamma)^\perp \parallel \mathcal{U}$ is a strategy, and the interpretation of the syntactic $+_p$ is simply $\llbracket M +_p N \rrbracket^\Gamma = \llbracket M \rrbracket^\Gamma +_p \llbracket N \rrbracket^\Gamma$.

For general reasons [Sco80], because the reflexive object $(\mathcal{U}, \text{app}, \lambda)$ is extensional, the semantics validates the β and η equations: for any $M, N \in \boldsymbol{\lambda}^+$,

$$\llbracket (\lambda x.M)N \rrbracket \cong \llbracket M[N/x] \rrbracket \quad \text{and} \quad \llbracket \lambda x.Mx \rrbracket \cong \llbracket M \rrbracket.$$

However more work needs to be done in order to make sense of the probabilistic interpretation. In the next section, we define *sequential innocence* for probabilistic concurrent strategies. Then (in Section 4.3) we will prove an adequacy result relating syntax and semantics, stating that the *probability of convergence* of a term (this will be defined) can be recovered from the first level of Player moves in the (sequential innocent) strategy representing it.

4.2 Probabilistic innocence

Innocence is a condition on strategies which captures definability by a purely functional program; we will see (in this chapter and the next) the precise sense in which this holds. In ‘traditional’ game semantics, innocent strategies are those in which Player’s behaviour at any point of the game depends only on a part of the current execution trace called the *P-view*. In concurrent game semantics, the explicit causal dependency relation allows for a clean definition of innocence, in which *P-views* are replaced by *grounded causal chains*.

Definition 4.5. A **grounded causal chain (gcc)** in a strategy $\sigma : \mathcal{S} \rightarrow \mathcal{A}$ is a set of events $\rho = \{s_0, \dots, s_n\} \subseteq S$ such that $s_0 \rightarrow \dots \rightarrow s_n$ and s_0 is initial in S .

Innocence is there to forbid interference between Player moves appearing in distinct gccs, where by *interference* we mean either causal dependence, conflict, or probabilistic dependence. Note that innocence already appears in [CCW15] for strategies without probability. The contribution of this thesis is essentially limited to an extra constraint on the valuation.

4.2.1 Justifiers and visibility


By itself, innocence (and *a fortiori* probabilistic innocence) is not stable under composition of strategies, without first restricting to *visible* strategies. Visibility appears already in Hyland-Ong games [HO00], and involves the notion of *justification pointers*.

Definition 4.6. Let $\sigma : \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$ be a strategy, and $s \in S$ be a non-initial move. The **justifier** of s , denoted $\text{just}(s)$, is defined as follows.

- If σs is an initial A -move, then $\text{just}(s)$ is the unique $s' \in S$ such that $s' \rightarrow s$. ($\sigma s'$ is necessarily an initial B -move. Uniqueness is because σ is well-threaded.)
- Otherwise, $\text{just}(s)$ is the unique $s' \in S$ such that $\sigma s' \rightarrow \sigma s$. (Uniqueness is because arenas are forest-shaped and σ is locally injective.)

This is depicted using dashed lines (**justification pointers**) from every non-initial Player move to its justifier. We do not need to specify the justifier of an Opponent move s : by the receptivity condition on strategies, this must necessarily be the unique t such that $t \rightarrow s$.

Example 4.7. The identity strategy on the game $\mathcal{A} = \{\ominus \rightarrow \oplus\}$ in \mathbf{PG}_1 is the strategy $\varepsilon_{\mathcal{A}}$ on $!\mathcal{A}^\perp \parallel \mathcal{A}$ defined by:



in which the valuation is 1 everywhere.

We can finally define:

Definition 4.8. A strategy $\sigma : \mathcal{S} \rightarrow \mathcal{A}$ is **visible** if for every gcc ρ in S , $\sigma\rho \in \mathcal{C}(A)$.

An equivalent requirement, closer in spirit to the visibility condition in [HO00], is that any gcc containing a non-initial move must also contains its justifier.

Visible strategies are closed under composition [CCW15], pairing, tensor and currying. They also include copycat and all of the structural morphisms of Chapter 2.

4.2.2 Sequential innocence

Visibility should be understood informally as a way of enforcing that each gcc is a well-defined thread. In this point of view, innocence prevents interference between certain concurrent threads of a strategy, namely those spawned by Opponent.

The innocence condition for concurrent strategies appears in [CCW15, Cas17], and this can likely be extended to probabilistic strategies by enforcing an independence constraint on the valuation. But it will not be necessary in this thesis to carry out this extension in full generality. We are concerned here with programming languages without any concurrency primitives, so we only give a probabilistic version of *sequential* innocent strategies. This makes the presentation significantly simpler.

Definition 4.9. A strategy $\sigma : \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$ is sequential innocent if it is visible and satisfies the following properties:

- a subset $X \subseteq S$ is a configuration if and only if $\sigma X \in \text{Con}_{\mathcal{A}^\perp \parallel \mathcal{B}}$ and it is an *Opponent-branching forest* (that is, causality is forest-shaped and if $a \rightarrow b$ and $a \rightarrow c$ in X then $\text{pol}(a) = +$).
- the valuation v_S is Markov.

The first condition is equivalent to the sequential innocence condition for non-deterministic strategies given in [CCW15]. If $\sigma : \mathcal{S} \rightarrow \mathcal{A}$ is sequential innocent, then causality in S is itself forest-shaped, and if $a \rightarrow b$ and $a \rightarrow c$ with $\text{pol}(a) = -$, then necessarily $b \sim c$. In this sense the behaviour of Player is sequential, and any concurrency in S is induced by Opponent.

Additionally, concurrent gccs in S (necessarily Opponent-branching) cannot interfere with each other without violating the requirements. In a Markov strategy causal independence implies probabilistic independence; so there is no “probabilistic interference” between concurrent gccs.

Sequential innocent strategies are stable under composition:

Lemma 4.10. *If $\sigma : \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$ and $\tau : \mathcal{T} \rightarrow \mathcal{B}^\perp \parallel \mathcal{C}$ are sequential innocent strategies, then their composition $\tau \odot \sigma$ is sequential innocent.*

Proof. The first condition is preserved by composition: one can show that every configuration of $T \otimes S$ is an Opponent-branching tree, and this is preserved by hiding. See [CCW15] for details. We show that $v_{T \otimes S}$ is Markov.

Let $x, x' \in \mathcal{C}(T \odot S)$ and $x \dashv\!\! \dashv^s x'$. To avoid ambiguity in the rest of the proof, we denote by $[-]_\odot$ and $[-]_\otimes$ the down-closure operations in $T \odot S$, and $[-]_\otimes$ for the down-closure in $T \otimes S$.

We must show that $v_{T \circledast S}(x') \cdot v_{T \circledast S}([s]_{\circledast}) = v_{T \circledast S}(x) \cdot v_{T \circledast S}([s]_{\circledast})$, *i.e.* $v_{T \circledast S}([x']_{\circledast}) \cdot v_{T \circledast S}([s]_{\circledast}) = v_{T \circledast S}([x]_{\circledast}) \cdot v_{T \circledast S}([s]_{\circledast})$, observing that $[[s]_{\circledast}]_{\circledast} = [s]_{\circledast}$.

We have $[[s]_{\circledast}]_{\circledast} \subseteq [[s]_{\circledast}]_{\circledast}$ in $\mathcal{C}(T \circledast S)$, and because the causality relation in $T \circledast S$ is forest-shaped, there is a unique covering chain

$$[[s]_{\circledast}]_{\circledast} \text{---}^{e_1} u_1 \text{---}^{e_2} \dots \text{---}^{e_{n-1}} u_{n-1} \text{---}^s [s]_{\circledast},$$

with $u_i = [e_i]_{\circledast}$ for every i . There is a covering chain

$$[x]_{\circledast} \text{---}^{e_1} w_1 \text{---}^{e_2} \dots \text{---}^{e_{n-1}} w_{n-1} \text{---}^s [x']_{\circledast}$$

and since by Lemma 3.11 the interaction satisfies the Markov property, we have

$$\begin{aligned} & v_{T \circledast S}([x']_{\circledast} \mid [x]_{\circledast}) \\ &= v_{T \circledast S}([x']_{\circledast} \mid w_{n-1}) \cdot v_{T \circledast S}(w_{n-1} \mid w_{n-2}) \dots v_{T \circledast S}(w_1 \mid [x]_{\circledast}) \\ &= v_{T \circledast S}([s]_{\circledast} \mid u_{n-1}) \cdot v_{T \circledast S}(u_{n-1} \mid u_{n-2}) \dots v_{T \circledast S}(u_1 \mid [[s]_{\circledast}]_{\circledast}) \\ &= v_{T \circledast S}([s]_{\circledast} \mid [[s]_{\circledast}]_{\circledast}) \end{aligned}$$

which concludes the proof. \square

For any arena \mathcal{A} , the copycat strategy is sequential innocent, so there is a sub-bicategory \mathbf{PG}^{si} of \mathbf{PG} whose morphisms are sequential innocent strategies. Furthermore, \mathbf{PG}^{si} retains all the structure of \mathbf{PG} : it is a symmetric monoidal closed bicategory, with finite products and supporting a linear exponential pseudo-comonad. In particular, the Kleisli bicategory $\mathbf{PG}_!^{\text{si}}$ is cartesian closed.

Verifying the above properties is straightforward, with the key observation that the tensor product creates no additional probabilistic dependencies. Details are omitted.

4.3 Adequacy

Observe that the reflexive object \mathcal{U} is still reflexive in $\mathbf{PG}_!^{\text{si}}$, because the morphisms app and λ are sequential innocent.

Further, the definition of $[[M]]$ for $M \in \boldsymbol{\lambda}^+$ is not affected by the restriction: all strategies involved in the definition (4.1.3) are sequential innocent, and the sum $\sigma +_p \tau$ of sequential innocent strategies $\sigma, \tau \in \mathbf{PG}_!^{\text{si}}$ is also sequential innocent.

So in this section we prove an adequacy theorem bridging the operational and denotational semantics of $\boldsymbol{\lambda}^+$. We begin by defining notions of convergence for terms and strategies.

4.3.1 Convergence

The normal forms for the reduction \xrightarrow{p} defined in 4.1 are terms of the form

$$\lambda x_0 \dots x_{n-1}. y M_0 \dots M_{k-1},$$

for non-zero $n, k \in \mathbb{N}$ and $M_i \in \Lambda^+$ for all i . Such terms are called **head-normal forms** (hnfs). In the pure λ -calculus, each term has at most one head-normal form. This is of course not the case in $\boldsymbol{\lambda}^+$.

The **probability of convergence** of a term M , denoted $\Pr_{\Downarrow}(M)$, is the probability of M reducing to some hnf: $\Pr_{\Downarrow}(M) = \sum_{H \text{ hnf}} \Pr(M \rightarrow H)$. More generally for any set \mathcal{N} of terms we write $\Pr(M \rightarrow \mathcal{N}) = \sum_{N \in \mathcal{N}} \Pr(M \rightarrow N)$.

Finally we say that two terms M and N are **observationally equivalent**, written $M \cong N$, if for all contexts $C[\]$, $\Pr_{\Downarrow}(C[M]) = \Pr_{\Downarrow}(C[N])$.

Now, given a strategy $\sigma : \mathcal{S} \rightarrow (!\mathcal{U}^{\Gamma})^{\perp} \parallel \mathcal{U}$, we obtain its **probability of convergence** by summing the coefficients assigned to the first level of Player moves. Opponent plays the initial move, so it is equivalent to consider the configurations containing exactly 2 events:

$$\Pr_{\Downarrow}(\sigma) = \sum_{\substack{x \in \mathcal{C}(\mathcal{S}) \\ |x|=2}} v_{\mathcal{S}}(x).$$

We show an adequacy theorem for the interpretation of λ^+ -terms in \mathcal{U} :

Theorem 4.11 (Adequacy). *For each $M \in \lambda^+_{\Gamma}$,*

$$\Pr_{\Downarrow}(\llbracket M \rrbracket^{\Gamma}) \leq \Pr_{\Downarrow}(M).$$

The proof involves a relation \triangleleft between terms and strategies which we obtain using a fixed point construction. This is an untyped alternative to the “logical relations” technique traditionally used to show adequacy. This method was put forward by Pitts [Pit93] and recently used to obtain adequacy results for λ^+ – see [EPT11, LP19], from which we draw much inspiration.

4.3.2 The relation \triangleleft

The relation \triangleleft is constructed so as to contain pairs (σ, M) where $\sigma \in \mathbf{PG}_{!}^{\text{si}}[\mathcal{U}^{\Gamma}, \mathcal{U}]$ for some Γ , and $M \in \lambda^+$ has free variables in Γ . So we consider the set of relations R with

$$R \subseteq \bigcup_{\Gamma} \mathbf{PG}_{!}^{\text{si}}[\mathcal{U}^{\Gamma}, \mathcal{U}] \times \lambda^+_{\Gamma},$$

where Γ ranges over finite sets of variables (and recall $\lambda^+_{\Gamma} = \{M \in \lambda^+ \mid \text{fv}(M) \subseteq \Gamma\}$).

For each such relation R we will define a new relation $\phi(R)$, and the relation \triangleleft will arise as a fixed point for ϕ . We first introduce some notation. If Γ and Δ are sets of variables with $\Gamma \subseteq \Delta$, then any strategy $\sigma : \mathcal{S} \rightarrow (\mathcal{U}^{\Gamma})^{\perp} \parallel \mathcal{U}$ gives rise to a strategy $\sigma^{\uparrow\Delta}$ defined by the map

$$\sigma^{\uparrow\Delta} : \mathcal{S} \rightarrow (\mathcal{U}^{\Gamma})^{\perp} \parallel \mathcal{U} \rightarrow (\mathcal{U}^{\Delta})^{\perp} \parallel \mathcal{U}$$

where the second arrow is the canonical injection $(\mathcal{U}^{\Gamma})^{\perp} \parallel \mathcal{U} \rightarrow (\mathcal{U}^{\Delta})^{\perp} \parallel \mathcal{U}$.

Then, given $\sigma \in \mathbf{PG}_{!}^{\text{si}}[\mathcal{U}^{\Gamma}, \mathcal{U}]$ and $\tau \in \mathbf{PG}_{!}^{\text{si}}[\mathcal{U}^{\Delta}, \mathcal{U}]$, define

$$\sigma @ \tau = \text{Ev} \odot_{!} \langle \text{app} \odot_{!} \sigma^{\uparrow\Delta}, \tau \rangle$$

and we write $\sigma @ \tau_1 \dots \tau_n$ for $(\sigma @ \tau_1) \dots @ \tau_n$.

Now fix an arbitrary relation R . For each Γ , define the relation $\phi(R)$ to contain the pair $(\sigma, M) \in \mathbf{PG}_!^{\text{si}}[\mathcal{U}^\Gamma, \mathcal{U}] \times \boldsymbol{\lambda}^+_\Gamma$ if the following condition holds: $\forall \Delta \supseteq \Gamma, \forall n \in \mathbb{N}$, if $\tau_1, \dots, \tau_n \in \mathbf{PG}_!^{\text{si}}[\mathcal{U}^\Delta, \mathcal{U}]$ and $N_1, \dots, N_n \in \boldsymbol{\lambda}^+_\Delta$ with all $(\tau_i, N_i) \in R$, then

$$\text{Pr}_\Downarrow(\sigma @ \tau_1 \dots \tau_n) \leq \text{Pr}_\Downarrow(MN_1 \dots N_n).$$

The action of ϕ is anti-monotone with respect to inclusion of relations, so in order to find a fixed point for it we consider the map ψ defined for each pair of relations (R^+, R^-) as

$$\psi(R^+, R^-) = (\phi(R^-), \phi(R^+)).$$

The map ψ is monotone for the order on pairs defined as $(R_1^+, R_1^-) \sqsubseteq (R_2^+, R_2^-)$ iff $R_1^+ \subseteq R_2^+$ and $R_1^- \subseteq R_2^-$, and moreover the order \sqsubseteq makes the set of pairs of relations a complete lattice.

Consider the set of pre-fixed points for ψ :

$$\{(R_1, R_2) \mid \psi(R_1, R_2) \sqsubseteq (R_1, R_2)\}$$

We write $(\triangleleft^+, \triangleleft^-)$ for its glb, which by Tarski's Theorem is a least fixed point for ψ : $\psi(\triangleleft^+, \triangleleft^-) = (\triangleleft^+, \triangleleft^-)$, or in other words $\phi(\triangleleft^+) = \triangleleft^-$ and $\phi(\triangleleft^-) = \triangleleft^+$. From this we deduce easily that $(\triangleleft^-, \triangleleft^+)$ is also a fixed point for ψ , so that $(\triangleleft^+, \triangleleft^-) \sqsubseteq (\triangleleft^-, \triangleleft^+)$, *i.e.* $\triangleleft^+ \subseteq \triangleleft^-$. As we will see now, the reverse inclusion holds too, and we will take $\triangleleft = \triangleleft^- = \triangleleft^+$. We first note the following general fact:

Lemma 4.12. *For any relation R , if $M \in \boldsymbol{\lambda}^+_\Gamma$, then the set $\{\sigma \in \mathbf{PG}_!^{\text{si}}[\mathcal{U}^\Gamma, \mathcal{U}] \mid (\sigma, M) \in \phi(R)\}$ is closed under ω -colimits of strong embeddings.*

Proof. Observe first that if strategies $\eta_i \in \mathbf{PG}_!^{\text{si}}[\mathcal{U}^\Delta, \mathcal{U}]$ form an ω -chain, then

$$\text{Pr}_\Downarrow(\bigvee_{i \in \omega} \eta_i) = \sup_{i \in \omega} \text{Pr}_\Downarrow(\eta_i).$$

Now, let $\sigma_0 \xrightarrow{f_0} \sigma_1 \xrightarrow{f_1} \dots$ be a chain of strong embeddings in $\mathbf{PG}_!^{\text{si}}[\mathcal{U}^\Gamma, \mathcal{U}]$, and suppose $(\sigma_i, M) \in \phi(R)$ for all σ . Let $\Delta \supseteq \Gamma$, $\tau_1, \dots, \tau_n \in \mathbf{PG}_!^{\text{si}}[\mathcal{U}^\Delta, \mathcal{U}]$, and N_1, \dots, N_n with each $(\tau_k, N_k) \in R$.

Since colimits commute with all constructions, we have $(\bigvee_{i \in \omega} \sigma_i) @ \tau_1 \dots \tau_n = \bigvee_{i \in \omega} (\sigma_i @ \tau_1 \dots \tau_n)$, so that by the remark above,

$$\text{Pr}_\Downarrow((\bigvee_{i \in \omega} \sigma_i) @ \tau_1 \dots \tau_n) = \sup_{i \in \omega} \text{Pr}_\Downarrow(\sigma_i @ \tau_1 \dots \tau_n).$$

By assumption, $\text{Pr}_\Downarrow(\sigma_i @ \tau_1 \dots \tau_n) \leq \text{Pr}_\Downarrow(MN_1 \dots N_n)$ for each i , so

$$\sup_{i \in \omega} \text{Pr}_\Downarrow(\sigma_i @ \tau_1 \dots \tau_n) \leq \text{Pr}_\Downarrow(MN_1 \dots N_n),$$

hence $(\bigvee_{i \in \omega} \sigma_i, M) \in \phi(R)$ and we are done. \square

Then, to show $\triangleleft^- \subseteq \triangleleft^+$ we consider for each strategy $\sigma \in \mathbf{PG}_!^{\text{si}}[\mathcal{U}^\Gamma, \mathcal{U}]$ a sequence $(\sigma_i)_{i \in \omega}$ of *finite-depth* approximants. Let us say first what we mean by the *depth* of a strategy:

Definition 4.13. The **depth** of a sequential innocent strategy $\sigma : \mathcal{S} \rightarrow \mathcal{U}^\perp \parallel \mathcal{U}$, $\text{depth}(\sigma)$, is the maximum number of Player moves in a gcc of S , and ∞ when all gccs have unbounded length.

Given a strategy $\sigma : \mathcal{S} \rightarrow (\mathcal{U}^\Gamma)^\perp \parallel \mathcal{U}$, and $d \in \mathbb{N}$, let σ_d be the largest sub-strategy of σ with $\text{depth} \leq d$. It is clear that the σ_d approximate σ in the sense that the latter is the colimit in $\mathbf{PG}_!^{\text{si}}[\mathcal{U}^\Gamma, \mathcal{U}]$ of the chain $\sigma_0 \Rightarrow \sigma_1 \Rightarrow \dots$ (where all arrows are inclusion maps).

Lemma 4.14. *The relations \triangleleft^- and \triangleleft^+ satisfy $\triangleleft^+ \subseteq \triangleleft^-$, so that $\triangleleft^+ = \triangleleft^-$.*

Proof. We show by induction that for all $d \in \mathbb{N}$, if $\sigma \triangleleft^- M$ then $\sigma_d \triangleleft^+ M$. By Lemma 4.12, this shows that $\sigma \triangleleft^+ M$.

For $d = 0$, $\sigma_d \cong \perp$, and so for any τ_1, \dots, τ_n , $\sigma_d @ \tau_1 \dots \tau_n \cong \perp$. In particular, $\text{Pr}_\Downarrow(\sigma_d @ \tau_1 \dots \tau_n) = 0$; from this we easily deduce $(\sigma_d, M) \in \phi(\triangleleft^-) = \triangleleft^+$.

For the inductive step, observe first the following general fact: if $\sigma \Rightarrow \sigma'$ is an embedding, then $\text{Pr}_\Downarrow(\sigma) \leq \text{Pr}_\Downarrow(\sigma')$. Furthermore if $(\sigma', M) \in \phi(R)$ for some R , then $(\sigma, M) \in \phi(R)$.

Assume $\sigma_d \triangleleft^+ M$. Because $\sigma \triangleleft^- M$, we have $\sigma_{d+1} \triangleleft^- M$ by the remark of the previous paragraph. To show that $\sigma_{d+1} \triangleleft^+ M$, we use that $\triangleleft^+ = \phi(\triangleleft^-)$. Let $\Delta \supseteq \Gamma$, and let $\tau_1, \dots, \tau_n \in \mathbf{PG}_!^{\text{si}}[\mathcal{U}^\Delta, \mathcal{U}]$ and N_1, \dots, N_n with $\tau_i \triangleleft^- N_i$. We show that

$$\text{Pr}_\Downarrow(\sigma_{d+1} @ \tau_1 \dots \tau_n) \leq \text{Pr}_\Downarrow(MN_1 \dots N_n).$$

Applying the induction hypothesis to the τ_i , we get that $(\tau_i)_d \triangleleft^+ N_i$ for each i . Since $\sigma_{d+1} \triangleleft^- M$ and $\triangleleft^- = \phi(\triangleleft^+)$, we deduce that

$$\text{Pr}_\Downarrow(\sigma_{d+1} @ (\tau_1)_d \dots (\tau_n)_d) \leq \text{Pr}_\Downarrow(MN_1 \dots N_n).$$

It suffices now to show that $\text{Pr}_\Downarrow(\sigma_{d+1} @ (\tau_1)_d \dots (\tau_n)_d) = \text{Pr}_\Downarrow(\sigma_{d+1} @ \tau_1 \dots \tau_n)$, and this is a straightforward induction on n , observing that in a composition of the form $\text{Ev} \odot \langle \text{app} \odot \sigma_{d+1}, \tau \rangle$, a τ -event of depth $> d$ cannot occur. \square

So we define $\triangleleft = \triangleleft^+ = \triangleleft^-$. We move to the proof of the adequacy theorem.

4.3.3 Closure properties and proof of adequacy

At this point, we have defined a relation $\triangleleft \subseteq \bigcup_\Gamma \mathbf{PG}_!^{\text{si}}[\mathcal{U}^\Gamma, \mathcal{U}] \times \lambda^+_\Gamma$ such that $\phi(\triangleleft) = \triangleleft$. We now aim to show that $\llbracket M \rrbracket^\Gamma \triangleleft M$ for every M , which implies the adequacy theorem. This involves checking that \triangleleft satisfies a number of properties.

Lemma 4.15. *The relation \triangleleft satisfies the following properties:*

- (1) *If $\sigma \triangleleft M$ and $\sigma \cong \sigma'$, then $\sigma' \triangleleft M$.*
- (2) *$\pi_x \triangleleft x$, where $\pi_x \in \mathbf{PG}_!^{\text{si}}[\mathcal{U}^{x, \Gamma}, \mathcal{U}]$ is the x^{th} projection.*
- (3) *$\sigma \triangleleft (\lambda x.M)N$ iff $\sigma \triangleleft M[N/x]$.*
- (4) *If $\sigma \triangleleft M$ and $\tau \triangleleft N$ then $\sigma +_p \tau \triangleleft M +_p N$.*

Proof. (1) Straightforward, because $\sigma \cong \sigma'$ implies $\text{Pr}_{\Downarrow}(\sigma) = \text{Pr}_{\Downarrow}(\sigma')$.

(2) Notice that for any N_1, \dots, N_n , $\text{Pr}_{\Downarrow}(xN_1 \dots N_n) = 1$, so necessarily $\text{Pr}_{\Downarrow}(\pi_x @ \tau_1 \dots \tau_n) \leq \text{Pr}_{\Downarrow}(xN_1 \dots N_n)$. So $(\pi_x, x) \in \phi(\triangleleft) = \triangleleft$.

(3) This follows directly from the fact that $\text{Pr}_{\Downarrow}((\lambda x.M)NN_1 \dots N_n) = \text{Pr}_{\Downarrow}(M[N/x]N_1 \dots N_n)$ for any N_1, \dots, N_n .

(4) For any τ_1, \dots, τ_n , $(\sigma +_p \tau) @ \tau_1 \dots \tau_n \cong (\sigma @ \tau_1 \dots \tau_n) +_p (\tau @ \tau_1 \dots \tau_n)$ so in particular $\text{Pr}_{\Downarrow}((\sigma +_p \tau) @ \tau_1 \dots \tau_n) = p \cdot \text{Pr}_{\Downarrow}(\sigma @ \tau_1 \dots \tau_n) + (1-p) \cdot \text{Pr}_{\Downarrow}(\tau @ \tau_1 \dots \tau_n)$. It is clear that $\text{Pr}_{\Downarrow}((M +_p N)N_1 \dots N_n) = p \cdot \text{Pr}_{\Downarrow}(MN_1 \dots N_n) + (1-p) \cdot \text{Pr}_{\Downarrow}(NN_1 \dots N_n)$, so the result follows. \square

We will also need the following characterisation:

Lemma 4.16. *For every $(\sigma, M) \in \mathbf{PG}_!^{\text{si}}[\mathcal{U}^\Gamma, \mathcal{U}] \times \boldsymbol{\lambda}^+_\Gamma$, the following are equivalent:*

(1) $\sigma \triangleleft M$;

(2) for every $\Delta \supseteq \Gamma$ and $(\tau, N) \in \mathbf{PG}_!^{\text{si}}[\mathcal{U}^\Delta, \mathcal{U}] \times \boldsymbol{\lambda}^+_\Delta$, if $\tau \triangleleft N$ then $\sigma @ \tau \triangleleft MN$.

Proof. The (1) \Rightarrow (2) direction is immediate, since $\phi(\triangleleft) = \triangleleft$. Assume (2) holds, and we show $(\sigma, M) \in \phi(\triangleleft)$. We must show that for every $\Delta \subset \Gamma$, $n \in \mathbb{N}$, $\tau_1, \dots, \tau_n \in \mathbf{PG}_!^{\text{si}}[\mathcal{U}^\Delta, \mathcal{U}]$ and $N_1, \dots, N_n \in \boldsymbol{\lambda}^+_\Delta$, if $\tau_i \triangleleft N_i$ for all i then $\text{Pr}_{\Downarrow}(\sigma @ \tau_1 \dots \tau_n) \leq \text{Pr}_{\Downarrow}(MN_1 \dots N_n)$. Because by assumption, $\sigma @ \tau_1 \triangleleft MN_1$, the property holds directly for every $n \geq 1$. It remains to treat the case $n = 0$, i.e. show that $\text{Pr}_{\Downarrow}(\sigma) \leq \text{Pr}_{\Downarrow}(M)$.

But by Lemma 4.15(2), $\pi_x \triangleleft x$, and so $\sigma @ \pi_x \triangleleft Mx$ and in particular $\text{Pr}_{\Downarrow}(\sigma @ \pi_x) \leq \text{Pr}_{\Downarrow}(Mx)$. But standard reasoning shows $\sigma \cong \lambda \odot_! \text{Cur}(\sigma @ \pi_x)$, hence $\text{Pr}_{\Downarrow}(\sigma @ \pi_x) = \text{Pr}_{\Downarrow}(\sigma)$. An easy argument shows $\text{Pr}_{\Downarrow}(Mx) = \text{Pr}_{\Downarrow}(M)$, so $\text{Pr}(\sigma) \leq \text{Pr}(M)$, and we are done. \square

We are now able to show the following:

Lemma 4.17. *Let $\Gamma = x_1, \dots, x_n$, and let $M \in \boldsymbol{\lambda}^+_\Gamma$. Let $\Delta \supseteq \Gamma$ and suppose $\tau_1, \dots, \tau_n \in \mathbf{PG}_!^{\text{si}}[\mathcal{U}^\Delta, \mathcal{U}]$ and $N_1, \dots, N_n \in \boldsymbol{\lambda}^+_\Delta$ with $\tau_i \triangleleft N_i$ for each i . Then,*

$$\llbracket M \rrbracket^\Gamma \odot_! \langle \tau_1, \dots, \tau_n \rangle \triangleleft M[N_1/x_1, \dots, N_n/x_n].$$

Proof. The proof is by induction on M .

Case $M = x$. We must have $x = x_j$ for some j , then $M[N_i/x_i] = N_j$ and $\llbracket M \rrbracket^\Gamma \odot_! \langle \sigma_1, \dots, \sigma_n \rangle \cong \sigma_j$, so the result holds using Lemma 4.15(1).

Case $M = \lambda x.N$. We use Lemma 4.16 and show that for each $\Theta \supseteq \Delta$, $\rho \in \mathbf{PG}_!^{\text{si}}[\mathcal{U}^\Theta, \mathcal{U}]$ and $P \in \boldsymbol{\lambda}^+_\Theta$ with $\rho \triangleleft P$,

$$(\llbracket M \rrbracket^\Gamma \odot_! \langle \tau_1, \dots, \tau_n \rangle) @ \rho \triangleleft (M[N_i/x_i])P.$$

But $(\llbracket M \rrbracket^\Gamma \odot_! \langle \tau_1, \dots, \tau_n \rangle) @ \rho \cong \llbracket N \rrbracket^{\Gamma, x} \odot_! \langle \tau_1, \dots, \tau_n, \rho \rangle$, and by the induction hypothesis, $\llbracket N \rrbracket^{\Gamma, x} \odot_! \langle \tau_1, \dots, \tau_n, \rho \rangle \triangleleft N[N_i/x_i, P/x]$. But $(M[N_i/x_i])P \rightarrow N[N_i/x_i, P/x]$ so we conclude by Lemma 4.15(3).

Case $M = NN'$. We show $(\llbracket NN' \rrbracket^\Gamma \odot_! \langle \tau_1, \dots, \tau_n \rangle) \triangleleft (NN')[N_i/x_i]$. This is equivalent to $(\llbracket N \rrbracket^\Gamma \odot_! \langle \tau_i \rangle) @ (\llbracket N' \rrbracket^\Gamma \odot_! \langle \tau_i \rangle) \triangleleft N[N_i/x_i]N'[N_i/x_i]$, which follows from the IH for N and N' .

Case $M = N +_p P$. Direct consequence of Lemma 4.15(4). \square

The adequacy property is a direct consequence of the lemma:

Theorem 4.18. *For any $M \in \lambda^+_{\Gamma}$,*

$$\Pr_{\Downarrow}(\llbracket M \rrbracket^{\Gamma}) \leq \Pr_{\Downarrow}(M).$$

Proof. By Lemma 4.17, $\llbracket M \rrbracket^{\Gamma} \triangleleft M$, so $(\llbracket M \rrbracket^{\Gamma}, M) \in \phi(\triangleleft)$, and thus $\Pr_{\Downarrow}(\llbracket M \rrbracket^{\Gamma}) \leq \Pr_{\Downarrow}(M)$. \square

4.4 Strategies and Nakajima trees

In this section we refine the computational meaning of the semantics through a connection between innocent strategies and the probabilistic Nakajima trees of Leventis [Lev16]. We start by defining the latter.

4.4.1 Probabilistic Nakajima trees

Nakajima trees. The Nakajima tree [Bar84] of a pure λ -term M is in general an infinite tree, which can be defined as the limit of a sequence of finite-depth approximants. In fact those approximants will suffice for our purposes: given a λ -term M and $d \in \mathbb{N}$, the tree $\text{NT}^d(M)$ is \perp if $d = 0$ or if M has no head-normal form, and

$$\begin{array}{c} \lambda z_0 \dots z_{n-1} x_0 x_1 \dots \bullet y \\ \swarrow \quad \downarrow \quad \searrow \quad \searrow \\ \text{NT}^{d-1}(P_0) \quad \dots \quad \text{NT}^{d-1}(P_{k-1}) \quad \text{NT}^{d-1}(x_0) \quad \text{NT}^{d-1}(x_1) \quad \dots \end{array}$$

if $d > 0$ and M has hnf $\lambda z_0 \dots z_{n-1}.y P_0 \dots P_{k-1}$.

In order to deal with issues of α -renaming, we adopt a convention also used in [Lev16], whereby the infinite sequence of abstracted variables at the root of a tree of depth $d > 0$ is labelled x_0^d, x_1^d, \dots so that any tree is determined by the pair $(y, (T_n)_{n \in \mathbb{N}})$ of its head variable and sequence of subtrees.

Leventis' probabilistic trees Nakajima trees for the λ -calculus have striking properties: they characterise observational equivalence of terms, and as a model they yield the *maximal consistent sensible* λ -theory (see [Bar84] for details). In his PhD thesis, Leventis [Lev16] proposes a notion of *probabilistic* Nakajima tree which plays the same role for λ^+ . Intuitively, because a term of the form $\lambda x_0 \dots x_{n-1}.z P_0 \dots P_{k-1} +_p \lambda y_0 \dots y_{m-1}.w Q_0 \dots Q_{l-1}$ has two hnf's, it may be represented by a probability distribution over trees of the form of that above. Accordingly, two different kinds of trees are considered: **value trees**, representing head-normal forms (without probability distribution at top-level), and **probabilistic Nakajima trees**, representing general terms:

Definition 4.19. For each $d \in \mathbb{N}$, the sets \mathcal{PT}^d of **probabilistic Nakajima trees of depth d** and \mathcal{VT}^d of **value trees of depth d** are defined by $\mathcal{VT}^0 = \emptyset$,

$$\begin{aligned} \mathcal{VT}^{d+1} &= \{(y, (T_n)_{n \in \mathbb{N}}) \mid y \in \text{Var} \text{ and } \forall n \in \mathbb{N}, T_n \in \mathcal{PT}^d\} \text{ and} \\ \mathcal{PT}^d &= \{T : \mathcal{VT}^d \rightarrow [0, 1] \mid \sum_{t \in \mathcal{VT}^d} T(t) \leq 1\}. \end{aligned}$$

We can then assign trees to individual terms:

Definition 4.20. Given $M \in \lambda^+$ and $d \in \mathbb{N}$, its **probabilistic Böhm tree of depth d** is the tree $\text{PT}^d(M) \in \mathcal{PT}^d$ defined as follows:

$$\begin{aligned} \text{PT}^d(M) : \mathcal{VT}^d &\longrightarrow [0, 1] \\ t &\longmapsto \Pr(M \rightarrow \{H \text{ hnf} \mid \text{VT}^d(H) = t\}) \end{aligned}$$

where for any hnf $H = \lambda z_0 \dots z_{n-1}.y P_0 \dots P_{k-1}$, the **value tree of depth d of H** is defined as

$$\text{VT}^d(H) = (y, (\text{PT}^{d-1}(P_0), \dots, \text{PT}^{d-1}(P_{k-1}), \text{PT}^{d-1}(x_n^d), \dots)).$$

Consider for example the term $M_1 = \lambda xy.x (y + \frac{1}{3}(\lambda z.z))$, a head-normal form. The first steps in the construction of its value tree of depth d , for some fixed $d \geq 2$ are given as follows (where we use the symbol δ_t to denote the Dirac distribution at t : $\delta_t(t') = 1$ if $t = t'$, and 0 otherwise):


$$\begin{array}{c} \lambda x_0^d x_1^d x_2^d \dots \bullet x_0^d \\ \swarrow \quad \downarrow \quad \searrow \\ \frac{1}{3} \delta_{\text{VT}^{d-1}(x_1^d)} + \frac{2}{3} \delta_{\text{VT}^{d-1}(\lambda z.z)} \delta_{\text{VT}^{d-1}(x_2^d)} \delta_{\text{VT}^{d-1}(x_3^d)} \dots \end{array}$$

where $\text{VT}^{d-1}(x_l^d)$ (for $l \in \mathbb{N}$) and $\text{VT}^{d-1}(\lambda z.z)$ are

$$\begin{array}{cc} \begin{array}{c} \lambda x_0^{d-1} x_1^{d-1} x_2^{d-1} \dots \bullet x_l^d \\ \swarrow \quad \downarrow \\ \delta_{\text{VT}^{d-2}(x_0^{d-1})} \delta_{\text{VT}^{d-2}(x_1^{d-1})} \dots \end{array} & \begin{array}{c} \lambda x_0^{d-1} x_1^{d-1} x_2^{d-1} \dots \bullet x_0^{d-1} \\ \swarrow \quad \downarrow \\ \delta_{\text{VT}^{d-2}(x_1^{d-1})} \delta_{\text{VT}^{d-2}(x_2^{d-1})} \dots \end{array} \end{array}$$

and so on.

As an aside, compare the above with the interpretation of M_1 as a strategy on \mathcal{U} :



The strategy is an alternative representation of the term; but in this particular example it corresponds precisely to the Nakajima tree representation. Ignoring the depth superscript on the variable names in the latter, we see a correspondence between the negative branching in the strategy and the branches in the value trees, and between the positive branching and the variables indices at the head of probabilistic trees. This correspondence will be formalised below. Note that the strategy is pictured in “reduced form” [CCW15], meaning that the (countably many) symmetric copies of each Opponent move are not pictured.

Probabilistic Nakajima trees precisely characterise observational equivalence in λ^+ ; writing $M =_{PT} N$ if for every $d \in \mathbb{N}$, $PT^d(M) = PT^d(N)$, we have:

Theorem 4.21 (Leventis [Lev16]). *For any $M, N \in \Lambda^+$, $M \cong N$ if and only if $M =_{PT} N$.*

In this sense probabilistic Nakajima trees provide a *fully abstract* interpretation of the probabilistic λ -calculus.

4.4.2 Nakajima-like strategies

In [KNO02], the authors prove an *exact correspondence theorem* for the pure λ -calculus: Nakajima trees precisely correspond to deterministic innocent strategies on the universal arena.

For λ^+ however, the correspondence is not so exact: although terms M and $M +_p M$ have the same Nakajima tree, they have different interpretations in $\mathbf{PG}_1^{\text{si}}$, where each probabilistic choice is recorded as an explicit branching point.¹ In what follows, we identify a class of *Nakajima-like* probabilistic strategies for which the exact correspondence does hold, and we show that any strategy can be reduced to a Nakajima-like one, essentially by quotienting out the “redundant” branching. This yields a notion of equivalence between strategies, defined as reduction to the same Nakajima-like strategy.

¹In particular, $\mathbf{PG}_1^{\text{si}}$ does not yield a *probabilistic λ -theory* in the sense of Leventis [Lev16].

First, given a sequential innocent strategy $\sigma : \mathcal{S} \rightarrow \mathcal{U}$, define a relation \approx on the events of S as the smallest equivalence relation such that if $s_1 \approx s'_1$, $s_1 \rightarrow s_2$, $s'_1 \rightarrow s'_2$ and there is an order-isomorphism $\varphi : \{s \in S \mid s_2 \leq s\} \cong \{s' \in S \mid s'_2 \leq s'\}$ such that for all $s \geq s_2$,

- $\sigma s \sim^+ (\sigma \circ \varphi) s$, and
- $v_S([s] \mid [s_2]) = v_S([\varphi(s)] \mid [s'_2])$,

then $s_2 \approx s'_2$. Informally, \approx identifies events coming from the same syntactic construct in two copies of a term in an idempotent probabilistic sum, as in $M +_p M$ (where Opponent has played the same copy indices).

Definition 4.22. We say $\sigma : \mathcal{S} \rightarrow \mathcal{U}$ is **Nakajima-like** if v_S is non-vanishing and for every $s, s' \in S$, if $s \approx s'$ then $s = s'$.

In other words, a Nakajima-like strategy is one with no redundant branches. Many λ^+ -strategies do not satisfy this property, but all can be reduced to one that does.

Given an innocent sequential strategy $\sigma : \mathcal{S} \rightarrow \mathcal{U}^\Gamma \rightarrow \mathcal{U}$, construct a sub-strategy $\sigma^{>0} : \mathcal{S}^{>0} \rightarrow \mathcal{U}^\Gamma$ with events $S^{>0} = \{s \in S \mid v([s]) > 0\}$. Then, let S_{nak} be the set of \approx -equivalence classes in $S^{>0}$. We now show that this inherits a partial order structure from $\mathcal{S}^{>0}$; we will then turn this into an essp.

Lemma 4.23. (1) *If $s \approx s'$ then there is a bijection $l : [s] \cong [s']$ such that $t \approx l(t)$ for every $t \in [s]$. (n.b. $[s]$ is the down-closure, not the equivalence class of s .)*

(2) *The relation defined on \approx -equivalence classes as $\mathbf{s} \leq \mathbf{t}$ iff there exists $s \in \mathbf{s}$ and $t \in \mathbf{t}$ with $s \leq t$ is a partial order.*

(3) *The order \leq is tree-shaped.*

Proof. (1) By induction on the definition of \approx . If $s = s'$ then this is clearly true, taking l to be the identity. Moreover the property is closed under applications of transitivity and symmetry for \approx . Now suppose $s_1 \approx s'_1$ and there is such a bijection $l_1 : [s_1] \rightarrow [s'_1]$. If $s_1 \rightarrow s_2$, $s'_1 \rightarrow s'_2$ with conditions implying $s_2 \approx s'_2$. Then, because causality in $S^{>0}$ is forest-shaped, $[s_2] = [s_1] \uplus \{s_2\}$, and $[s'_2] = [s'_1] \uplus \{s'_2\}$, we can define $l_2 : [s_2] \rightarrow [s'_2]$ as the extension of l_1 sending s_2 to s'_2 .

(2) Reflexivity is clear. For transitivity, suppose $\mathbf{s} \leq \mathbf{t}$ and $\mathbf{t} \leq \mathbf{u}$, so there is $s \in \mathbf{s}$, $t, t' \in \mathbf{t}$ and $u \in \mathbf{u}$ with $s \leq t$ and $t' \leq u$. Since $t \approx t'$, by (1) there is $s' \in \mathbf{s}$ such that $s' \leq t'$. So $s' \leq u$ and thus $\mathbf{s} \leq \mathbf{u}$. For antisymmetry, suppose $s, s' \in \mathbf{s}$ and $t, t' \in \mathbf{t}$ are such that $s \leq t$ and $t' \leq s'$. By (1), there is $t'' \leq s$ such that $t'' \approx t' \approx t$. But $t \leq t''$, so $t = t'' = s$ and therefore $\mathbf{s} = \mathbf{t}$.

(3) Suppose $\mathbf{s} \leq \mathbf{t}$ and $\mathbf{u} \leq \mathbf{t}$, so there is $s \in \mathbf{s}$, $t, t' \in \mathbf{t}$ and $u \in \mathbf{u}$ with $s \leq t$ and $u \leq t'$. Since $t \approx t'$, by (1) there is $s' \in \mathbf{s}$ such that $s' \leq t'$. Since the order in $S^{>0}$ is tree-shaped, $s \leq u$ or $u \leq s$, so $\mathbf{s} \leq \mathbf{u}$ or $\mathbf{u} \leq \mathbf{s}$. \square

With respect to the order defined in (2) above, there is a unique initial move $\varepsilon \in S_{\text{nak}}$. Define a map $\sigma_{\text{nak}} : S_{\text{nak}} \rightarrow \mathcal{U}$ on $\mathbf{s} \in S_{\text{nak}}$ by induction on the length of the unique chain $\varepsilon \rightarrow \dots \rightarrow \mathbf{s}$. Define $\sigma_{\text{nak}}\varepsilon = \varepsilon$. Now suppose $\sigma_{\text{nak}}\mathbf{s} = u \in (U^\Gamma)^\perp \parallel U$

and $\mathbf{s} \rightarrow \mathbf{t}$. Then $\sigma_{\text{nak}}\mathbf{t}$ is taken to be any successor u' of u such that $u' \sim^+ \sigma t$ for every $t \in \mathbf{t}$ (this always exists and is in fact unique when $\text{pol } \mathbf{t} = -$). We define the rest of the structure of the essp \mathcal{S}_{nak} , writing $f : \mathcal{S}^{>0} \rightarrow \mathcal{S}_{\text{nak}}$ for the quotient map:

- *consistency*: $X \subseteq \mathcal{S}_{\text{nak}}$ iff $[X]$ is an Opponent branching-tree and $\sigma_{\text{nak}}X \in \text{Con}_{(U^\Gamma)^\perp \parallel U}$.
- *symmetry*: $\theta : x \cong_{\mathcal{S}_{\text{nak}}} y$ if there is $\varphi : z \cong_S w$ with $f\varphi = \theta$.
- *probability*: taken as the push-forward of v_S under the quotient map $f : \mathcal{S}^{>0} \rightarrow \mathcal{S}_{\text{nak}}$.

Lemma 4.24. *The map $\sigma_{\text{nak}} : \mathcal{S}_{\text{nak}} \rightarrow \mathcal{U}$ is a sequential innocent strategy.*

Proof. We first check that σ_{nak} is indeed a strategy. It is courteous, because if $\mathbf{s} \rightarrow \mathbf{t}$ and $\sigma_{\text{nak}}\mathbf{s} \vdash \sigma_{\text{nak}}\mathbf{t}$, then there are $s \in \mathbf{s}$ and $t \in \mathbf{t}$ with $s \rightarrow_S t$ but $s \not\vdash_S t$, so $\text{pol}(\mathbf{s}) = \text{pol}(s) = -$ and $\text{pol}(\mathbf{t}) = \text{pol}(t) = +$. To show it is receptive, let $\mathbf{x} \in \mathcal{C}(\mathcal{S}_{\text{nak}})$ and $\sigma_{\text{nak}}\mathbf{x} \subseteq^- y \in \mathcal{C}(U)$. The map $f : \mathcal{C}(S) \rightarrow \mathcal{C}(\mathcal{S}_{\text{nak}})$ is surjective, so there is $x \in \mathcal{C}(S)$ such that $fx = \mathbf{x}$, so there is a unique extension $x \subseteq^- x'$ such that $\sigma x' = y$. Then $\mathbf{x} \subseteq^- fy$, and (by surjectivity of f) this extension is unique among those mapping to y , so σ_{nak} is receptive. The argument for \sim -receptivity is the same. Finally, if $\text{id}_{fx} \subseteq \theta' : \mathbf{x}' \cong_{\mathcal{S}_{\text{nak}}} \mathbf{x}''$, then by definition of $\cong_{\mathcal{S}_{\text{nak}}}$, $\text{id}_x \subseteq^+ \varphi$ with $f\varphi = \theta'$. \mathcal{S} is thin, so $\varphi = \text{id}_{x'}$ for some x' , so $\theta' = \text{id}_{\mathbf{x}'}$ and \mathcal{S}_{nak} is thin.

By construction, $f : \mathcal{S}^{>0} \rightarrow \mathcal{S}_{\text{nak}}$ is a weak map of strategies, so by Lemma 3.25, the pushforward f_*v_S is a valuation on \mathcal{S}_{nak} . We use the characterisation of Lemma 3.9 to show that the valuation $v_{\text{nak}} (= f_*v_S)$ is Markov. (In what follows we use that it is non-vanishing.)

The proof uses the following observation. Let $\mathbf{x} \prec^s \mathbf{x}'$ in $\mathcal{C}(\mathcal{S}_{\text{nak}})$, and for $x \in f^{-1}\{\mathbf{x}\}$, define a set $A_x = \{s \in \mathbf{s} \mid x \cup \{s\} \in \mathcal{C}(S)\}$. Then for any $x, y \in f^{-1}\{\mathbf{x}\}$, there is a bijection $\phi : A_x \cong A_y$ such that for all $s \in A_x$,

$$\frac{v([s])}{v(\uparrow s)} = \frac{v([\phi(s)])}{v(\uparrow \phi(s))}.$$

To see why such a ϕ exists, consider for each $s \in A_x$ its unique predecessor $t \in x$. Then let u be the unique event of y such that $f(t) = f(u)$; since f is a quotient map, $t \approx u$, and so $\phi(s)$ is taken as the appropriate successor of u .

We use this to show v_{nak} is Markov. Let $\mathbf{x} \prec^s \mathbf{x}'$ in $\mathcal{C}(\mathcal{S}_{\text{nak}})$; then

$$\begin{aligned} \frac{v_{\text{nak}}(\mathbf{x}')}{v_{\text{nak}}(\mathbf{x})} &= \frac{\sum_{x \in f^{-1}\{\mathbf{x}\}} \sum_{s \in A_x} v(x \cup \{s\})}{\sum_{x \in f^{-1}\{\mathbf{x}\}} v(x)} \\ &= \frac{\sum_{x \in f^{-1}\{\mathbf{x}\}} v(x) \sum_{s \in A_x} \frac{v([s])}{v(\uparrow s)}}{\sum_{x \in f^{-1}\{\mathbf{x}\}} v(x)} \\ &= \sum_{s \in A_x} \frac{v([s])}{v(\uparrow s)} \quad \text{for any } x \text{ (see above)} \\ &= \frac{v_{\text{nak}}([\mathbf{s}])}{v_{\text{nak}}(\uparrow \mathbf{s})}, \end{aligned}$$

where the last step uses the observation of the previous paragraph (instantiated to $[s] \text{--} \text{c}^s [s]$) and that $[s] = [s']$ for any $s, s' \in A_x$. \square

So in summary, for any sequential innocent $\sigma \in \mathbf{PG}_!^{\text{si}}[\mathcal{U}^\Gamma, \mathcal{U}]$, we construct a Nakajima-like strategy $\sigma_{\text{nak}} \in \mathbf{PG}_!^{\text{si}}[\mathcal{U}^\Gamma, \mathcal{U}]$. Write $\sigma =_{\text{PT}} \tau$ when $\sigma_{\text{nak}} = \tau_{\text{nak}}$.

4.4.3 Nakajima trees correspond to Nakajima-like strategies

We can now make formal the connection between sequential innocent strategies and probabilistic Nakajima trees. To do so we define a bijective map from the set of Nakajima-like strategies of depth d on $(\mathcal{U}^\Gamma)^\perp \parallel \mathcal{U}$, to the set \mathcal{PT}_Γ^d of probabilistic Nakajima trees of depth d with free variables in Γ . Since probabilistic Nakajima trees are defined together with value trees, it will be necessary to also consider a class of *value strategies*:

Definition 4.25. A strategy $\sigma : \mathcal{S} \rightarrow (\mathcal{U}^\Gamma)^\perp \parallel \mathcal{U}$ is a **value strategy** if the initial move $\varepsilon \in S$ has a unique successor s , such that $v_S(\{\varepsilon, s\}) = 1$.

We can now define maps going both ways:

Lemma 4.26. For every $d \in \mathbb{N}$ and every $\Gamma \subseteq_{\text{fin}} \text{Var}$ there are maps

$$\left\{ \sigma \in \mathbf{PG}_!^{\text{si}}[\mathcal{U}^\Gamma, \mathcal{U}] \mid \sigma \text{ is Nakajima-like and } \text{depth}(\sigma) \leq d \right\} \xrightleftharpoons[\langle \Psi_\Gamma^d \rangle^\bullet]{\Psi_\Gamma^d} \mathcal{PT}_\Gamma^d$$

and

$$\left\{ \sigma \in \mathbf{PG}_!^{\text{si}}[\mathcal{U}^\Gamma, \mathcal{U}] \mid \sigma \text{ is a Nakajima-like value strategy and } \text{depth}(\sigma) \leq d \right\} \xrightleftharpoons[\langle \Phi_\Gamma^d \rangle^\bullet]{\Phi_\Gamma^d} \mathcal{VT}_\Gamma^d$$

which are inverses up to isomorphism of strategies.

Proof. By induction on d . The case for $d = 0$ is straightforward: the domain and codomain of Φ_Γ^d are empty, and those of Ψ_Γ^d are singletons.

In the general case, suppose σ is a Nakajima-like value strategy $\mathcal{S} \rightarrow (\mathcal{U}^\Gamma)^\perp \parallel \mathcal{U}$ of depth $\leq d$. We define a value tree $t = (y, (T_n)_{n \in \mathbb{N}}) \in \mathcal{VT}_\Gamma^d$. The head-variable y is determined by the (unique) minimal positive move s of S . Either σs is the initial move in the copy of U^\perp corresponding to some variable $x \in \Gamma$: in this case we set $x = y$; or σs is mapped to a minimal positive move of \mathcal{U} , in which case we set $x = x_i^d$ where $\sigma s = (i, k)$ for some k .

By receptivity, s has ω immediate successors t_0, t_1, \dots . Each of them induces an essp \mathcal{S}_i , the subtree of S with root t_i , which can be turned into a strategy $\sigma_i : \mathcal{S}_i \rightarrow (\mathcal{U}^{\Gamma + \{x_k^d \mid k \in \mathbb{N}\}})^\perp \parallel \mathcal{U}$, such that σ_i is Nakajima-like (because σ is) and has depth $\leq d - 1$. Then, for each i , we define $T_i = \Psi_{\Gamma + \{x_k^d \mid k \in \mathbb{N}\}}^{d-1}(\sigma_i)$, and we have defined $\Phi_\Gamma^d(\sigma)$ as $t = (y, (T_n)_{n \in \mathbb{N}})$.

Suppose now that $\tau : \mathcal{T} \rightarrow (\mathcal{U}^\Gamma)^\perp \parallel \mathcal{U}$ is an arbitrary Nakajima-like strategy. Then it is of the form $\sum_{i \in I} p_i \cdot \sigma_i$, where for each $i \in I$, $\sigma_i : \mathcal{S}_i \rightarrow (\mathcal{U}^\Gamma)^\perp \parallel \mathcal{U}$ is a value Nakajima-like strategy of depth d , and moreover all the σ_i are distinct.

So, by the drop condition, a Nakajima-like probabilistic strategy corresponds to a probability distribution on value strategies. Using that Φ_{Γ}^d is a bijection, a probability distribution on value strategies is the same thing as one on value trees, and therefore Ψ_{Γ}^d is a bijection. \square

In the rest of the chapter, we show that this bijection preserves the interpretation of λ^+ -terms, in the following sense:

Theorem 4.27 (Correspondence theorem). *For any $M \in \Lambda^+$ and $d \in \mathbb{N}$, $\Psi_{\Gamma}^d(\llbracket M \rrbracket_{\text{nak}}^d) = \text{PT}^d(M)$, where $\llbracket M \rrbracket^d$ is the maximal sub-strategy of $\llbracket M \rrbracket$ with depth $\leq d$.*

To alleviate notation we will simply write Φ and Ψ for Φ_{Γ}^d and Ψ_{Γ}^d , when the context is clear. The next section gives a detailed proof and some immediate consequences.

4.5 The correspondence theorem

The proof is by induction on d , and follows a similar argument as in the non-probabilistic case [KNO02], with the additional difficulty of dealing with *infinite width*: a probabilistic Nakajima tree may be a probability distribution with infinite support, and the first level of Player moves in a probabilistic strategy may be infinite. The proof must therefore make use of finite-width *approximations*; but the need to approximate in both depth and width means that the order on probabilistic Nakajima trees is more subtle than the naive probabilistic extension of the “subtree” order. (For instance, $\text{PT}^d(M)$ is *not* in general a subtree of $\text{PT}^{d+1}(M)$.)

The intricacies of Nakajima tree approximation are dealt with in the work of Leventis in [Lev16]. In what follows, we describe the steps of his argument, and reproduce them through game semantics, to get the desired correspondence result.

4.5.1 Leventis’ method for approximating trees

For each $d \in \mathbb{N}$, Leventis [Lev16] proposes a method for approximating the tree $\text{PT}^d(M)$ using a sequence of trees $T_i \in \mathcal{PT}^d$ (for $i \in \mathbb{N}$) with finite support. (That is, the probability distributions $T_i : \mathcal{VT} \rightarrow [0, 1]$ have finite support. In this sense the T_i have “finite width”.)

There are three steps in the design of the method:

- For $d \in \mathbb{N}$, define a partial order \leq_d on \mathcal{PT}^d .
- Define for each term M a finite-width tree $\text{pt}^d(M) \in \mathcal{PT}^d$ called its *local tree*.
- Show that for each term M there is a sequence of terms $(L^i(M))_{i \in \omega}$ such that the sequence of trees $\text{pt}^d(L^i(M))$ approximates $\text{PT}^d(M)$ in the order \leq_d .

Because of the mutually recursive definition, each of the above must be done also for value trees. The details follow.

Ordering trees. Let $\varepsilon \in \mathbb{R}$ with $\varepsilon > 0$. Define by induction on $d \in \mathbb{N}$ two relations $\leq_{d,\varepsilon}$ and $\leq_{d,\varepsilon}^v$ on \mathcal{PT}^d and \mathcal{VT}^d , respectively, as follows:

$$\frac{\forall i < m, T_i \leq_{d,\varepsilon} T'_i}{(y, (T_1, \dots, T_m, x_0^d, \dots)) \leq_{d+1,\varepsilon}^v (y, (T'_1, \dots, T'_m, x_0^d, \dots))}$$

$$\frac{\forall A \subseteq \mathcal{VT}^d, \sum_{t \in \uparrow A} T(t) \leq \sum_{t \in \uparrow_\varepsilon A} T'(t) + \varepsilon}{T \leq_{d,\varepsilon} T'}$$

where $\uparrow_\varepsilon A = \{t \in \mathcal{VT}^d \mid \exists t' \in A \text{ with } t' \leq_{d,\varepsilon}^v t\}$ and $\uparrow A = \uparrow_0 A$.

Define $\leq_d = \leq_{d,0}$ and $\leq_d^v = \leq_{d,0}^v$. Leventis shows the two relations are indeed partial orders, and he gives sufficient conditions for $T \in \mathcal{PT}^d$ to be the lub (w.r.t. \leq_d) of a sequence $(T_n)_{n \in \omega}$ of trees, namely:

- for all $n \in \omega$, $T_n \leq_d T$; and
- for all $\varepsilon > 0$, there exists $n \in \omega$ such that $T \leq_{d,\varepsilon} T_n$.

Local trees. Recall that for a term M and value tree $t \in \mathcal{VT}^d$, $\text{PT}^d(M)(t)$ is the probability that M converges (by head-reduction) to a hnf h with $\text{VT}^d(h) = t$. In contrast, the *local* Nakajima tree $\text{pt}^d(M)$ does not model the convergence behaviour of M , but instead gathers static information about the term structure. To obtain this we consider the **canonical form** of M , denoted $\text{can}(M)$ and defined to be its normal form with respect to the reduction

$$\lambda x.(M +_p N) \mapsto \lambda x.M +_p \lambda x.N \quad (M +_p N)P \mapsto MP +_p NP$$

which is strongly normalising and confluent [Lev16]. This pushes some instances of $+_p$ to the top level, and indeed **canonical terms** M, N are obtained as sums of **values**, as in the following grammar:

$$M, N ::= v \mid M +_p N$$

$$v ::= x \mid vM \mid \lambda x.v$$

Note that head-normal forms are values, but not all values are in head-normal form.

To enforce associativity, commutativity and idempotence of $+_p$ one can consider canonical terms up to the following notion of *syntactic equivalence*:

Definition 4.28. Let \equiv_{syn} be the equivalence relation on λ^+ -terms generated by:

$$M +_p N \equiv_{\text{syn}} N +_{1-p} M$$

$$(M +_p N) +_q P \equiv_{\text{syn}} M +_{pq} (N +_{\frac{(1-p)q}{1-pq}} P) \quad \text{if } pq \neq 1$$

$$M +_p M \equiv_{\text{syn}} M$$

$$M +_1 N \equiv_{\text{syn}} M +_1 P$$

Note that this is not a congruence; we are only concerned with “top-level” sums. Canonical terms up to \equiv_{syn} correspond precisely to finitely-supported probability distributions on values, written as convex sums $\sum_{i \in I} p_i v_i$. (This representation will be used in the definition of local trees.)

We note in passing that the semantics of λ^+ -terms is invariant under \mapsto , and that \equiv_{syn} -equivalent terms have Nakajima-equivalent interpretation:

Lemma 4.29. *Let M and N be terms with free variables in Γ . Then:*

- *If $M \mapsto N$, then $\llbracket M \rrbracket^\Gamma \cong \llbracket N \rrbracket^\Gamma$, and so $\llbracket M \rrbracket_{\text{nak}}^\Gamma \cong \llbracket N \rrbracket_{\text{nak}}^\Gamma$.*
- *If $M \equiv_{\text{syn}} N$, then $\llbracket M \rrbracket_{\text{nak}}^\Gamma \cong \llbracket N \rrbracket_{\text{nak}}^\Gamma$.*

Proof. Using Lemma 2.27, it is routine to check that $\lambda \odot (\sigma +_p \tau) \cong (\lambda \odot \sigma) +_p (\lambda \odot \tau)$ and $\text{app} \odot (\sigma' +_p \tau') \cong (\text{app} \odot \sigma') +_p (\text{app} \odot \tau')$ for any σ, τ, σ' and τ' (with the appropriate type). Similarly $\text{Cur}(\sigma +_p \tau) = \text{Cur}(\sigma) +_p \text{Cur}(\tau)$, and from this we easily deduce $\llbracket \lambda x.M +_p N \rrbracket^\Gamma \cong \llbracket \lambda x.M +_p \lambda x.N \rrbracket^\Gamma$. To show that $\llbracket (M +_p N)P \rrbracket^\Gamma \cong \llbracket (MP) +_p (NP) \rrbracket^\Gamma$, it is enough to verify that $\text{Ev} \odot \langle \text{app} \odot \llbracket M \rrbracket^\Gamma +_p \text{app} \odot \llbracket N \rrbracket^\Gamma, \llbracket P \rrbracket^\Gamma \rangle \cong \text{Ev} \odot \langle \text{app} \odot \llbracket M \rrbracket^\Gamma, \llbracket P \rrbracket^\Gamma \rangle +_p \text{Ev} \odot \langle \text{app} \odot \llbracket N \rrbracket^\Gamma, \llbracket P \rrbracket^\Gamma \rangle$, which can be done by a straightforward inspection.

For \equiv_{syn} , we go through the equations in the definition above. There are obvious isomorphisms of strategies $\sigma +_p \tau \cong \tau +_{(1-p)} \sigma$ and $(\sigma +_p \tau) +_q \rho \Rightarrow \sigma +_{pq} (\tau +_{\frac{(1-p)q}{1-pq}} \rho)$. Moreover, the Nakajima quotient σ_{nak} was defined specifically so that $(\sigma +_p \sigma)_{\text{nak}} \cong \sigma_{\text{nak}}$ and $(\sigma +_1 \tau)_{\text{nak}} = \sigma_{\text{nak}}$, from which the result follows. \square

We now define local Nakajima trees and value trees. Leventis’ original presentation is slightly different, but our definition requires fewer technical tools.

Definition 4.30. Suppose $M \in \lambda^+$, and let $\sum_{i \in I} p_i v_i$ be the probability distribution corresponding to $\text{can}(M)$. Then, for $d \in \mathbb{N}$, the **local probabilistic Nakajima tree of M of depth d** is the tree $\text{pt}^d(M) \in \mathcal{PT}^d$ defined as follows:

$$\begin{aligned} \text{pt}^d(M) : \mathcal{VT}^d &\longrightarrow [0, 1] \\ t &\longmapsto \sum_{\substack{i \in I, v_i \text{ hnf} \\ \text{vt}^d(v_i) = t}} p_i, \end{aligned}$$

where for any hnf $H = \lambda z_0 \dots z_{n-1}.y P_0 \dots P_{k-1}$, the **local value tree of depth d of H** is defined as

$$\text{vt}^d(H) = (y, (\text{pt}^{d-1}(P_0), \dots, \text{pt}^{d-1}(P_{k-1}), \text{pt}^{d-1}(x_n^d), \dots)).$$

The finite-width approximants. The final step is to identify a sequence of trees t_i which approximate $\text{PT}^d(M)$ in the order \leqslant_d defined above. These are obtained as the local Nakajima trees for a sequence of canonical terms $L^i(M)$, defined below. The convergence proof uses that for any $\varepsilon \geqslant 0$ there is $k \in \mathbb{N}$ such that $L^k(M)$ is in normal form “up to ε ”:

Definition 4.31. For $d \in \mathbb{N}$ and $\varepsilon \geq 0$, the sets of **canonical d, ε -head-normal forms** and **d, ε -normal values** are defined by:

$$\begin{aligned} \text{NF}_v^{0,\varepsilon} &= \{v \mid v \text{ value}\} \\ \text{NF}_v^{d+1,\varepsilon} &= \{\lambda x_0 \dots x_{n-1}.yP_0 \dots P_{m-1} \mid \forall i, P_i \in \text{NF}^{d,\varepsilon}\} \\ \text{NF}^{d,\varepsilon} &= \{M \mid M \equiv_{\text{syn}} (\sum_{i \in I} p_i v_i) + (1 - \sum_i p_i)P, \\ &\quad \text{with } v_i \in \text{NF}_v^{d,\varepsilon} \text{ and } (1 - \sum_i p_i) \text{Pr}_{\downarrow}(P) \leq \varepsilon\}, \end{aligned}$$

using the n -ary sum notation.

The terms $L^i(M)$ are defined by means of the reduction \rightarrow_L between canonical terms defined as follows:

$$\frac{\frac{\frac{M_1 \rightarrow_L N_1 \quad M_2 \rightarrow_L N_2}{M_1 +_p M_2 \rightarrow_L N_1 +_p N_2}}{\forall i. P_i \rightarrow_L Q_i}}{\lambda x_0 \dots x_{n-1}.yP_0 \dots P_{m-1} \rightarrow_L \lambda x_0 \dots x_{n-1}.yQ_0 \dots Q_{m-1}}}{\lambda x_0 \dots x_{n-1}.(\lambda y.M)PQ_0 \dots P_{m-1} \rightarrow_L \text{can}(\lambda x_0 \dots x_{n-1}.(M[P/y])Q_0 \dots Q_{m-1})}$$

Terms of λ^+ have the following property.

Lemma 4.32 ([Lev16]). *For every canonical term N there is a unique term $L(N)$ such that $N \rightarrow_L L(N)$. For a (not necessarily canonical) term M , let $L^0(M) \triangleq \text{can}(M)$ and $L^{i+1}(M) \triangleq L(L^i(M))$. Then, for every $d \in \mathbb{N}$ and $\varepsilon \geq 0$, there exists $k \in \mathbb{N}$ such that $L^k(M) \in \text{NF}^{d,\varepsilon}$.*

One can then show that the necessary convergence conditions laid out above are satisfied, so that:

Proposition 4.33 ([Lev16]). *For any term M and $d \in \mathbb{N}$,*

$$PT^d(M) = \bigvee_{i \in \omega} \text{pt}^d(L^i(M)).$$

We now reproduce the steps above with Nakajima-like strategies, aiming to show a correspondence theorem. It will be helpful to note the following:

Lemma 4.34. *For every $M \in \lambda^+_{\Gamma}$ and $i \in \omega$, $\llbracket M \rrbracket^{\Gamma} = \llbracket L^i(M) \rrbracket^{\Gamma}$.*

Proof. Routine verification. □

4.5.2 Local semantics

We proceed to define for each canonical $M \in \boldsymbol{\lambda}^+$ with free variables in Γ a strategy $[M]^\Gamma \in \mathbf{PG}_!^{\text{si}}[\mathcal{U}^\Gamma, \mathcal{U}]$. The definition is similar to that of $\llbracket M \rrbracket^\Gamma$, with the difference that sub-terms of the form $(\lambda x.M)N$ are assigned a trivial strategy. Thus $[M]^\Gamma$ provides no information about the behaviour of M under β -reduction. This is reminiscent of the definition of the local tree $\text{pt}^d(M)$, and indeed $([M]_d^\Gamma)_{\text{nak}}$ will be its semantic counterpart; see Lemma 4.37 below.

For a canonical $M \in \boldsymbol{\lambda}^+_\Gamma$ we define:

$$[M]^\Gamma = \begin{cases} \pi_x & \text{if } M = x \\ \lambda \odot_! (\Lambda([N]^{x,\Gamma})) & \text{if } M = \lambda x.N \\ \text{Ev} \odot_! \langle \text{app} \odot_! [N]^\Gamma, [P]^\Gamma \rangle & \text{if } M = NP \text{ with } N \neq \lambda x.N' \\ [N]^\Gamma +_p [P]^\Gamma & \text{if } M = N +_p P \\ \perp & \text{otherwise.} \end{cases}$$

The definition makes use of the strategy \perp , which is initial in the category $\mathbf{PG}_!^{\text{si}}[\mathcal{U}^\Gamma, \mathcal{U}]$. This is only defined up to isomorphism, so for a concrete construction we may take $\perp : \{\varepsilon\} \rightarrow !(\mathcal{U}^\Gamma)^\perp \parallel \mathcal{U}$, where ε is the unique initial move of \mathcal{U} , and the map has evident action. We note the following property:

Lemma 4.35. *If $M, N \in \boldsymbol{\lambda}^+_\Gamma$ with $M \equiv_{\text{syn}} N$, then $[M]^\Gamma =_{PT} [N]^\Gamma$.*

Proof. Straightforward. \square

We will show that the strategy $([M]_d^\Gamma)_{\text{nak}}$ corresponds to the local tree $\text{pt}^d(M)$ under the map Ψ^d . Since the strategies and Nakajima trees involved have finite width, the proof is a straightforward extension of Ker, Nickau and Ong's "exact correspondence theorem" for the λ -calculus [KNO02].

We first prove a technical lemma.

Lemma 4.36. *If $v \in \boldsymbol{\lambda}^+$ is a value which is not in head-normal form, then $[v]^\Gamma \cong \perp$.*

Proof. If $v = x$ then it is in hnf. Suppose $v = v'M$ for M a canonical term. If v is not in hnf, then either $v' = (\lambda x.N)$ in which case $[M]^\Gamma = \perp$ by definition, or v' is itself not in hnf, so by the induction hypothesis $[v']^\Gamma \cong \perp$. Then $[v'M]^\Gamma \cong \text{Ev} \odot_! \langle \text{app} \odot_! \perp, [M]^\Gamma \rangle \cong \text{Ev} \odot_! \langle \perp, [M]^\Gamma \rangle$ which we can easily check to be isomorphic to \perp .

Finally if $v = \lambda x.v'$ and v is not in hnf, then v' is not in hnf so $[v']^{\Gamma,x} \cong \perp$ and we conclude as above because $\lambda \odot_! (\text{Cur}(\perp)) \cong \perp$. \square

We are ready to state the correspondence result for local trees and strategies:

Lemma 4.37. *For every $d \in \mathbb{N}$, the following hold:*

- For every hnf $H \in \boldsymbol{\lambda}^+_\Gamma$, $[H]_d^\Gamma$ (and therefore $([H]_d^\Gamma)_{\text{nak}}$) is a value strategy, and $\Phi^d(([H]_d^\Gamma)_{\text{nak}}) = \text{vt}^d(H)$.
- For every canonical term $M \in \boldsymbol{\lambda}^+_\Gamma$, $\Psi^d(([M]_d^\Gamma)_{\text{nak}}) = \text{pt}^d(M)$.

Proof. The proof is by induction on d , and the base case ($d = 0$) is immediate. Suppose $d > 0$.

If H is a hnf, say $H = \lambda x_0 \dots x_{n-1}. y P_0 \dots P_{m-1}$, then the strategy $[H]^\Gamma$ is obtained from $[y P_0 \dots P_{m-1}]^{\Gamma, x_0, \dots, x_{n-1}}$ by repeated currying and composition with the strategy λ associated with the reflexive object \mathcal{U} . Write $\Delta = \Gamma, x_0, \dots, x_{n-1}$. Then, the strategy $[y P_0 \dots P_{m-1}]^\Delta : \mathcal{S} \rightarrow !(\mathcal{U}^\Delta)^\perp \parallel \mathcal{U}$ is a value strategy whose top-level Player move corresponds to the variable y . The subsequent ω Opponent moves correspond to the arguments given to y : the first m branches are the $[P_i]^\Delta$, and the following branches are the appropriate copycat strategies. This corresponds precisely to $\text{vt}^d(H)$. (Remark: this deterministic step is the content of the correspondence result in [KNO02], of which we have given a very informal account.)

Now let M be an arbitrary canonical term. We show the result by induction on the term structure: M is either a value or of the form $N +_p P$ for canonical terms N and P . If M is a value not in head-normal form, then by Lemma 4.36 $[M]^\Gamma \cong \perp$, so in particular $[M]_{\text{nak}}^\Gamma \cong \perp$. It follows from the definition that $\text{pt}^d(M)(t) = 0$ for all t , so the result holds. If M is a value in head-normal form, then the result follows from the reasoning of the previous paragraph, using that $\text{pt}^d(M)(t) = 1$ if $t = \text{vt}^d(M)$, 0 otherwise.

Finally suppose $M = N +_p P$ for canonical terms N and P . By the induction hypothesis, $\Psi(([N]_d^\Gamma)_{\text{nak}}) = \text{pt}^d(N)$ and $\Psi(([P]_d^\Gamma)_{\text{nak}}) = \text{pt}^d(P)$. Let $t \in \mathcal{VT}^d$. Recall from Lemma 4.26 that to obtain the coefficient $\Psi(([M +_p P]_d^\Gamma)_{\text{nak}})(t)$, we regard $([N +_p P]_d^\Gamma)_{\text{nak}}$ as a convex sum of value strategies, and sum over those corresponding to t under Φ^d . Because the Nakajima quotient only identifies branches corresponding to the same value tree, this process can be done directly on the strategy $[N +_p P]_d^\Gamma$. But $[N +_p P]_d^\Gamma = [N]_d^\Gamma +_p [P]_d^\Gamma$, and so $\Psi(([N +_p P]_d^\Gamma)_{\text{nak}})(t) = p \cdot \Psi(([N]_d^\Gamma)_{\text{nak}})(t) + (1 - p) \cdot \Psi(([P]_d^\Gamma)_{\text{nak}})(t)$. It is clear that $\text{pt}^d(N +_p P)(t) = p \cdot \text{pt}^d(N)(t) + (1 - p) \cdot \text{pt}^d(P)(t)$, so this concludes the proof. \square

Lemma 4.38. *For every canonical $M \in \lambda^+_\Gamma$, there is an embedding $[M]^\Gamma \Rightarrow [[M]^\Gamma]$.*

Proof. Strong embeddings are closed under horizontal composition and pairing, and there is an embedding $\perp \Rightarrow \sigma$ for any strategy σ , so the result holds by induction on M , inspecting the cases in the definition of $[M]^\Gamma$. \square

Lemma 4.39. *If $\sigma, \tau \in \mathbf{PG}_1^{\text{si}}[\mathcal{U}^\Gamma, \mathcal{U}]$ and there is an embedding $\sigma \Rightarrow \tau$, then $\Psi((\sigma_d)_{\text{nak}}) \leq_d \Psi((\tau_d)_{\text{nak}})$ for each d . If $\sigma, \tau \in \mathbf{PG}_1^{\text{si}}[\mathcal{U}^\Gamma, \mathcal{U}]$ are value strategies and there is an embedding $\sigma \Rightarrow \tau$, then $\Phi((\sigma_d)_{\text{nak}}) \leq_d^v \Phi((\tau_d)_{\text{nak}})$ for each d .*

Proof. Mutual induction on d , where the base case holds immediately.

Let $d > 0$ and let σ and τ be value strategies. The embedding must preserve the label of the unique Player move at depth 1, and that of its immediate successors. Restricting the embedding to any branch, using that each branch is a substrategy of depth $\leq d - 1$, we can apply the induction hypothesis and get the desired result, since by definition \leq_d^v is a branch-wise \leq_d .

Now, let σ and τ be arbitrary strategies, and let $A \subseteq \mathcal{VT}^d$. Write σ_d as a convex sum $\sum_{i \in I} p_i \cdot \sigma_i$ of value sub-strategies, and similarly let $\tau_d = \sum_{j \in I} q_j \cdot \tau_j$. The embedding of the statement can then be seen as consisting of an injection $\iota : I \rightarrow J$,

together with embeddings $\sigma_i \Rightarrow \tau_{\iota(i)}$ such that $p_i \leq q_{\iota(i)}$ for each $i \in I$. By the reasoning in the value case above, we have for each i that $\Phi((\sigma_i)_{\text{nak}}) \leq_d^v \Phi((\tau_{\iota(i)})_{\text{nak}})$. So if $\Phi((\sigma_i)_{\text{nak}}) \in \uparrow A$, then also $\Phi((\tau_{\iota(i)})_{\text{nak}}) \in \uparrow A$. Thus,

$$\begin{aligned} \sum_{t \in \uparrow A} \Psi((\sigma_d)_{\text{nak}})(t) &= \sum_{\substack{i \in I \\ \Phi((\sigma_i)_{\text{nak}}) \in \uparrow A}} p_i \\ &\leq \sum_{\substack{i \in I \\ \Phi((\tau_{\iota(i)})_{\text{nak}}) \in \uparrow A}} q_{\iota(i)} \\ &\leq \sum_{t \in \uparrow A} \Psi((\tau_d)_{\text{nak}})(t). \end{aligned}$$

This shows that $\Psi((\sigma_d)_{\text{nak}}) \leq_d \Psi((\tau_d)_{\text{nak}})$. \square

Hence, using the embedding of Lemma 4.38, we get:

Corollary 4.40. *For every M , $\Psi^d((\llbracket M \rrbracket_d^\Gamma)_{\text{nak}}) \leq_d \Psi^d((\llbracket M \rrbracket_d^\Gamma)_{\text{nak}})$, and for every hnf H , $\Phi^d((\llbracket H \rrbracket_d^\Gamma)_{\text{nak}}) \leq_d \Phi^d((\llbracket H \rrbracket_d^\Gamma)_{\text{nak}})$.*

4.5.3 The strategy $\llbracket M \rrbracket_{\text{nak}}^\Gamma$ as a lub

Lemma 4.41. *If $M \in \text{NF}^{d,\varepsilon}$, then $\Psi^d((\llbracket M \rrbracket_d^\Gamma)_{\text{nak}}) \leq_{d,\varepsilon} \Psi^d((\llbracket M \rrbracket_d^\Gamma)_{\text{nak}})$.*

Proof. We show the statement by induction on d , together with the corresponding property for values: if $v \in \text{NF}_v^{d,\varepsilon}$, then $\Phi^d((\llbracket v \rrbracket_d^\Gamma)_{\text{nak}}) \leq_{d,\varepsilon} \Phi^d((\llbracket v \rrbracket_d^\Gamma)_{\text{nak}})$.

The case $d = 0$ is immediate. For $d > 0$ we consider values first: for $v = \lambda x_0 \dots x_{n-1}. y P_0 \dots P_{m-1}$ we have:

$$\begin{aligned} \Phi((\llbracket v \rrbracket_d^\Gamma)_{\text{nak}}) &= (y, \Psi((\llbracket P_0 \rrbracket_{d-1}^\Gamma)_{\text{nak}}), \dots, \Psi((\llbracket P_{m-1} \rrbracket_{d-1}^\Gamma)_{\text{nak}}), \text{PT}^{d-1}(x_n^d), \dots) \\ \Phi^d((\llbracket v \rrbracket_d^\Gamma)_{\text{nak}}) &= (y, \Psi^d((\llbracket P_0 \rrbracket_{d-1}^\Gamma)_{\text{nak}}), \dots, \Psi^d((\llbracket P_{m-1} \rrbracket_{d-1}^\Gamma)_{\text{nak}}), \text{PT}^{d-1}(x_n^d), \dots). \end{aligned}$$

Therefore, by the induction hypothesis and the definition of $\leq_{d,\varepsilon}^v$,

$$\Phi((\llbracket v \rrbracket_d^\Gamma)_{\text{nak}}) \leq_{d,\varepsilon}^v \Phi^d((\llbracket v \rrbracket_d^\Gamma)_{\text{nak}}).$$

We move to the case of a general $M \in \text{NF}^{d,\varepsilon}$. By definition, $M \equiv_{\text{syn}} (\sum_{i \in I} p_i v_i) + (1 - \sum_i p_i) P$ where all $v_i \in \text{NF}_v^{d,\varepsilon}$ and $(1 - \sum_i p_i) \text{Pr}_\downarrow(P) \leq \varepsilon$. By Lemma 4.29, we have

$$\llbracket M \rrbracket_d^\Gamma =_{\text{PT}} \sum_{i \in I} p_i (\llbracket v_i \rrbracket_d^\Gamma)_{\text{nak}} + (1 - \sum_i p_i) (\llbracket P \rrbracket_d^\Gamma)_{\text{nak}}, \quad (4.1)$$

and the analogous statement for $\llbracket M \rrbracket_d^\Gamma$. Write $T = \Psi((\llbracket M \rrbracket_d^\Gamma)_{\text{nak}})$ and $T' = \Psi^d((\llbracket M \rrbracket_d^\Gamma)_{\text{nak}})$. To show $T \leq_{d,\varepsilon} T'$, let A be an arbitrary subset of $\mathcal{V}\mathcal{T}^d$. We show that $\sum_{t \in \uparrow A} T(t) \leq \sum_{t \in \uparrow_{\varepsilon} A} T'(t) + \varepsilon$.

We easily deduce from (4.1) that

$$\sum_{t \in \uparrow A} T(t) = \sum_{\substack{i \in I \\ \Phi((\llbracket v_i \rrbracket_d^\Gamma)_{\text{nak}}) \in \uparrow A}} p_i + (1 - \sum_{i \in I} p_i) \sum_{t \in \uparrow A} \Psi((\llbracket P \rrbracket_d^\Gamma)_{\text{nak}})(t).$$

Recall that the definition of Ψ involves the first layer of Player moves, so it is clear that

$$\sum_{t \in \uparrow A} \Psi(\llbracket P \rrbracket_d^\Gamma)_{\text{nak}}(t) \leq \text{Pr}_\downarrow(\llbracket P \rrbracket_d^\Gamma)_{\text{nak}} = \text{Pr}_\downarrow(\llbracket P \rrbracket^\Gamma)$$

where the equality is because $d > 0$. By the adequacy property (Theorem 4.11), $\text{Pr}_\downarrow(\llbracket P \rrbracket^\Gamma) \leq \text{Pr}_\downarrow(P)$, so that

$$(1 - \sum_{i \in I} p_i) \sum_{t \in \uparrow A} \Psi(\llbracket P \rrbracket_d^\Gamma)_{\text{nak}}(t) \leq \varepsilon.$$

Thus

$$\sum_{t \in \uparrow A} T(t) \leq \sum_{\substack{i \in I \\ \Phi(\llbracket v_i \rrbracket_d^\Gamma)_{\text{nak}} \in \uparrow A}} p_i + \varepsilon.$$

Now, the property we have shown above for d, ε -normal values indicates that for every v_i ,

$$\Phi(\llbracket v_i \rrbracket_d^\Gamma)_{\text{nak}} \leq_{d, \varepsilon}^v \Phi(\llbracket v_i \rrbracket_d^\Gamma)_{\text{nak}},$$

and therefore if $\Phi(\llbracket v_i \rrbracket_d^\Gamma)_{\text{nak}} \in \uparrow A$ then $\Phi(\llbracket v_i \rrbracket_d^\Gamma)_{\text{nak}} \in \uparrow_\varepsilon A$. From this we deduce that

$$\sum_{t \in \uparrow A} T(t) \leq \sum_{\varphi(\llbracket v_i \rrbracket_d^\Gamma)_{\text{nak}} \in \uparrow_\varepsilon A} p_i + \varepsilon,$$

and since

$$\sum_{t \in \uparrow_\varepsilon A} T'(t) = \sum_{\substack{i \in I \\ \Phi(\llbracket v_i \rrbracket_d^\Gamma)_{\text{nak}} \in \uparrow_\varepsilon A}} p_i + (1 - \sum_{i \in I} p_i) \sum_{t \in \uparrow_\varepsilon A} \Psi(\llbracket P \rrbracket_d^\Gamma)_{\text{nak}}(t),$$

we have $\sum_{t \in \uparrow A} T(t) \leq \sum_{t \in \uparrow_\varepsilon A} T'(t) + \varepsilon$, as required. \square

We are finally in a position to show the following crucial step in the proof of the correspondence theorem:

Theorem 4.42. *For any $M \in \lambda^+_\Gamma$,*

$$\Psi(\llbracket M \rrbracket_d^\Gamma)_{\text{nak}} = \bigvee_{i \in \omega} \Psi(\llbracket L^i(M) \rrbracket_d^\Gamma)_{\text{nak}}$$

Proof. First, for every $i \in \omega$, $\Phi(\llbracket L^i(M) \rrbracket_d^\Gamma)_{\text{nak}} \leq_d \Phi(\llbracket L^i(M) \rrbracket_d^\Gamma)_{\text{nak}}$ by Corollary 4.40, and since $\llbracket L^i(M) \rrbracket^\Gamma = \llbracket M \rrbracket^\Gamma$ (Lemma 4.34), $\Phi(\llbracket L^i(M) \rrbracket_d^\Gamma)_{\text{nak}} \leq_d \Phi(\llbracket M \rrbracket_d^\Gamma)_{\text{nak}}$. This shows that $\Phi(\llbracket M \rrbracket_d^\Gamma)_{\text{nak}}$ is an upper bound for the chain.

To show it is the *least*, by [Lev16], it suffices to show that for any $\varepsilon > 0$, there exists $k \in \mathbb{N}$ such that $\Phi(\llbracket M \rrbracket_d^\Gamma)_{\text{nak}} \leq_{d, \varepsilon} \Phi(\llbracket L^k(M) \rrbracket_d^\Gamma)_{\text{nak}}$. By Lemma 4.32, there exists $k \in \mathbb{N}$ such that $L^k(M) \in \text{NF}^{d, \varepsilon}$. Lemma 4.41 says $\Phi(\llbracket L^k(M) \rrbracket_d^\Gamma)_{\text{nak}} \leq_{d, \varepsilon} \Phi(\llbracket L^k(M) \rrbracket_d^\Gamma)_{\text{nak}}$, and we conclude since $\llbracket L^k(M) \rrbracket^\Gamma = \llbracket M \rrbracket^\Gamma$. \square

4.5.4 Wrapping up and full abstraction

Theorem 4.43 (Correspondence theorem). *For any $M \in \mathbf{\lambda}^+_{\Gamma}$ and $d \in \mathbb{N}$, $\Psi(\llbracket M \rrbracket_d^{\Gamma})_{\text{nak}} = PT^d(M)$.*

Proof. For any M ,

$$\begin{aligned} \Psi(\llbracket M \rrbracket_d^{\Gamma})_{\text{nak}} &= \bigvee_{i \in \omega} \Psi(\llbracket L^i(M) \rrbracket_d^{\Gamma})_{\text{nak}} && \text{(Theorem 4.42)} \\ &= \bigvee_{i \in \omega} \text{pt}^d(L^i(M)) && \text{(Lemma 4.37)} \\ &= PT^d(M) && \text{(Proposition 4.33)}. \end{aligned}$$

□

We combine this with Leventis' result, to get:

Theorem 4.44 (Full abstraction). *For any $M, N \in \Lambda^+$, the following are equivalent:*

- $M \cong N$;
- $M =_{PT} N$;
- $\llbracket M \rrbracket =_{PT} \llbracket N \rrbracket$.

Chapter 5

Probabilistic PCF

Following the game semantics tradition, we investigate an application of concurrent games to the functional language PCF [Plo77]. Specifically we see how to obtain using probabilistic strategies a model for terms of *Probabilistic PCF*, *i.e.* PCF extended with a probabilistic primitive. This leads to an “intensional” full abstraction result.

5.1 Syntax and operational semantics

Probabilistic PCF (PPCF) extends PCF with a probabilistic Boolean **coin** which gives **tt** or **ff** with equal probability. The choice of **coin** over another probabilistic primitive (such as the $+_p$ operator of Chapter 4), is not particularly significant: although there are differences in expressivity, it can be shown ([ETP14]) that this does not affect contextual equivalence.

The types of PPCF are those of PCF, so

$$A, B ::= \mathbf{Bool} \mid \mathbf{Nat} \mid A \rightarrow B.$$

Its terms are those of PCF, augmented with the Boolean primitive **coin**:

$$\begin{aligned} M, N_1, N_2 ::= & x \mid \lambda x.M \mid M N \mid \underline{n} \mid \underline{b} \mid \mathbf{succ} M \mid \mathbf{pred} M \mid \mathbf{if} M N_1 N_2 \\ & \mid \mathbf{iszero} M \mid Y M \mid \mathbf{coin} \end{aligned}$$

where n and b range over natural numbers and Booleans, respectively. Typing rules are standard; we only show that for **coin**:

$$\overline{\Gamma \vdash \mathbf{coin} : \mathbf{Bool}}$$

As for λ^+ , the operational semantics is given by means of a weighted reduction relation, defined by the following rules:

$$\begin{aligned} (\lambda x.M)N &\xrightarrow{1} M[N/x] & \mathbf{if} \mathbf{tt} N_1 N_2 &\xrightarrow{1} N_1 & \mathbf{if} \mathbf{ff} N_1 N_2 &\xrightarrow{1} N_2 \\ \mathbf{iszero} \underline{0} &\xrightarrow{1} \mathbf{tt} & \mathbf{iszero} \underline{n+1} &\xrightarrow{1} \mathbf{ff} & Y M &\xrightarrow{1} M(Y M) \\ \mathbf{pred} (\underline{n+1}) &\xrightarrow{1} \underline{n} & \mathbf{succ} \underline{n} &\xrightarrow{1} \underline{n+1} & \mathbf{coin} &\xrightarrow{\frac{1}{2}} \underline{b} \text{ if } b \in \{\mathbf{tt}, \mathbf{ff}\} \end{aligned}$$

$$\begin{array}{c}
\frac{M \xrightarrow{p} M'}{M N \xrightarrow{p} M' N} \quad \frac{M \xrightarrow{p} M'}{\mathbf{succ} M \xrightarrow{p} \mathbf{succ} M'} \quad \frac{M \xrightarrow{p} M'}{\mathbf{pred} M \xrightarrow{p} \mathbf{pred} M'} \\
\\
\frac{M \xrightarrow{p} M'}{\mathbf{iszero} M \xrightarrow{p} \mathbf{iszero} M'} \quad \frac{M \xrightarrow{p} M'}{\mathbf{if} M N P \xrightarrow{p} \mathbf{if} M' N P}
\end{array}$$

This gives for every M and N a probability of reduction $\Pr(M \rightarrow N)$, got by summing over all reduction paths (as done in 4.1).

In λ^+ , observational equivalence was defined with as observables the head-normal forms; in PPCF, observables are ground type values:

Definition 5.1. Let M and N be PPCF terms such that $\Gamma \vdash M : A$ and $\Gamma \vdash N : A$. We write $M \lesssim_{\text{ctx}} N$ if for every context $C[\cdot]$ such that $\vdash C[P] : \mathbf{Bool}$ for every $\Gamma \vdash P : A$,

$$\Pr(C[M] \rightarrow \underline{b}) \leq \Pr(C[N] \rightarrow \underline{b})$$

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

5.2 Arenas with questions and answers

PPCF types are interpreted as arenas. There are two ground type arenas,

$$\begin{array}{c}
\mathbf{q}^- \\
\swarrow \quad \searrow \\
\mathbf{tt}^+ \sim \mathbf{ff}^+
\end{array}
\quad
\begin{array}{c}
\mathbf{q}^- \\
\swarrow \quad \searrow \quad \nearrow \\
0^+ \sim 1^+ \sim 2^+ \dots
\end{array}$$

from which all types are interpreted inductively using $\llbracket A \rightarrow B \rrbracket = \llbracket A \rrbracket \rightarrow \llbracket B \rrbracket$ ($= !\llbracket A \rrbracket \multimap \llbracket B \rrbracket$). Observe, in both cases, the contrast between the initial *question* \mathbf{q} , indicating the start of the computation, and the *answers* $\mathbf{tt}, \mathbf{ff}, 0, 1, \text{etc.}$, indicating that the computation is terminated and returns with a particular value. This intuition is still valid for higher-order types, see for example the arena for $\mathbf{Bool} \rightarrow \mathbf{Bool}$:

$$\begin{array}{c}
\mathbf{q}^- \\
\swarrow \quad \searrow \quad \nearrow \\
\mathbf{q}_0^+ \quad \mathbf{q}_1^+ \quad \dots \quad \mathbf{q}^- \\
\swarrow \quad \searrow \quad \swarrow \quad \searrow \quad \swarrow \quad \searrow \\
\mathbf{tt}_0^- \sim \mathbf{ff}_0^- \quad \mathbf{tt}_1^- \sim \mathbf{ff}_1^- \quad \mathbf{tt}^+ \sim \mathbf{ff}^+
\end{array}$$

On the argument side, the roles of Player and Opponent are reversed: Player starts the computation and Opponent provides a value.

This duality was not apparent in the semantics of λ^+ . Since the language has no constants, there are no ‘return values’ and it is sufficient for the model to describe the dynamics of variable calls. (In other words, all moves are questions.)

Definition 5.2. A **arena with questions and answers** is an arena \mathcal{A} equipped with a labelling function $\text{lbl} : A \rightarrow \{\mathcal{Q}, \mathcal{A}\}$ making each move either a **question** or an **answer**, and such that:

- initial moves are questions;
- answers are maximal;
- bijections in \cong_A preserve the labelling;

The addition of questions and answers to the objects of **PG** poses no difficulty; all constructions extend in the obvious way. (Note that in the dual game \mathcal{A}^\perp polarity is reversed but the question/answer labelling remains the same as in \mathcal{A} .) The definition of a strategy $\sigma : \mathcal{S} \rightarrow \mathcal{A}$ is not affected, but the labelling can be lifted: $s \in S$ is assigned the same label as its image $\sigma(s)$.

With this extra structure in place we can apply a *well-bracketing* condition on strategies, which ensures that Player respects the call/return discipline. The condition is the same as in the concurrent games model of PCF [CCW15]; it is not affected by the presence of probability.

Before we give the condition let us set some terminology. If $\sigma : \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$ is a strategy and the arenas \mathcal{A} and \mathcal{B} have questions and answers (assume from now on all arenas have questions and answers), for $X \in \text{Con}_S$ a question $s \in X$ is **answered** in X if there is an answer $s' \in X$ such that $\sigma(s) \rightarrow \sigma(s')$. We say X is **complete** if every question in X is answered in X . For a gcc ρ in S , the **pending question** of ρ is, if it exists, the latest (*i.e.* maximal) unanswered question $\rho_i \in \rho$.

Definition 5.3. A visible strategy $\sigma : \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$ is **well-bracketed** if the following two conditions are met:

- For every gcc $\rho = \{\rho_1 \rightarrow \dots \rightarrow \rho_n \rightarrow \rho_{n+1}\}$ in S , if ρ_{n+1} is an answer, then it answers the pending question in $\rho_1 \rightarrow \dots \rightarrow \rho_n$.
- If ρ and ζ are gccs of S , forking at $\rho_i = \zeta_i$, and merging at s as follows, with both branches disjoint,

$$\begin{array}{c}
 \rho_1 \triangleright \dots \triangleright \rho_i \begin{array}{l} \nabla \\ \triangle \end{array} \begin{array}{l} \rho_{i+1} \rightarrow \dots \triangleright \rho_n \\ \zeta_{i+1} \rightarrow \dots \triangleright \zeta_m \end{array} \begin{array}{l} \triangleright \\ \triangleright \end{array} s
 \end{array}$$

then the sets $\{\rho_{i+1}, \dots, \rho_n\}$ and $\{\zeta_{i+1}, \dots, \zeta_m\}$ are complete.

Well-bracketing is stable under composition [CCW15], so that we can define a bicategory $\mathbf{PG}_1^{\text{bsi}}$ having arenas (with questions and answers) as objects and well-bracketed and sequential innocent strategies as morphisms. Copycat, and therefore all structural strategies, are well-bracketed, so that $\mathbf{PG}_1^{\text{bsi}}$ is cartesian closed.

All strategies are probabilistic, but most PPCF primitives (all but **coin**) have a deterministic interpretation; when this is the case we omit the valuation from the graphical representation.

5.3 Semantics of PPCF

We now give the semantics of PPCF terms as morphisms in $\mathbf{PG}_!^{\text{bsi}}$, following the standard methodology, whereby a type A is assigned an arena $\llbracket A \rrbracket$ as defined above, and an open term $\Gamma \vdash M : A$ a strategy $\llbracket M \rrbracket^\Gamma : !\llbracket \Gamma \rrbracket \multimap \llbracket A \rrbracket$. The latter uses the usual semantics of contexts: $\llbracket x_1 : A_1, \dots, x_n : A_n \rrbracket = \&_{i=1}^n \llbracket A_i \rrbracket$. (This is the terminal object if $n = 0$.)

We give the interpretation of PPCF terms as well-bracketed, sequential innocent strategies.

Constants. For $b \in \mathbb{B}$ and $n \in \mathbb{N}$, define strategies $\llbracket b \rrbracket : \mathbf{1} \multimap \llbracket \mathbf{Bool} \rrbracket$ and $\llbracket n \rrbracket : \mathbf{1} \multimap \llbracket \mathbf{Nat} \rrbracket$ as

$$\begin{array}{ccc} \llbracket \mathbf{Bool} \rrbracket & & \llbracket \mathbf{Nat} \rrbracket \\ \text{q}^- & & \text{q}^- \\ \downarrow & & \downarrow \\ \text{b}^+ & & \text{n}^+ \end{array}$$

For a nonempty Γ , $\llbracket b \rrbracket^\Gamma$ and $\llbracket n \rrbracket^\Gamma$ are got by precomposing $\llbracket b \rrbracket$ and $\llbracket n \rrbracket$ with the unique strategy $\llbracket \Gamma \rrbracket \multimap \mathbf{1}$.

Probabilistic choice. The strategy $\llbracket \text{coin} \rrbracket$ is the following:

$$\begin{array}{c} \llbracket \mathbf{Bool} \rrbracket \\ \text{q}^- \\ \swarrow \quad \searrow \\ \frac{1}{2} \text{tt}^+ \quad \sim \quad \text{ff}^+ \quad \frac{1}{2} \end{array}$$

with $\llbracket \text{coin} \rrbracket^\Gamma : \llbracket \Gamma \rrbracket \multimap \llbracket \mathbf{Bool} \rrbracket$ defined as for constants.

λ -Calculus constructions. The interpretation of variables, applications, and abstractions uses the cartesian closed structure in a standard way:

$$\begin{aligned} \llbracket x \rrbracket^\Gamma &= \varpi_x, \text{ the } x^{\text{th}} \text{ projection} \\ \llbracket \lambda x.M \rrbracket^\Gamma &= \text{Cur}(\llbracket M \rrbracket^{\Gamma, x:A}) \\ \llbracket MN \rrbracket^\Gamma &= \text{Ev} \odot! \langle \llbracket M \rrbracket^\Gamma, \llbracket N \rrbracket^\Gamma \rangle \end{aligned}$$


Conditionals. For every ground type $\mathbb{X} \in \{\mathbf{Bool}, \mathbf{Nat}\}$ we define a strategy $\text{if}_{\mathbb{X}} : !(\llbracket \mathbf{Bool} \rrbracket \& \llbracket \mathbb{X} \rrbracket \& \llbracket \mathbb{X} \rrbracket) \multimap \llbracket \mathbb{X} \rrbracket$, and whenever $\Gamma \vdash \text{if } M \ N_1 \ N_2 : \mathbb{X}$ we set

$$\llbracket \text{if } M \ N_1 \ N_2 \rrbracket^\Gamma = \text{if}_{\mathbb{X}} \odot! \langle \llbracket M \rrbracket^\Gamma, \llbracket N_1 \rrbracket^\Gamma, \llbracket N_2 \rrbracket^\Gamma \rangle.$$

The strategy $\text{if}_{\mathbf{Nat}}$ is displayed below. Note that to make the representation more convenient we draw it as a strategy $!\llbracket \mathbf{Bool} \rrbracket \otimes !\llbracket \mathbf{Nat} \rrbracket \otimes !\llbracket \mathbf{Nat} \rrbracket \multimap \llbracket \mathbf{Nat} \rrbracket$ which is not a morphism of the Kleisli bicategory $\mathbf{PG}_!$. But the two games $!\llbracket \mathbf{Bool} \rrbracket \otimes$


$![[\mathbf{Nat}]] \otimes ![[\mathbf{Nat}]]$ and $!([[\mathbf{Bool}]] \& [[\mathbf{Nat}]] \& [[\mathbf{Nat}]])$ are equivalent in \mathbf{PG} , so this suffices.

$$![[\mathbf{Bool}]] \otimes ![[\mathbf{Nat}]] \otimes ![[\mathbf{Nat}]] \rightarrow [[\mathbf{Nat}]]$$



Operations on natural numbers. The interpretation of **pred** is given by:

$$![[\mathbf{Nat}]] \rightarrow [[\mathbf{Nat}]]$$



and we do not give the diagrams for **succ** and **iszero**, which should be easy to recover.

Fixpoints. Finally, for any M , the term $Y M$ is interpreted as the colimit of the ω -chain

$$Y_0 [[M]]^\Gamma \Rightarrow Y_1 [[M]]^\Gamma \Rightarrow \dots$$

where $Y_0 [[M]]^\Gamma = \perp$ and $Y_{n+1} [[M]]^\Gamma = \text{Ev} \odot \langle \sigma, Y_n [[M]]^\Gamma \rangle$. The colimit exists, since there is a strong embedding $\perp \Rightarrow Y_1 [[M]]^\Gamma$ and strong embeddings are preserved by horizontal composition and pairing.

5.4 Definability

We continue with a definability result for *finite* sequential innocent strategies. We define what finite means here:

Definition 5.4. Say a sequential innocent strategy $\sigma : \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$ is **finite** when:

- for every negative $s \in \mathcal{S}$, the set $\{s' \in \mathcal{S} \mid s \rightarrow s'\}$ is finite;
- for every positive question $s \in \mathcal{S}$, all but finitely many answers in the set $\{s' \in \mathcal{S} \mid s \rightarrow s'\}$ are maximal.

- there is a bound to the length of gccs in S ;
- For every $x \in \mathcal{C}(S)$, $v(x) \in [0, 1] \cap \mathbb{Q}$;

As in the game model for PCF [HO00], the definability result does not hold for PPCF but for PPCF extended with a family of definition-by-cases primitives

$$\mathbf{case}_k M [N_1 \mid \cdots \mid N_k]$$

with operational semantics given by

$$\frac{\frac{M \xrightarrow{p} M'}{\mathbf{case}_k M [N_1 \mid \cdots \mid N_k] \xrightarrow{p} \mathbf{case}_k M' [N_1 \mid \cdots \mid N_k]}_{i \leq k}}{\mathbf{case}_k i [N_1 \mid \cdots \mid N_k] \xrightarrow{1} N_i}$$

and typing rule

$$\frac{\Gamma \vdash M : \mathbf{Nat} \quad \Gamma \vdash N_i : \mathbb{X}}{\Gamma \vdash \mathbf{case}_k M [N_1 \mid \cdots \mid N_k] : \mathbb{X}}$$

where \mathbb{X} can be either \mathbf{Nat} or \mathbf{Bool} . We define

$$\llbracket \mathbf{case}_k M [N_1 \mid \cdots \mid N_k] \rrbracket^\Gamma = \mathbf{case}_k^{\mathbb{X}} \odot! \langle \llbracket M \rrbracket^\Gamma, \llbracket N_1 \rrbracket, \dots, \llbracket N_k \rrbracket \rangle$$

where

$$\mathbf{case}_k^{\mathbb{X}} : !(\llbracket \mathbf{Nat} \rrbracket \& \bigotimes_{i=1}^k \llbracket \mathbb{X} \rrbracket) \rightarrow \llbracket \mathbb{X} \rrbracket$$

is a strategy which inspects its first argument and, upon return of a value i , continues as the $(i + 1)^{\text{st}}$ argument if $i \leq k$, and stops otherwise.

Now, for $p \in (0, 1) \cap \mathbb{Q}$, the binary choice operator $M +_p N$ is definable from \mathbf{coin} . From this we can easily encode, for every finite convex sum of rationals $\sum_{i=1}^n p_i$, an operator of natural number type, returning each i with probability p_i . Combining this with the definition-by-cases construct, we can define finite rational convex sums of terms: $\sum_{i \in I} p_i \cdot M_i$, in such a way that

$$\left[\left[\sum_{i \in I} p_i \cdot M_i \right] \right]^\Gamma = \sum_{i \in I} p_i \cdot \llbracket M_i \rrbracket^\Gamma.$$

From this we derive:

Theorem 5.5 (Finite definability). *Let A be a PPCF type, and let $\sigma : \mathcal{S} \rightarrow \llbracket A \rrbracket$ be a finite, innocent sequential strategy such that v_S is nonvanishing. Then there is a (PPCF + \mathbf{case}) term M such that $\vdash M : A$ and $\llbracket M \rrbracket \cong \sigma$.*

Proof. The proof is by induction on the bound on the gccs of σ . In the base case (in which all gccs have length 1), σ must be trivial, and therefore we take M to be any diverging term.

In the general case, because σ is sequential innocent and finite, there are k Player moves s_1, \dots, s_k , pairwise inconsistent, immediately following the initial Opponent move. By the axioms on valuations we can write σ as a convex sum $\sum_{i=1}^k p_i \sigma_k$, where $p_i = v_S(\llbracket s_i \rrbracket)$ for each i , and each σ_k is a *value strategy* (in the sense of Definition 4.25, adapted to the arena $\llbracket A \rrbracket$). By finiteness of σ , each p_i is rational, so that by the remarks immediately preceding the theorem it suffices to give for each i a term M_i such that $\llbracket M_i \rrbracket \cong \sigma_i$.

So we show the definability result holds for any finite (sequential innocent) value strategy $\tau : \mathcal{T} \rightarrow \llbracket A \rrbracket$. The proof is exactly that of [HO00]. Write $A = A_1 \rightarrow \dots \rightarrow A_n \rightarrow \mathbf{Nat}$. The case where the return type is **Bool** treated similarly. Consider Player's first move t after Opponent has started the game (this is unique and has probability 1, since τ is a value strategy). If t answers the initial question directly with an integer l , then $M = \lambda \vec{f}.l$ will do.

If instead t is an initial question in one of the A_i , then we write $A_i = C_1 \rightarrow \dots \rightarrow C_m \rightarrow \mathbf{Nat}$. By receptivity of τ , the move t is followed by each Opponent move available at this stage of the game, namely: the return values in **Nat**, and ω symmetric copies of the initial question in each $\llbracket C_j \rrbracket$, which we write q^{C_j} . The copies are redundant, since by the axioms of symmetry each copy has the same future. So for each $j \leq m$, write $C_j = D_{j,1} \rightarrow \dots \rightarrow D_{j,k_j} \rightarrow \mathbf{Nat}$ and consider the branch starting at q^{C_j} as a strategy τ_j on

$$\llbracket D_{j,1} \rightarrow \dots \rightarrow D_{j,k_j} \rightarrow A_1 \rightarrow \dots \rightarrow A_n \rightarrow \mathbf{Nat} \rrbracket.$$

This has bound on gccc strictly smaller than that of σ , so by the induction hypothesis there is a term M_j such that $\llbracket M_j \rrbracket \cong \tau_j$. We abstract away the variables $D_{j,i}$ to get a term $\lambda \vec{y}_j.M_j$ of type $A_1 \rightarrow \dots \rightarrow A_n \rightarrow \mathbf{Nat}$.

Going back to the move t , we consider the possible answers to it in **Nat**: let k be the greatest one with a successor in the strategy (this exists by the second condition on finite strategies). For each $l \leq k$, isolating the “ l -branch”, we obtain as above a strategy on $A_1 \rightarrow \dots \rightarrow A_n \rightarrow \mathbf{Nat}$, and thus, a term N_l of type A . We define

$$M = \lambda \vec{f}.\mathbf{case}_k (f_i(\lambda \vec{y}_1.M_1) \dots (\lambda \vec{y}_m.M_m)) [N_1 \mid \dots \mid N_k],$$

and leave out the verification that $\llbracket M \rrbracket \cong \tau$. □

We conclude with a few remarks.

- From here the path to full abstraction is relatively standard, although the addition of the definition-by-cases construct is more problematic than in [HO00, DH02], since it is not clear that the term $\mathbf{case}_k M [N_1 \mid \dots \mid N_k]$ is definable up to observational equivalence in PPCF. This does not affect the full abstraction result, since it follows from the main theorem in [ETP14] that PPCF + **case** terms have no more distinguishing power than standard PPCF terms.
- In fact, full abstraction can be obtained from the results of [ETP14] even without a finite definability result. This was carried out in collaboration with Simon Castellan, Pierre Clairambault, and Glynn Winskel [CCPW18].

- We have seen that it is possible to give a “sequential innocence” condition for probabilistic strategies in concurrent games. This is known to be difficult with Hyland-Ong games, so the authors of [DH02] instead construct a model for a probabilistic version of *Idealised Algol*, which does not require innocence. Although Idealised Algol, which includes first-order references, can be modelled with concurrent games [CCW19], its probabilistic extension is problematic: the interaction of concurrency and state already involves nondeterministic behaviour due to scheduling issues outside of Player’s control. It is notoriously difficult to mix nondeterminism and probability, so this is a limitation of the concurrent games model presented here. Recent work by Marc de Visme on event structures with *mixed choice* [dV19] seems like a promising solution, which remains to be fully investigated in the context of games.

Chapter 6

Measurable concurrent games

The preceding chapters were concerned with the addition of *discrete* probability to concurrent games, and we have seen that to make a strategy $\sigma : \mathcal{S} \rightarrow \mathcal{A}$ probabilistic it suffices to equip it with a function $v : \mathcal{C}(S) \rightarrow [0, 1]$ assigning a coefficient to each configuration.

We now aim for a generalised model, supporting probabilistic programs with continuous datatypes (*e.g.* with a type of real numbers). It is well-known that the “naive” approach to probability, in which coefficients are assigned directly to individual elements of a set, is not satisfying in the situation where this set is uncountable. To remedy this problem is the purpose of the *measure-theoretic* approach to probability, where coefficients are instead assigned to certain subsets of elements. This leads to the following notion. (A standard textbook on measure and probability theory is [Bil08].)

Definition 6.1. A **measurable space** is a set X equipped with a σ -**algebra**, that is, a set Σ_X of subsets of X containing X itself, and closed under complements and countable unions. (This implies Σ_X is also closed under countable intersections.) The elements of Σ_X are called **measurable subsets** of X .

If we are to generalise probabilistic concurrent strategies to a continuous setting, the set of configurations $\mathcal{C}(S)$ must therefore be turned into a measurable space. It is tempting to consider a fully generalised model in which any σ -algebra on $\mathcal{C}(S)$ gives a valid strategy. But this approach quickly proves too abstract in the context of game semantics, where in order to get notions of composition and identity, one requires well-understood connections between the measurable space structure, on one hand, and causality, consistency, polarity, and symmetry, on the other.

Accordingly, we are led to a more involved notion of *measurable event structure with symmetry*. This is the topic of Section 6.1. We will see that this can be used to construct an appropriate bicategory of “measurable games and strategies”, using the same method as in the first part of this thesis. Sections 6.2 to 6.4 are devoted to the development of this model. We defer to the next chapter the addition of probability.




Figure 6.1: A fibred event structure $f : E \rightarrow E'$.

6.1 Measurable event structures

The key idea is best conveyed using *fibrations* of event structures, so this will be our starting point. The presentation we give in 6.1.1 agrees with that given in [PW18], where measurable event structures were introduced.

In 6.1.2, we will investigate an alternative, more general notion of measurable event structures. Roughly speaking, to do this we move from discrete fibrations to *presheaves* via the standard correspondence – but there are additional subtleties. We compare the two notions in the spirit of existing work connecting event structures and presheaves [Win99, SW10].

This new notion of measurable event structure is more easily enriched with *symmetry*, which we add in 6.1.3. We build a symmetric monoidal category **Mess** (of “measurable event structures with symmetry”), in which the usual category **Ess** of event structures with symmetry embeds fully and faithfully.

6.1.1 Fibrations of event structures

Consider a process outputting two real numbers r_1 and r_2 consecutively, each chosen non-deterministically in \mathbb{R} . An event structure representation of it is pictured as E on the left of the arrow in Figure 6.1. Each ‘real line’ represents an uncountable set of events, all pairwise in immediate conflict. Only a portion of the event structure is displayed — there are in fact uncountably many such “ r_2 ” real lines, one for each $r_1 \in \mathbb{R}$.

Configurations of E can have one of three forms: \emptyset , $\{r_1\}$, or $\{r_1, r_2\}$ where $r_1, r_2 \in E$ and $r_1 \rightarrow r_2$. Our approach involves projecting them to the configurations of a *base event structure* E' , displayed on the right of the figure. The goal is to *encapsulate* the uncountable non-deterministic branching in E' in *fibres* over the configurations of E' : \emptyset , $\{a_1\}$ and $\{a_1, a_2\}$.

Observe that the projection map $E \rightarrow E'$ is rigid (*i.e.* preserves causal dependency, *cf.* Definition 3.23). Rigid maps are appropriate in this context because they provide a well-behaved notion of *fibre*:

Lemma 6.2. *If $f : E \rightarrow E'$ is a map of event structures, then f is rigid if and only if the induced map $\mathcal{C}(E) \rightarrow \mathcal{C}(E')$ is a discrete fibration of partial orders, i.e. for every $x \in \mathcal{C}(E)$, if $y \subseteq fx$ for some $y \in \mathcal{C}(E')$, then there exists a unique $x' \in \mathcal{C}(E)$ such that $x' \subseteq x$ and $fx' = y$.*

Proof. (Only if). Suppose that $f : E \rightarrow E'$ is rigid and that we have $x \in \mathcal{C}(E)$ and $y \in \mathcal{C}(E')$ such that $y \subseteq fx$. The restriction of f to x is injective by assumption, and $(f|_x)^{-1}y$ is necessarily the only $x' \subseteq x$ such that $fx' = y$. It is a configuration,

since it is consistent (as a subset of x) and down-closed (f preserves and reflects causal dependency, and y is down-closed).

(If). Suppose now that f is not rigid, so there are $e, e' \in E$ such that $e \rightarrow e'$ but $f(e) \not\rightarrow f(e')$. Consider $x = [e']$ and $y = [f(e')]$. Then $y \subseteq fx$, but since $f(e') \in y$ and $f(e) \notin y$, there can be no $x' \in \mathcal{C}(E)$ such that $x' \subseteq x$ and $fx' = y$ (such an x' would not be down-closed). \square

Accordingly, given a rigid map $f : E \rightarrow E'$ and a configuration $x \in \mathcal{C}(E')$, the **fibre over** x is the preimage $f^{-1}\{x\} = \{z \in \mathcal{C}(E) \mid fz = x\}$. If $x \subseteq y \in \mathcal{C}(E')$, we write $r_{x,y} : f^{-1}\{y\} \rightarrow f^{-1}\{x\}$ for the **restriction map** determined by Lemma 6.2; $r_{x,y}$ sends $z \in f^{-1}\{y\}$ to the unique $w \subseteq z$ such that $fw = x$.

Once the configurations of E are organised as fibres, we turn them individually into measurable spaces. (The measurable space structure on $\mathcal{C}(E)$ can then be obtained via a coproduct construction, but it will not play any useful role in the development.) We then require that the restriction maps are *measurable functions* in the standard sense [Bil08]:

Definition 6.3. A **measurable function** from (X, Σ_X) to (Y, Σ_Y) is a function $m : X \rightarrow Y$ such that every $U \in \Sigma_Y$ has a measurable preimage: $m^{-1}U \in \Sigma_X$.

The *measurable event structures* of [PW18] are defined as follows. To avoid confusion with the alternative presentation we use in this thesis (to be introduced below), we call them *measurable fibrations* of event structures.

Definition 6.4. A **measurable fibration** of event structures consists of a rigid map $f : E \rightarrow E'$, and for each $x \in \mathcal{C}(E')$, a σ -algebra Σ_x on the fibre over x , such that for every $x \subseteq y \in \mathcal{C}(E)$, the restriction map $r_{x,y}$ is measurable.

6.1.2 Event structures and presheaves

For any category \mathbb{D} , there is a well-known categorical equivalence between discrete fibrations $\mathbb{C} \rightarrow \mathbb{D}$ and contravariant functors $\mathbb{D}^{\text{op}} \rightarrow \mathbf{Set}$ to the category of sets, better known as **presheaves** over \mathbb{D} .

With this in mind, the following is not surprising:

Lemma 6.5. *For every measurable fibration $f : E' \rightarrow E$, the following defines a functor $\mathcal{M}_f : \mathcal{C}(E)^{\text{op}} \rightarrow \mathbf{Meas}$, where $\mathcal{C}(E)$ is a partial order seen as a category and \mathbf{Meas} is the category of measurable spaces and measurable functions:*

$$\begin{aligned} \mathcal{M}_f : x &\longmapsto (f^{-1}\{x\}, \Sigma_x) \\ (x \subseteq y) &\longmapsto r_{x,y} \end{aligned}$$

Proof. Direct verification. \square

The category **Meas** inherits many properties from **Set**; in particular it is complete and cocomplete. (But note that it is not cartesian closed [A⁺61].) Thus we refer to functors $\mathbb{C}^{\text{op}} \rightarrow \mathbf{Meas}$ as **measurable presheaves**, which in this thesis we often simply call presheaves; there should be no ambiguity.

The converse to Lemma 6.5 is not true, as many presheaves on E are not *representable* in this way. The representable ones can be characterised via a “separatedness” condition:

Definition 6.6. A presheaf $\mathcal{M} : \mathcal{C}(E)^{\text{op}} \rightarrow \mathbf{Meas}$ is **separated** if it satisfies the following condition: for every $x \in \mathcal{C}(E)$ and $u, v \in \mathcal{M}(x)$, if $\mathcal{M}([e] \subseteq x)(u) = \mathcal{M}([e] \subseteq x)(v)$ for every $e \in x$, then $u = v$.

(Note that \mathcal{M} is separated if and only if, writing $U : \mathbf{Meas} \rightarrow \mathbf{Set}$ for the forgetful functor, the presheaf $U \circ \mathcal{M} : \mathcal{C}(E)^{\text{op}} \rightarrow \mathbf{Set}$ is *separated*, in the topos-theoretic sense [MM12], with respect to the Grothendieck topology on $\mathcal{C}(E)$ generated by covering families $\{[e] \subseteq x \mid e \in x\}$ for each x . The thesis does not make use of it.)

We give a characterisation of those measurable presheaves arising from measurable fibrations. The result is a straightforward adaptation of the representation theorem in [Win99]. By a **nonempty presheaf** we mean one for which there exists $x \in \mathcal{C}(E)^{\text{op}}$ with $\mathcal{M}(x) \neq \emptyset$. Observe that a nonempty, separated presheaf is necessarily **rooted**, in the sense that $\mathcal{M}(\emptyset)$ is a singleton space. (All representable presheaves must be rooted, since for any $f : E \rightarrow E'$, $f^{-1}\{\emptyset\} = \{\emptyset\}$.)

Lemma 6.7 (Adapted from [Win99]). *A presheaf $\mathcal{M} : \mathcal{C}(E)^{\text{op}} \rightarrow \mathbf{Meas}$ is nonempty and separated if and only if there is a measurable fibration $f : E' \rightarrow E$ such that $\mathcal{M} \cong \mathcal{M}_f$.*

Although representable presheaves have a more intuitive operational behaviour in terms of event structures, it is interesting to keep the extra generality, since the model supports it. Thus we define:

Definition 6.8. A **measurable event structure** consists of an event structure E and a presheaf $\mathcal{M} : \mathcal{C}(E)^{\text{op}} \rightarrow \mathbf{Meas}$. Say (E, \mathcal{M}) is **representable** if \mathcal{M} is nonempty and separated in the sense of Definition 6.6.

Before introducing maps of measurable event structures, we introduce symmetry.

6.1.3 Symmetry in measurable event structures

Fix $\mathcal{E} = (E, \cong_E)$ an event structure with symmetry. To make \mathcal{E} measurable we take care to ensure that symmetric configurations of E have isomorphic fibres. A natural solution is to make \mathcal{M} functorial with respect to the bijections in the isomorphism family \cong_E . We begin by adding them to the category $\mathcal{C}(E)$:

Definition 6.9. The **category of configurations** $\mathcal{C}(\mathcal{E})$ of an ess \mathcal{E} is the subcategory of \mathbf{Set} with objects the elements of $\mathcal{C}(E)$, and morphisms generated by inclusions maps $x \hookrightarrow y$ and symmetries $\theta : x \cong_E y$.

It is worth remarking that any morphism $x \rightarrow y$ in $\mathcal{C}(\mathcal{E})$ can be factored uniquely as $x \cong_E x' \hookrightarrow y$, by the axioms of symmetry.

Definition 6.10. A **measurable event structure with symmetry** is $(\mathcal{E}, \mathcal{M})$ where \mathcal{E} is an ess and $\mathcal{M} : \mathcal{C}(\mathcal{E})^{\text{op}} \rightarrow \mathbf{Meas}$ is a functor. Say $(\mathcal{E}, \mathcal{M})$ is **representable** when $(E, \mathcal{M} \circ \iota)$ is, writing $\iota : \mathcal{C}(E)^{\text{op}} \hookrightarrow \mathcal{C}(\mathcal{E})^{\text{op}}$ for the canonical inclusion functor.

In the presence of symmetry, the issue of representability deserves further discussion. Indeed the correspondence between separated presheaves and rigid maps of event structures does not extend to this setting, *i.e.* there are rigid maps $\mathcal{E}' \rightarrow \mathcal{E}$ for which the induced functor $\mathcal{C}(\mathcal{E}') \rightarrow \mathcal{C}(\mathcal{E})$ is not a discrete fibration (and so does not give a presheaf).

Nonetheless, if $\mathcal{M} : \mathcal{C}(\mathcal{E})^{\text{op}} \rightarrow \mathbf{Meas}$ is nonempty and separated, and $f : E' \rightarrow E$ is the measurable fibration obtained (via Lemma 6.7) from the underlying measurable event structure, then E' can be equipped with an isomorphism family, so that $f : \mathcal{E}' \rightarrow \mathcal{E}$ gives a discrete fibration $\mathcal{C}(\mathcal{E}') \rightarrow \mathcal{C}(\mathcal{E})$ with corresponding presheaf isomorphic to \mathcal{M} . The isomorphism family $\cong_{E'}$ is set to contain bijections $\theta : x \cong y$ meeting the following conditions:

- $f\theta \in \cong_E$ and θ is the bijection $x \cong fx \stackrel{f\theta}{\cong} fy \cong y$;
- The map $\mathcal{M}(f\theta) : \mathcal{M}(fy) \rightarrow \mathcal{M}(fx)$ satisfies $\mathcal{M}(f\theta)(y) = x$.

This construction is significant, because the world of measurable fibrations (with symmetry) is generally more intuitive. (There are other ways to relate event structures with symmetry and presheaf models, see [SW10].)

We proceed to define a category \mathbf{Mess} of measurable event structures with symmetry. We begin by observing that because maps of event structures with symmetry preserve both inclusion and symmetries, a map $f : \mathcal{E} \rightarrow \mathcal{E}'$ induces a functor $\mathcal{C}(\mathcal{E}) \rightarrow \mathcal{C}(\mathcal{E}')$. We also use f to denote this functor, and f^{op} for the corresponding functor $\mathcal{C}(\mathcal{E})^{\text{op}} \rightarrow \mathcal{C}(\mathcal{E}')^{\text{op}}$. Then we define:

Definition 6.11. Let $(\mathcal{E}, \mathcal{M})$ and $(\mathcal{E}', \mathcal{M}')$ be measurable event structures with symmetry. A **map** $(\mathcal{E}, \mathcal{M}) \rightarrow (\mathcal{E}', \mathcal{M}')$ is a map $f : \mathcal{E} \rightarrow \mathcal{E}'$ of ess, together with a natural transformation $\text{nt}^f : \mathcal{M} \rightarrow \mathcal{M}' \circ f^{\text{op}}$ between functors $\mathcal{C}(\mathcal{E})^{\text{op}} \rightarrow \mathbf{Meas}$.

We can compose maps: for

$$(\mathcal{E}_1, \mathcal{M}_1) \xrightarrow{(f, \text{nt}^f)} (\mathcal{E}_2, \mathcal{M}_2) \xrightarrow{(g, \text{nt}^g)} (\mathcal{E}_3, \mathcal{M}_3)$$

the composition $(\mathcal{E}_1, \mathcal{M}_1) \xrightarrow{(g, \text{nt}^g) \circ (f, \text{nt}^f)} (\mathcal{E}_3, \mathcal{M}_3)$ consists of the map $g \circ f$ and the natural transformation $\text{nt}^{g \circ f}$ defined as $(\text{id}_{f^{\text{op}}} \star \text{nt}^g) \circ \text{nt}^f$, as in the following pasting diagram:

$$\begin{array}{ccc}
 & \mathcal{M}_1 & \\
 & \curvearrowright & \\
 \mathcal{C}(\mathcal{E}_1)^{\text{op}} & & \mathbf{Meas} \\
 & \text{nt}^f \Downarrow & \\
 & \mathcal{M}_2 & \\
 & \curvearrowright & \\
 \mathcal{C}(\mathcal{E}_2)^{\text{op}} & & \mathbf{Meas} \\
 & \text{nt}^g \Downarrow & \\
 & \mathcal{M}_3 & \\
 & \curvearrowright & \\
 \mathcal{C}(\mathcal{E}_3)^{\text{op}} & & \mathbf{Meas}
 \end{array}$$

There is an obvious notion of identity $(\mathcal{E}, \mathcal{M}) \rightarrow (\mathcal{E}, \mathcal{M})$, and associativity and unit laws pose no problem:

Lemma 6.12. *Measurable event structures with symmetry and maps between them form a category, called **Mess**.*

We note the following basic fact about isomorphisms in **Mess**:

Lemma 6.13. *If $(\mathcal{E}, \mathcal{M})$ and $(\mathcal{E}', \mathcal{M}')$ are objects of **Mess**, $f : \mathcal{E} \rightarrow \mathcal{E}'$ is an isomorphism of event structures with symmetry, and $\text{nt}^f : \mathcal{M} \xrightarrow{\bullet} \mathcal{M}' \circ f^{\text{op}}$ is invertible, then $(f, \text{nt}^f) : (\mathcal{E}, \mathcal{M}) \rightarrow (\mathcal{E}', \mathcal{M}')$ is an isomorphism.*

Proof. Its inverse is given by $(f^{-1}, \text{nt}^{f^{-1}})$ where

$$\text{nt}^{f^{-1}} : \mathcal{M}' = (\mathcal{M}' \circ f^{\text{op}}) \circ (f^{\text{op}})^{-1} \xrightarrow{(\text{nt}^f)^{-1} \circ (f^{\text{op}})^{-1}} \mathcal{M} \circ (f^{\text{op}})^{-1}.$$

The verification is straightforward. \square

Objects of **Mess** support a notion of parallel composition which extends¹ that of **Ess** using the cartesian product in **Meas**, and the fact that $\mathcal{C}(\mathcal{E} \parallel \mathcal{E}') \cong \mathcal{C}(\mathcal{E}) \times \mathcal{C}(\mathcal{E}')$.

When (X, Σ_X) and (Y, Σ_Y) are measurable spaces the **product space** $(X \times Y, \Sigma_{X \times Y})$ is the product of X and Y as sets, with σ -algebra the smallest one containing $\{U \times V \mid U \in \Sigma_X, V \in \Sigma_Y\}$. This is a categorical product in **Meas** (in particular there are measurable projections $X \times Y \rightarrow X$ and $X \times Y \rightarrow Y$). We make use of the canonical functor $\times : \mathbf{Meas} \times \mathbf{Meas} \rightarrow \mathbf{Meas}$ in the next definition.

Definition 6.14. The **parallel composition** of $(\mathcal{E}, \mathcal{M})$ and $(\mathcal{E}', \mathcal{M}')$ is $\mathcal{E} \parallel \mathcal{E}'$ equipped with the functor $\mathcal{M} \parallel \mathcal{M}'$ defined to be the composite

$$\mathcal{C}(\mathcal{E} \parallel \mathcal{E}')^{\text{op}} \cong \mathcal{C}(\mathcal{E})^{\text{op}} \times \mathcal{C}(\mathcal{E}')^{\text{op}} \xrightarrow{\mathcal{M} \times \mathcal{M}'} \mathbf{Meas} \times \mathbf{Meas} \xrightarrow{\times} \mathbf{Meas}.$$

Given $(\mathcal{E}_1, \mathcal{M}_1), (\mathcal{E}_2, \mathcal{M}_2), (\mathcal{E}'_1, \mathcal{M}'_1), (\mathcal{E}'_2, \mathcal{M}'_2)$ and maps $(f_i, \text{nt}^{f_i}) : (\mathcal{E}_i, \mathcal{M}_i) \rightarrow (\mathcal{E}'_i, \mathcal{M}'_i)$ for $i = 1, 2$, define $(f_1, \text{nt}^{f_1}) \parallel (f_2, \text{nt}^{f_2}) = (f_1 \parallel f_2, \text{nt}^{f_1 \parallel f_2})$ where the map $\text{nt}_{x_1 \parallel x_2}^{f_1 \parallel f_2} : \mathcal{M}_1(x_1) \times \mathcal{M}_2(x_2) \rightarrow \mathcal{M}'_1(f_1 x_1) \times \mathcal{M}'_2(f_2 x_2)$ is given by $\text{nt}_{x_1}^{f_1} \times \text{nt}_{x_2}^{f_2}$.

The empty ess \emptyset is assigned the terminal presheaf $\mathcal{M}_{\emptyset} : \mathcal{C}(\emptyset)^{\text{op}} \rightarrow \mathbf{Meas}$ mapping the empty configuration \emptyset to the singleton space $\{*\}$. (Note that this is the only nonempty, separated presheaf on $\mathcal{C}(\emptyset)$.) It is easy then to check that **Mess** is symmetric monoidal:

Lemma 6.15. *(**Mess**, \parallel, \emptyset) is a symmetric monoidal category, where for objects $(\mathcal{A}, \mathcal{M}_{\mathcal{A}}), (\mathcal{B}, \mathcal{M}_{\mathcal{B}}), (\mathcal{C}, \mathcal{M}_{\mathcal{C}})$, the structural isomorphisms $\underline{a}_{\mathcal{A}, \mathcal{B}, \mathcal{C}}, \underline{r}_{\mathcal{A}}, \underline{l}_{\mathcal{A}}$ and $\underline{b}_{\mathcal{A}, \mathcal{B}}$ associated with the smc **Ess** (see Lemma 2.40) are equipped with natural transformations*

$$\begin{aligned} \text{nt}^{\underline{a}_{\mathcal{A}, \mathcal{B}, \mathcal{C}}} : (\mathcal{M}_{\mathcal{A}} \parallel \mathcal{M}_{\mathcal{B}}) \parallel \mathcal{M}_{\mathcal{C}} &\xrightarrow{\bullet} (\mathcal{M}_{\mathcal{A}} \parallel (\mathcal{M}_{\mathcal{B}} \parallel \mathcal{M}_{\mathcal{C}})) \circ \underline{a}_{\mathcal{A}, \mathcal{B}, \mathcal{C}}^{\text{op}} \\ \text{nt}^{\underline{l}_{\mathcal{A}}} : \mathcal{M}_{\emptyset} \parallel \mathcal{M}_{\mathcal{A}} &\xrightarrow{\bullet} \mathcal{M}_{\mathcal{A}} \circ \underline{l}_{\mathcal{A}}^{\text{op}} \\ \text{nt}^{\underline{r}_{\mathcal{A}}} : \mathcal{M}_{\mathcal{A}} \parallel \mathcal{M}_{\emptyset} &\xrightarrow{\bullet} \mathcal{M}_{\mathcal{A}} \circ \underline{r}_{\mathcal{A}}^{\text{op}} \\ \text{nt}^{\underline{b}_{\mathcal{A}, \mathcal{B}}} : \mathcal{M}_{\mathcal{A}} \parallel \mathcal{M}_{\mathcal{B}} &\xrightarrow{\bullet} \mathcal{M}_{\mathcal{B}} \parallel \mathcal{M}_{\mathcal{A}} \circ \underline{b}_{\mathcal{A}, \mathcal{B}}^{\text{op}} \end{aligned}$$

¹This extension is canonical, and it might be informative to deduce this from some general theory, recognising that **Mess** is a *lax comma object* of the form $(\mathcal{C}(-)^{\text{op}} \downarrow \mathbf{Meas})$ in the 2-category **Cat**. Then, because $\mathcal{C}(-)^{\text{op}} : \mathbf{Ess} \rightarrow \mathbf{Cat}$ is a symmetric monoidal functor and **Meas** is a (cartesian) monoid in **Cat**, **Mess** inherits a canonical symmetric monoidal structure from **Ess**. The presentation we give is more concrete.

all of which are made up of canonical morphisms using the cartesian structure of **Meas**.

The category **Mess** will form the basis of the model of measurable games and strategies developed in the next section. Before proceeding with the development, we briefly remark that **Ess** occurs as a subcategory of **Mess**; later we will use this to show that the games model in the first part of this thesis is a special case of the forthcoming “measurable games” model.

6.1.4 Embedding **Ess** into **Mess**

Every **ess** \mathcal{E} can be turned into a measurable **ess** with trivial structure: consider the presheaf $\mathbf{1}_{\mathcal{E}} : \mathcal{C}(\mathcal{E}) \rightarrow \mathbf{Meas}$, defined by $\mathbf{1}_{\mathcal{E}}(x) = \{*\}$ for every x . Then the pair $(\mathcal{E}, \mathbf{1}_{\mathcal{E}})$ is representable, and has corresponding fibration the identity function $\mathcal{E} \rightarrow \mathcal{E}$. A measurable event structure of this form is called **discrete**, because the induced σ -algebra on $\mathcal{C}(E)$ is *discrete* [Bil08], *i.e.* every subset is measurable.

Accordingly, we write $\mathbf{disc}(\mathcal{E}) = (\mathcal{E}, \mathbf{1}_{\mathcal{E}})$, and it is easy to check (using that $\mathbf{1}$ is terminal in the presheaf category) that this defines an embedding $\mathbf{disc} : \mathbf{Ess} \rightarrow \mathbf{Mess}$.

6.2 Measurable games and strategies

We finally get to the development of a model of games and strategies based on measurable event structures with symmetry. The story relies on the development in Chapter 2, and unfolds in essentially the same order. The first step is the addition of polarity, which is straightforward: a **measurable essp** is a measurable **ess** $(\mathcal{A}, \mathcal{M}_{\mathcal{A}})$ where \mathcal{A} is also an **essp**.

We call **measurable game with symmetry** a tuple $\mathcal{A} = (A, \mathcal{M}_{\mathcal{A}}, \cong_A, \cong_A^-, \cong_A^+)$ such that $(A, \cong_A, \cong_A^-, \cong_A^+)$ is a game with symmetry (*cf.* Definition 2.12) and $((A, \cong_A), \mathcal{M}_{\mathcal{A}})$ is a measurable **essp**. Say \mathcal{A} is a **measurable arena** if the underlying game with symmetry is an arena.

With the notion of game in place, we define strategies.

Definition 6.16. A **measurable strategy** on a measurable game \mathcal{A} consists of a measurable **essp** $(\mathcal{S}, \mathcal{M}_{\mathcal{S}})$, and a map $(\sigma, \mathbf{nt}^{\sigma}) : (\mathcal{S}, \mathcal{M}_{\mathcal{S}}) \rightarrow (\mathcal{A}, \mathcal{M}_{\mathcal{A}})$, such that:

- the underlying map $\sigma : \mathcal{S} \rightarrow \mathcal{A}$ is a strategy;
- $(\sigma, \mathbf{nt}^{\sigma})$ is **measurably receptive**: if $x \subseteq^- y$ in $\mathcal{C}(\mathcal{S})$, then

$$\begin{array}{ccc} \mathcal{M}_{\mathcal{S}}(y) & \xrightarrow{\mathcal{M}_{\mathcal{S}}(x \subseteq y)} & \mathcal{M}_{\mathcal{S}}(x) \\ \mathbf{nt}_y^{\sigma} \downarrow & & \downarrow \mathbf{nt}_x^{\sigma} \\ \mathcal{M}_{\mathcal{A}}(\sigma y) & \xrightarrow{\mathcal{M}_{\mathcal{A}}(\sigma x \subseteq \sigma y)} & \mathcal{M}_{\mathcal{A}}(\sigma x) \end{array}$$

is a pullback in **Meas**.

Receptivity is there to prevent Player from constraining Opponent's behaviour further than is allowed by the game. In the measurable setting, any negative extension of a measurable *fibre* must agree with the corresponding extension in the game. This condition serves to ensure a well-behaved interaction with *measurable copycat*, the identity on a measurable game, which we introduce below. We will shortly give a characterisation of pullbacks in **Meas**.

Measurable strategies compose much like strategies, via a pullback construction followed by a hiding step. We first investigate pullbacks in **Mess**.

6.2.1 Pullbacks in Mess

We begin with a characterisation of pullbacks in **Meas**. Suppose X, Y, Z are measurable spaces and $g : X \rightarrow Y$ and $h : Z \rightarrow Y$ are measurable functions. The pullback

$$\begin{array}{ccc} P & \xrightarrow{\Pi_2} & Z \\ \Pi_1 \downarrow & & \downarrow h \\ X & \xrightarrow{g} & Y \end{array}$$

exists and has underlying set the pullback in **Set**: $P = \{(x, z) \in X \times Z \mid g(x) = h(z)\}$, with Π_1 and Π_2 the usual projections. The associated Σ_P is the subspace σ -algebra induced by $\Sigma_{X \times Z}$, using that $P \subseteq X \times Z$.

We then deduce the following:

Lemma 6.17. *For $\mathcal{E} \in \mathbf{Ess}$, the category $[\mathcal{C}(\mathcal{E})^{\text{op}}, \mathbf{Meas}]$ of presheaves on $\mathcal{C}(\mathcal{E})$ (and natural transformations between them) has all pullbacks.*

Proof. It is standard that limits exist in a functor category whenever they exist in the codomain category. Note that they are computed pointwise. \square

We use this to construct pullbacks in **Mess**. Like **Ess**, **Mess** does not have all pullbacks. However when

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


is a pullback square in **Ess**, and $(\sigma, \text{nt}^\sigma) : (\mathcal{S}, \mathcal{M}_\mathcal{S}) \rightarrow (\mathcal{A}, \mathcal{M}_\mathcal{A})$ and $(\tau, \text{nt}^\tau) : (\mathcal{T}, \mathcal{M}_\mathcal{T}) \rightarrow (\mathcal{A}, \mathcal{M}_\mathcal{A})$, the ess $\mathcal{S} \wedge \mathcal{T}$ can be equipped with the functor $\mathcal{M}_{\mathcal{S} \wedge \mathcal{T}} : \mathcal{C}(\mathcal{S} \wedge \mathcal{T})^{\text{op}} \rightarrow \mathbf{Meas}$ obtained as the pullback

$$\begin{array}{ccc} & \mathcal{M}_{\mathcal{S} \wedge \mathcal{T}} & \\ \swarrow & & \searrow \\ \mathcal{M}_\mathcal{S} \circ \Pi_1^{\text{op}} & & \mathcal{M}_\mathcal{T} \circ \Pi_2^{\text{op}} \\ \text{nt}^\sigma \circ \Pi_1^{\text{op}} \searrow & & \swarrow \text{nt}^\tau \circ \Pi_2^{\text{op}} \\ & \mathcal{M}_\mathcal{A} \circ (\sigma \wedge \tau)^{\text{op}} & \end{array}$$


in $[\mathcal{C}(\mathcal{S} \wedge \mathcal{T})^{\text{op}}, \mathbf{Meas}]$, which always exists by the previous lemma. Call the projections nt^{Π_1} and nt^{Π_2} , respectively.

Lemma 6.18. *The measurable ess $(\mathcal{S} \wedge \mathcal{T}, \mathcal{M}_{\mathcal{S} \wedge \mathcal{T}})$ is the pullback of $(\sigma, \text{nt}^\sigma)$ and (τ, nt^τ) in \mathbf{Mess} , with projections $(\Pi_1, \text{nt}^{\Pi_1})$ and $(\Pi_2, \text{nt}^{\Pi_2})$.*


Proof. Suppose that




commutes in \mathbf{Mess} . Because $\mathcal{S} \wedge \mathcal{T}$ is a pullback of event structures with symmetry, there is a unique map $\omega : \mathcal{Q} \rightarrow \mathcal{S} \wedge \mathcal{T}$ making



commute. For abstract reasons [ML13], pre-composition with the functor $\omega^{\text{op}} : \mathcal{C}(\mathcal{Q})^{\text{op}} \rightarrow \mathcal{C}(\mathcal{S} \wedge \mathcal{T})^{\text{op}}$ preserves limits (and *a fortiori* pullbacks), so that the pullback cone for $\mathcal{M}_{\mathcal{S} \wedge \mathcal{T}}$ in $[\mathcal{C}(\mathcal{S} \wedge \mathcal{T})^{\text{op}}, \mathbf{Meas}]$ can be turned into one in $[\mathcal{C}(\mathcal{Q})^{\text{op}}, \mathbf{Meas}]$:



Finally, because the following commutes (using that $p_1 = \Pi_1 \circ \omega$ and $p_2 = \Pi_2 \circ \omega$)



there is a unique natural transformation $\mathcal{M}_{\mathcal{Q}} \xrightarrow{\cdot} \mathcal{M}_{\mathcal{S} \wedge \mathcal{T}} \circ \omega^{\text{op}}$ satisfying the necessary properties. \square

6.2.2 Composition of measurable strategies

A measurable strategy **from** $(\mathcal{A}, \mathcal{M}_{\mathcal{A}})$ **to** $(\mathcal{B}, \mathcal{M}_{\mathcal{B}})$ is a measurable strategy on $(\mathcal{A}, \mathcal{M}_{\mathcal{A}})^{\perp} \parallel (\mathcal{B}, \mathcal{M}_{\mathcal{B}})$, where the **dual** measurable game $(\mathcal{A}, \mathcal{M}_{\mathcal{A}})^{\perp}$ is $(\mathcal{A}^{\perp}, \mathcal{M}_{\mathcal{A}})$. (This is well-defined: $\mathcal{C}(\mathcal{A}^{\perp}) = \mathcal{C}(\mathcal{A})$.)

Interaction of measurable strategies. At this point it helps to introduce some lighter notation. From now on, we will use $\underline{\mathcal{A}}, \underline{\mathcal{B}}, \underline{\mathcal{S}}, \underline{\mathcal{T}}, \dots$ to denote measurable essps, where the underlying data $\mathcal{A}, \mathcal{B}, \mathcal{M}_{\mathcal{A}}, \mathcal{M}_{\mathcal{B}}, \dots$ is kept implicit. Similarly we write $\underline{\sigma}, \underline{\tau}, \dots$ for maps, omitting the components $\sigma, \tau, \text{nt}^{\sigma}, \text{nt}^{\tau}$, etc.

So to compose $\underline{\sigma} : \underline{\mathcal{S}} \rightarrow \underline{\mathcal{A}}^{\perp} \parallel \underline{\mathcal{B}}$ and $\underline{\tau} : \underline{\mathcal{T}} \rightarrow \underline{\mathcal{B}}^{\perp} \parallel \underline{\mathcal{C}}$, start with the pullback

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

in **Mess**. The measurable essp $\underline{\mathcal{T}} \otimes \underline{\mathcal{S}}$ is the **interaction** of $\underline{\sigma}$ and $\underline{\tau}$.

We note an important property of the interaction:

Lemma 6.19. *Suppose $x \text{--} c^e y \in \mathcal{C}(\underline{\mathcal{T}} \otimes \underline{\mathcal{S}})$, and assume e is a σ -action (i.e. $\Pi_1 e$ is a positive event of $\underline{\mathcal{S}}$). Then,*

$$\begin{array}{ccc}
 \mathcal{M}_{\underline{\mathcal{T}} \otimes \underline{\mathcal{S}}}(y) & \xrightarrow{\mathcal{M}_{\underline{\mathcal{T}} \otimes \underline{\mathcal{S}}}(x \text{--} c^e y)} & \mathcal{M}_{\underline{\mathcal{T}} \otimes \underline{\mathcal{S}}}(x) \\
 \text{nt}_y^{\Pi_1} \downarrow & & \downarrow \text{nt}_x^{\Pi_1} \\
 \mathcal{M}_{\underline{\mathcal{S}} \parallel \underline{\mathcal{C}}}(y_S \parallel y_C) & \longrightarrow & \mathcal{M}_{\underline{\mathcal{S}} \parallel \underline{\mathcal{C}}}(x_S \parallel x_C)
 \end{array}$$

is a pullback in **Meas**.

Proof. Consider the *interaction cube*:

$$\begin{array}{ccccc}
 & & \mathcal{M}_{\underline{\mathcal{T}} \otimes \underline{\mathcal{S}}}(y) & \longrightarrow & \mathcal{M}_{\underline{\mathcal{T}} \otimes \underline{\mathcal{S}}}(x) \\
 & & \downarrow \text{nt}^{\Pi_2} & & \downarrow \text{nt}^{\Pi_2} \\
 & & \mathcal{M}_{\underline{\mathcal{A}} \parallel \underline{\mathcal{T}}}(y_A \parallel y_T) & \xrightarrow{\text{nt}^{\Pi_1}} & \mathcal{M}_{\underline{\mathcal{A}} \parallel \underline{\mathcal{T}}}(x_A \parallel x_T) \\
 & \swarrow \text{nt}^{\Pi_1} & & \swarrow \text{nt}^{\Pi_1} & \\
 \mathcal{M}_{\underline{\mathcal{S}} \parallel \underline{\mathcal{C}}}(y_S \parallel y_C) & \xrightarrow{\text{nt}^{\mathcal{A} \parallel \tau}} & \mathcal{M}_{\underline{\mathcal{S}} \parallel \underline{\mathcal{C}}}(x_S \parallel x_C) & & \\
 \downarrow \text{nt}^{\sigma \parallel \mathcal{C}} & & \downarrow \text{nt}^{\sigma \parallel \mathcal{C}} & & \\
 \mathcal{M}_{\underline{\mathcal{A}} \parallel \underline{\mathcal{B}} \parallel \underline{\mathcal{C}}}(y_A \parallel y_B \parallel y_C) & \longrightarrow & \mathcal{M}_{\underline{\mathcal{A}} \parallel \underline{\mathcal{B}} \parallel \underline{\mathcal{C}}}(x_A \parallel x_B \parallel x_C) & &
 \end{array}$$

The left and right faces are pullbacks by definition. By measurable receptivity of τ (combined with the obvious fact that

$$\begin{array}{ccc} \mathcal{M}_{\mathcal{A}}(y_A) & \longrightarrow & \mathcal{M}_{\mathcal{A}}(x_A) \\ \downarrow & & \downarrow \\ \mathcal{M}_{\mathcal{A}}(y_A) & \longrightarrow & \mathcal{M}_{\mathcal{A}}(x_A) \end{array}$$

is a pullback), the bottom face of the cube is also a pullback. Then, standard categorical reasoning shows that the top face is a pullback, which is the desired result. \square

Of course, the dual result holds: if $x \text{--} \text{c}^e y$ for a τ -action e , then

$$\begin{array}{ccc} \mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(y) & \xrightarrow{\mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(x \text{--} y)} & \mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(x) \\ \text{nt}_y^{\Pi_2} \downarrow & & \downarrow \text{nt}_x^{\Pi_2} \\ \mathcal{M}_{\mathcal{A} \parallel \mathcal{T}}(y_A \parallel y_T) & \longrightarrow & \mathcal{M}_{\mathcal{A} \parallel \mathcal{T}}(x_A \parallel x_T) \end{array}$$

is a pullback.

Hiding. The ess $\mathcal{T} \odot \mathcal{S}$ is defined as usual, by restricting to the *visible* events of $\mathcal{T} \otimes \mathcal{S}$. There is a functor

$$[-] : \mathcal{C}(\mathcal{T} \odot \mathcal{S}) \rightarrow \mathcal{C}(\mathcal{T} \otimes \mathcal{S})$$

which assigns to each configuration its unique witness in the interaction. Then, the functor $\mathcal{M}_{\mathcal{T} \odot \mathcal{S}} : \mathcal{C}(\mathcal{T} \odot \mathcal{S})^{\text{op}} \rightarrow \mathbf{Meas}$ is simply defined as $\mathcal{M}_{\mathcal{T} \otimes \mathcal{S}} \circ [-]^{\text{op}}$. In other words, for every $x \in \mathcal{C}(\mathcal{T} \odot \mathcal{S})$, $\mathcal{M}_{\mathcal{T} \odot \mathcal{S}}(x) := \mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}([x])$.

We define the associated natural transformation $\text{nt}^{\tau \odot \sigma} : \mathcal{M}_{\mathcal{T} \odot \mathcal{S}} \Rightarrow \mathcal{M}_{\mathcal{A}^\perp \parallel \mathcal{C}} \circ (\tau \odot \sigma)^{\text{op}}$ as

$$\mathcal{M}_{\mathcal{T} \odot \mathcal{S}}(x) = \mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}([x]) \xrightarrow{\text{nt}_{[x]}^{\tau \otimes \sigma}} \mathcal{M}_{\mathcal{A} \parallel \mathcal{B} \parallel \mathcal{C}}((\tau \otimes \sigma)[x]) \xrightarrow{\pi_{\mathcal{A} \parallel \mathcal{C}}} \mathcal{M}_{\mathcal{A} \parallel \mathcal{C}}((\tau \odot \sigma)x).$$

Lemma 6.20. *The composition $\underline{\tau} \odot \underline{\sigma} : \underline{\mathcal{T}} \odot \underline{\mathcal{S}} \rightarrow \underline{\mathcal{A}}^\perp \parallel \underline{\mathcal{C}}$ is a measurable strategy.*

Proof. Because strategies are closed under composition, all that must be checked is the measurable receptivity axiom.

We first discuss negative extensions in the interaction $\underline{\mathcal{T}} \otimes \underline{\mathcal{S}}$. Suppose $x \text{--} \text{c}^e y \in \mathcal{C}(\underline{\mathcal{T}} \otimes \underline{\mathcal{S}})$ for e a negative (so visible) event; wlog assume e is an A -move, so

$x_S \text{--} \ulcorner^{\Pi_1 e} y_S$ and $x_A \text{--} \ulcorner^{(\tau \otimes \sigma) e} y_A$. We reason once again using the interaction cube:

$$\begin{array}{ccccc}
& & \mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(y) & \longrightarrow & \mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(x) \\
& & \downarrow & & \downarrow \\
& & \mathcal{M}_{\mathcal{A} \parallel \mathcal{T}}(y_A \parallel y_T) & \longrightarrow & \mathcal{M}_{\mathcal{A} \parallel \mathcal{T}}(x_A \parallel x_T) \\
& \swarrow & & \searrow & \\
\mathcal{M}_{\mathcal{S} \parallel \mathcal{C}}(y_S \parallel y_C) & \longrightarrow & \mathcal{M}_{\mathcal{S} \parallel \mathcal{C}}(x_S \parallel x_C) & & \\
\downarrow & & \downarrow & & \\
\mathcal{M}_{\mathcal{A} \parallel \mathcal{B} \parallel \mathcal{C}}(y_A \parallel y_B \parallel y_C) & \longrightarrow & \mathcal{M}_{\mathcal{A} \parallel \mathcal{B} \parallel \mathcal{C}}(x_A \parallel x_B \parallel x_C) & &
\end{array}$$

Since $y_T = x_T$, it is straightforward to check the bottom face is a pullback. The left and right faces are pullbacks, so the top face is a pullback. By measurable receptivity of $\underline{\sigma}$, the front face is a pullback. We paste the front and top pullback squares to get that

$$\begin{array}{ccc}
\mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(y) & \longrightarrow & \mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(x) \\
\downarrow \text{nt}_y^{\tau \otimes \sigma} & & \downarrow \text{nt}_x^{\tau \otimes \sigma} \\
\mathcal{M}_{\mathcal{A} \parallel \mathcal{B} \parallel \mathcal{C}}(y_A \parallel y_B \parallel y_C) & \longrightarrow & \mathcal{M}_{\mathcal{A} \parallel \mathcal{B} \parallel \mathcal{C}}(x_A \parallel x_B \parallel x_C)
\end{array}$$

is a pullback. We easily deduce that the following is also a pullback:

$$\begin{array}{ccc}
\mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(y) & \longrightarrow & \mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(x) \\
\downarrow \pi_{\mathcal{A} \parallel \mathcal{C}} \circ \text{nt}_y^{\tau \otimes \sigma} & & \downarrow \pi_{\mathcal{A} \parallel \mathcal{C}} \circ \text{nt}_x^{\tau \otimes \sigma} \\
\mathcal{M}_{\mathcal{A} \parallel \mathcal{C}}(y_A \parallel y_C) & \longrightarrow & \mathcal{M}_{\mathcal{A} \parallel \mathcal{C}}(x_A \parallel x_C)
\end{array}$$

Going back to the composition $\underline{\tau} \odot \underline{\sigma}$ — let $z, w \in \mathcal{C}(T \odot S)$ with $z \text{--} \ulcorner^- w$. By courtesy, no negative event can immediately depend on a non-visible one, so $[z] \text{--} \ulcorner^e [w]$ in $\mathcal{C}(T \otimes S)$ for some negative event e , and we are in the situation discussed above. The fact that

$$\begin{array}{ccc}
\mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}([w]) & \longrightarrow & \mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}([z]) \\
\downarrow \pi_{\mathcal{A} \parallel \mathcal{C}} \circ \text{nt}_y^{\tau \otimes \sigma} & & \downarrow \pi_{\mathcal{A} \parallel \mathcal{C}} \circ \text{nt}_x^{\tau \otimes \sigma} \\
\mathcal{M}_{\mathcal{A} \parallel \mathcal{C}}(w_A \parallel w_C) & \longrightarrow & \mathcal{M}_{\mathcal{A} \parallel \mathcal{C}}(z_A \parallel z_C)
\end{array}$$

is a pullback corresponds precisely to the measurable receptivity of $\underline{\tau} \odot \underline{\sigma}$. \square

6.3 The pseudo-double category \mathcal{MG}

Measurable games and strategies can be organised into a bicategory, which we will call \mathbf{MG} . This is a generalisation of the bicategory \mathbf{G} defined in Chapter 2, and we will see that it retains the necessary structural properties.

In the same way as for \mathbf{G} , we start by showing that there is a pseudo-double category \mathcal{MG} which is isofibrant and symmetric monoidal. We will then focus

on its horizontal bicategory (in the sense of Definition 2.29) and show it admits a sub-bicategory which is symmetric monoidal closed, has finite products and supports a linear exponential pseudo-comonad.

The purpose of this section is to give the details of its construction.

6.3.1 Measurable copycat

For a measurable game \underline{A} , the construction of $\mathbb{c}_{\underline{A}} : \mathbb{C}_{\underline{A}} \rightarrow \underline{A}^\perp \parallel \underline{A}$ goes in the following way.

Recall that configurations of $\mathbb{C}_{\underline{A}}$ are of the form $x \parallel y$ for some $x, y \in \mathcal{C}(\underline{A})$ such that $y \sqsubseteq x$ (that is, $y \sqsupseteq^- x \cap y \sqsubseteq^+ x$). Moreover, morphisms $x \parallel y \rightarrow z \parallel w$ in the category $\mathcal{C}(\mathbb{C}_{\underline{A}})$ are all of the form $f \parallel g$ for some $f : x \rightarrow z$ and $g : y \rightarrow w$ in $\mathcal{C}(\underline{A})$ which agree on $x \cap y$; so in particular from such a pair f, g we can define $f \cap g : x \cap y \rightarrow z \cap w$.

We define the presheaf $\mathcal{M}_{\mathbb{C}_{\underline{A}}} : \mathcal{C}(\mathbb{C}_{\underline{A}})^{\text{op}} \rightarrow \mathbf{Meas}$ as follows. For an object $x \parallel y$, $\mathcal{M}_{\mathbb{C}_{\underline{A}}}(x \parallel y)$ is defined as the pullback

$$\begin{array}{ccc} \mathcal{M}_{\mathbb{C}_{\underline{A}}}(x \parallel y) & \xrightarrow{\Pi_2} & \mathcal{M}_{\underline{A}}(y) \\ \downarrow \Pi_1 & & \downarrow \mathcal{M}_{\underline{A}}(x \cap y \hookrightarrow y) \\ \mathcal{M}_{\underline{A}}(x) & \xrightarrow{\mathcal{M}_{\underline{A}}(x \cap y \hookrightarrow x)} & \mathcal{M}_{\underline{A}}(x \cap y) \end{array}$$

Then, to define the action of $\mathcal{M}_{\mathbb{C}_{\underline{A}}}$ on morphisms, observe that for any $f \parallel g : x \parallel y \rightarrow z \parallel w$ the following diagram commutes:

$$\begin{array}{ccccc} \mathcal{M}_{\mathbb{C}_{\underline{A}}}(z \parallel w) & \longrightarrow & \mathcal{M}_{\underline{A}}(w) & & \\ \downarrow & \dashrightarrow & \downarrow & \searrow \mathcal{M}_{\underline{A}}(g) & \\ \mathcal{M}_{\underline{A}}(z) & \longrightarrow & \mathcal{M}_{\underline{A}}(z \cap w) & & \\ & \searrow \mathcal{M}_{\underline{A}}(f) & \downarrow & \searrow \mathcal{M}_{\underline{A}}(f \cap g) & \\ & & \mathcal{M}_{\mathbb{C}_{\underline{A}}}(x \parallel y) & \xrightarrow{\mathcal{M}_{\underline{A}}(f \cap g)} & \mathcal{M}_{\underline{A}}(y) \\ & & \downarrow & & \downarrow \\ & & \mathcal{M}_{\underline{A}}(x) & \longrightarrow & \mathcal{M}_{\underline{A}}(x \cap y) \end{array}$$

where the dashed arrow is the mediating map for the pullback $\mathcal{M}_{\mathbb{C}_{\underline{A}}}(x \parallel y)$; this is what we take as $\mathcal{M}_{\mathbb{C}_{\underline{A}}}(f \parallel g)$.

It remains to define $\text{nt}^{\mathbb{c}_{\underline{A}}} : \mathcal{M}_{\mathbb{C}_{\underline{A}}} \xrightarrow{\bullet} \mathcal{M}_{\underline{A}^\perp \parallel \underline{A}} \circ \mathbb{c}_{\underline{A}}^{\text{op}}$. For $x \parallel y \in \mathcal{C}(\mathbb{C}_{\underline{A}})$, $\text{nt}_{x \parallel y}^{\mathbb{c}_{\underline{A}}}$ is the pairing $\langle \Pi_1, \Pi_2 \rangle$ of the pullback projections, using that $\mathcal{M}_{\underline{A}^\perp \parallel \underline{A}}(\mathbb{c}_{\underline{A}}(x \parallel y)) = \mathcal{M}_{\underline{A}}(x) \times \mathcal{M}_{\underline{A}}(y)$, as shown below:

$$\begin{array}{ccc} \mathcal{M}_{\mathbb{C}_{\underline{A}}}(x \parallel y) & \xrightarrow{\Pi_2} & \mathcal{M}_{\underline{A}}(y) \\ \downarrow \Pi_1 & \dashrightarrow \text{nt}_{x \parallel y}^{\mathbb{c}_{\underline{A}}} & \downarrow \pi_2 \\ & \mathcal{M}_{\underline{A}}(x) \times \mathcal{M}_{\underline{A}}(y) & \xrightarrow{\pi_2} & \mathcal{M}_{\underline{A}}(y) \\ & \downarrow \pi_1 & & \\ & \mathcal{M}_{\underline{A}}(x) & & \end{array}$$

To generalise this we first define the symmetry relation between maps of measurable essps.

Lemma 6.22. *Let $f, g : \mathcal{E} \rightarrow \mathcal{E}'$. If $f \sim g$, the family $\varphi = \{\varphi_x\}_x$ (where $\varphi_x = \{(f(e), g(e)) \mid e \in x\}$) is a natural isomorphism between f and g seen as functors $\mathcal{C}(\mathcal{E}) \rightarrow \mathcal{C}(\mathcal{E}')$.*

Proof. For $x \in \mathcal{C}(\mathcal{E})$, φ_x is a morphism $fx \rightarrow gx$ in $\mathcal{C}(\mathcal{E}')$. Naturality is verified directly. \square

We can now define:

Definition 6.23. Maps of measurable essps $(f, \text{nt}^f), (g, \text{nt}^g) : (\mathcal{E}, \mathcal{M}_{\mathcal{E}}) \rightarrow (\mathcal{E}', \mathcal{M}_{\mathcal{E}'})$ are **symmetric** (written $(f, \text{nt}^f) \sim (g, \text{nt}^g)$) if $f \sim g$ and the following diagram commutes in $[\mathcal{C}(\mathcal{E})^{\text{op}}, \mathbf{Meas}]$:

$$\begin{array}{ccc} \mathcal{M}_{\mathcal{E}} & \xrightarrow{\text{nt}^f} & \mathcal{M}_{\mathcal{E}'} \circ f^{\text{op}} \\ \downarrow \text{nt}^g & \nearrow \mathcal{M}_{\mathcal{E}'} \circ \varphi & \\ \mathcal{M}_{\mathcal{E}'} \circ g^{\text{op}} & & \end{array}$$

(Note that $\varphi : f \overset{\bullet}{\rightarrow} g$ is also a natural transformation $g^{\text{op}} \overset{\bullet}{\rightarrow} f^{\text{op}}$.) They are **positively symmetric** if in addition, $f \sim^+ g$.

Then, maps between measurable strategies are given by:

Definition 6.24. For measurable strategies $\underline{\sigma} : \underline{\mathcal{S}} \rightarrow \underline{\mathcal{A}}^{\perp} \parallel \underline{\mathcal{B}}$ and $\underline{\sigma}' : \underline{\mathcal{S}} \rightarrow \underline{\mathcal{A}}'^{\perp} \parallel \underline{\mathcal{B}}'$, a **map from $\underline{\sigma}$ to $\underline{\sigma}'$** is $(\underline{f}, \underline{g}, \underline{h})$ where \underline{f} is a map of measurable essps and $\underline{g}, \underline{h}$ are maps of measurable games (*i.e.* maps of measurable essps whose underlying map is a map of games), such that

$$\begin{array}{ccc} \underline{\mathcal{S}} & \xrightarrow{f} & \underline{\mathcal{S}}' \\ \underline{\sigma} \downarrow & \sim^+ & \downarrow \underline{\sigma}' \\ \underline{\mathcal{A}}^{\perp} \parallel \underline{\mathcal{B}} & \xrightarrow{g^{\perp} \parallel h} & \underline{\mathcal{A}}'^{\perp} \parallel \underline{\mathcal{B}}' \end{array}$$

Vertical composition of maps of measurable strategies is done as in \mathcal{G} . We must verify the additional axiom in the definition of positive symmetry; this is not difficult and we omit the details.

Horizontal composition is more involved. Suppose we have maps

$$\begin{array}{ccc} \underline{\mathcal{S}} & \xrightarrow{f} & \underline{\mathcal{S}}' & \quad & \underline{\mathcal{T}} & \xrightarrow{g} & \underline{\mathcal{T}}' \\ \underline{\sigma} \downarrow & \sim^+ & \downarrow \underline{\sigma}' & & \underline{\sigma} \downarrow & \sim^+ & \downarrow \underline{\sigma}' \\ \underline{\mathcal{A}}^{\perp} \parallel \underline{\mathcal{B}} & \xrightarrow{h_1^{\perp} \parallel h_2} & \underline{\mathcal{A}}'^{\perp} \parallel \underline{\mathcal{B}}' & & \underline{\mathcal{B}}^{\perp} \parallel \underline{\mathcal{C}} & \xrightarrow{h_2^{\perp} \parallel h_3} & \underline{\mathcal{B}}'^{\perp} \parallel \underline{\mathcal{C}}' \end{array}$$

To obtain the necessary map $\underline{g} \odot \underline{f} : \underline{\mathcal{T}} \odot \underline{\mathcal{S}} \rightarrow \underline{\mathcal{T}'} \odot \underline{\mathcal{S}'}$, we must equip $g \odot f$ with a natural transformation

$$\text{nt}^{g \odot f} : \mathcal{M}_{\mathcal{T} \odot \mathcal{S}} \xrightarrow{\bullet} \mathcal{M}_{\mathcal{T}' \odot \mathcal{S}'} \circ (g \odot f)^{\text{op}}.$$

Explicitly, this is a family of maps

$$\text{nt}_{x_T \odot x_S}^{g \odot f} : \mathcal{M}_{\mathcal{T} \odot \mathcal{S}}(x_T \odot x_S) \rightarrow \mathcal{M}_{\mathcal{T}' \odot \mathcal{S}'}((g \odot f)(x_T \odot x_S))$$

and to give this is to give a map $\mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(x_T \otimes x_S) \rightarrow \mathcal{M}_{\mathcal{T}' \otimes \mathcal{S}'}((g \otimes f)(x_T \otimes x_S))$ for each $x_T \odot x_S$. Recall that $(g \otimes f)(x_T \otimes x_S)$ is a configuration $y_{T'} \otimes y_{S'} \in \mathcal{C}(T' \otimes S')$ with the property that there exist bijections

$$\phi : f x_S \parallel h_3 x_C \cong_{S' \parallel C'} y_{S'} \parallel y_{C'} \quad \psi : h_1 x_A \parallel g x_T \cong_{A' \parallel T'} y_{A'} \parallel y_{T'}$$

such that the following diagram commutes:

$$\begin{array}{ccc} & h_1 x_A \parallel h_2 x_B \parallel h_3 x_C & \\ \varphi_{x_S} \parallel h_3 x_C \swarrow & & \searrow h_1 x_A \parallel \varphi_{x_T} \\ \sigma'(f x_S) \parallel h_3 x_C & & h_1 x_A \parallel \tau'(g x_T) \\ (\sigma' \parallel C') \phi \searrow & & \swarrow (A' \parallel T') \psi \\ & y_{A'} \parallel y_{B'} \parallel y_{C'} & \end{array}$$

(where as usual the maps φ_{x_S} and φ_{x_T} are canonical bijections, obtained by the positive symmetry requirement in the definition of maps of strategies).

Consider the following diagram, where we write $h = h_1 \parallel h_2 \parallel h_3$:

$$\begin{array}{ccccc} & & \mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(x_T \otimes x_S) & & \\ & & \swarrow \text{nt}^{\Pi_1} & & \searrow \text{nt}^{\Pi_2} \\ & \mathcal{M}_{S \parallel C}(x_S \parallel x_C) & & & \mathcal{M}_{A \parallel T}(x_A \parallel x_T) \\ & \swarrow \text{nt}^f \parallel h_3 & & \swarrow \text{nt}^{h \circ (\sigma \parallel C)} & \swarrow \text{nt}^{h \circ (A \parallel \tau)} \\ \mathcal{M}_{S' \parallel C'}(f x_S \parallel h_3 x_C) & & \mathcal{M}_{A' \parallel B' \parallel C'}(h(x_A \parallel x_B \parallel x_C)) & & \mathcal{M}_{A' \parallel T'}(h_1 x_A \parallel g x_T) \\ & \swarrow \text{nt}^{\sigma' \parallel C'} & \swarrow \mathcal{M}_{A' B' C'}(\varphi_{x_S} \parallel C') & \swarrow \mathcal{M}_{A' B' C'}(A' \parallel \varphi_{x_T}) & \swarrow \text{nt}^{A' \parallel \tau'} \\ \mathcal{M}_{S' \parallel C'}(\phi) & \mathcal{M}_{A' \parallel B' \parallel C'}(\sigma'(f x_S) \parallel h_3 x_C) & & \mathcal{M}_{A' \parallel B' \parallel C'}(h_1 x_A \parallel \tau'(g x_T)) & \mathcal{M}_{A' \parallel T'}(\psi) \\ & \swarrow \mathcal{M}_{A' B' C'}((\sigma' \parallel C') \phi) & & \swarrow \mathcal{M}_{A' B' C'}((A' \parallel \tau') \psi) & \\ \mathcal{M}_{S' \parallel C'}(y_{S'} \parallel y_{C'}) & & \mathcal{M}_{A' \parallel B' \parallel C'}((\sigma' \parallel C') \phi) & & \mathcal{M}_{A' \parallel T'}(y_{A'} \parallel y_{T'}) \\ & \swarrow \text{nt}^{\sigma' \parallel C'} & & \swarrow \text{nt}^{A' \parallel \tau'} & \\ & & \mathcal{M}_{A' \parallel B' \parallel C'}(y_{A'} \parallel y_{B'} \parallel y_{C'}) & & \end{array}$$

This commutes because every sub-diagram does: the top square commutes by definition, the bottom-left and bottom-right squares are instances of the naturality axiom for $\text{nt}^{\sigma' \parallel C'}$ and $\text{nt}^{A' \parallel \tau'}$, the internal square commutes by the discussion of

the previous paragraph, and the remaining two commute because of the symmetry requirement in the definition of maps of strategies.

By definition, $\mathcal{M}_{\mathcal{T}' \otimes \mathcal{S}'}(y_{T'} \otimes y_{C'})$ is obtained as the pullback of the bottom-most two maps in the diagram, so the universal property gives a map $\mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(x_T \otimes x_C) \rightarrow \mathcal{M}_{\mathcal{T}' \otimes \mathcal{S}'}(y_{T'} \otimes y_{C'})$ which we take as $\mathbf{nt}_{x_T \otimes x_S}^{g \odot f}$.

Lemma 6.25. *The family of maps $\mathbf{nt}^{g \odot f}$ is a natural transformation, and $\underline{g} \odot \underline{f} = (g \odot f, \mathbf{nt}^{g \odot f})$ is a map of measurable strategies.*

Proof. For the sake of readability, we have given the componentwise definition of $\mathbf{nt}^{g \odot f}$ above, but observing that every element of the diagram only involves natural transformations, the family $\mathbf{nt}^{g \odot f}$ can be directly obtained as a mediating map in the functor category $[\mathcal{C}(\mathcal{T} \odot \mathcal{S})^{\text{op}}, \mathbf{Meas}]$. This makes it a natural transformation; we omit the details.

It remains to check that $\underline{g} \odot \underline{f}$ is a map of measurable strategies. By definition of $\mathbf{nt}^{g \odot f}$, for any $x_T \odot x_S$ we have that the following commutes:

$$\begin{array}{ccc} \mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(x_T \otimes x_S) & \xrightarrow{\mathbf{nt}^{g \odot f}} & \mathcal{M}_{\mathcal{T}' \otimes \mathcal{S}'}(y_{T'} \otimes y_{C'}) \\ \downarrow \mathbf{nt}^{h \odot (\tau \otimes \sigma)} & & \downarrow \mathbf{nt}^{\tau' \otimes \sigma'} \\ \mathcal{M}_{\mathcal{A}' \parallel \mathcal{B}' \parallel \mathcal{C}'}(h_1 x_A \parallel h_2 x_B \parallel h_3 x_C) & \longrightarrow & \mathcal{M}_{\mathcal{A}' \parallel \mathcal{B}' \parallel \mathcal{C}'}(y_{A'} \parallel y_{B'} \parallel y_{C'}) \end{array}$$

where the bottom arrow coincides, after hiding, with the image under of $\mathcal{M}_{\mathcal{A}' \parallel \mathcal{B}' \parallel \mathcal{C}'}$ of the canonical bijection φ obtained from the symmetry $(h_1^\perp \parallel h_3) \circ (\tau \odot \sigma) \sim^+ (\tau' \odot \sigma') \circ (g \odot f)$. \square

6.3.3 Structural isomorphisms

We will get a pseudo-double category after specifying some structural 2-cells: an *associator* $\underline{\alpha}_{\rho, \tau, \sigma} : (\underline{\rho} \odot \underline{\tau}) \odot \underline{\sigma} \Rightarrow \underline{\rho} \odot (\underline{\tau} \odot \underline{\sigma})$ for every $\underline{\sigma} : \underline{\mathcal{S}} \rightarrow \underline{\mathcal{A}}^\perp \parallel \underline{\mathcal{B}}$, $\underline{\tau} : \underline{\mathcal{T}} \rightarrow \underline{\mathcal{B}}^\perp \parallel \underline{\mathcal{C}}$ and $\underline{\rho} : \underline{\mathcal{R}} \rightarrow \underline{\mathcal{C}}^\perp \parallel \underline{\mathcal{D}}$, and *unitors* $\underline{\lambda}_\sigma : \underline{\alpha}_{\sigma, \tau, \rho} \Rightarrow \sigma$ and $\underline{\rho}_\sigma : \sigma \odot \underline{\alpha}_A \Rightarrow \sigma$.

Associator. We start with the associator, which has underlying map of essps the associator in \mathcal{G} , a globular map $\alpha_{\sigma, \tau, \rho} : (\sigma \odot \tau) \odot \rho \Rightarrow \sigma \odot (\tau \odot \rho)$. We adjoin to it a natural transformation

$$\mathbf{nt}^{\alpha_{\sigma, \tau, \rho}} : \mathcal{M}_{(\mathcal{R} \odot \mathcal{T}) \odot \mathcal{S}} \xrightarrow{\bullet} \mathcal{M}_{\mathcal{R} \odot (\mathcal{T} \odot \mathcal{S})} \circ \alpha_{\sigma, \tau, \rho}^{\text{op}}$$

whose components must be measurable functions of the form

$$\mathcal{M}_{(\mathcal{R} \otimes \mathcal{T}) \otimes \mathcal{S}}((x_R \otimes x_T) \otimes x_S) \rightarrow \mathcal{M}_{\mathcal{R} \otimes (\mathcal{T} \otimes \mathcal{S})}(x_R \otimes (x_T \otimes x_S)).$$

This is obtained canonically since both the domain and co-domain can be seen to rise as (ternary) pullbacks of the same diagram.

Unitors. Recall that the unitors in \mathbf{G} are strong isomorphisms

$$\begin{aligned}\rho_\sigma &: \mathcal{S} \odot \mathbb{C}_A \Rightarrow \mathcal{S} \\ \lambda_\sigma &: \mathbb{C}_B \odot \mathcal{S} \Rightarrow \mathcal{S}\end{aligned}$$

defined for every $\sigma : \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$, whose action on configurations is given in Lemma 2.22. Let $\underline{\sigma} : \underline{\mathcal{S}} \rightarrow \underline{\mathcal{A}}^\perp \parallel \underline{\mathcal{B}}$ be a measurable strategy. In the proof of the next lemma, we extend λ_σ to an isomorphism of measurable strategies:

Lemma 6.26. *For any $\sigma : \mathcal{S} \rightarrow \mathcal{A}^\perp \parallel \mathcal{B}$, the functors $\mathcal{M}_{\mathbb{C}_B \odot \mathcal{S}}$ and $\mathcal{M}_\mathcal{S} \circ \lambda_\sigma^{\text{op}}$ are naturally isomorphic, and the induced family of maps $\alpha_{\underline{\mathcal{B}}} \odot \underline{\sigma} \Rightarrow \underline{\sigma}$ is natural in $\underline{\sigma}$.*

Proof. We give an isomorphism $\mathcal{M}_{\mathbb{C}_B \odot \mathcal{S}} \circ \lambda_\sigma^{-1} \xrightarrow{\cdot} \mathcal{M}_\mathcal{S}$. Lemma 2.22 characterises the action of λ_σ on configurations; its inverse λ_σ^{-1} sends $x_S \in \mathcal{C}(\mathcal{S})$ to $(x_B^* \parallel x_B) \odot x_S^*$ where x_S^* is the maximal sub-configuration of x_S whose maximal B -moves are all positive, and $\sigma x_S^* = x_A \parallel x_B^*$. Note in particular that $x_B^* \subseteq^- x_B$ and $x_S^* \subseteq^- x_S$.

First, by definition of $\mathcal{M}_{\mathbb{C}_B}$ and using that $x_B^* \cap x_B = x_B^*$,

$$\begin{array}{ccc}\mathcal{M}_{\mathbb{C}_B}(x_B^* \parallel x_B) & \xrightarrow{\Pi_1} & \mathcal{M}_B(x_B^*) \\ \downarrow \Pi_2 & & \downarrow \text{id} \\ \mathcal{M}_B(x_B) & \xrightarrow{\mathcal{M}_B(x_B^* \hookrightarrow x_B)} & \mathcal{M}_B(x_B^*)\end{array}$$

is a pullback, so in particular Π_2 is an isomorphism. Observe also that

$$\begin{array}{ccc}\mathcal{M}_{\mathbb{C}_B}(x_B^* \parallel x_B) & \xrightarrow{\text{nt}_{x_B^* \parallel x_B}^{\alpha_B}} & \mathcal{M}_B(x_B^*) \times \mathcal{M}_B(x_B) \\ \downarrow \Pi_2 & \searrow \Pi_1 & \downarrow \pi_1 \\ \mathcal{M}_B(x_B) & \xrightarrow{\mathcal{M}_B(x_B^* \hookrightarrow x_B)} & \mathcal{M}_B(x_B^*)\end{array}$$

commutes, since nt^{α_B} is defined as the canonical injection of the pullback into the product.

By definition, we also have $\mathcal{M}_{\mathbb{C}_B \odot \mathcal{S}}((x_B^* \parallel x_B) \odot x_S^*) = \mathcal{M}_{\mathbb{C}_B \otimes \mathcal{S}}((x_B^* \parallel x_B) \otimes x_S^*)$ and this is the pullback

$$\begin{array}{ccc}\mathcal{M}_{\mathbb{C}_B \odot \mathcal{S}}((x_B^* \parallel x_B) \odot x_S^*) & & \\ \swarrow & & \searrow \\ \mathcal{M}_{\mathcal{S} \parallel \mathcal{B}}(x_S^* \parallel x_B) & & \mathcal{M}_{\mathcal{A} \parallel \mathbb{C}_B}(x_A \parallel x_B^* \parallel x_B) \\ \searrow \text{nt}^{\sigma \parallel \mathcal{B}} & & \swarrow \text{nt}^{\mathcal{A} \parallel \alpha_B} \\ & \mathcal{M}_{\mathcal{A} \parallel \mathcal{B} \parallel \mathcal{B}}(x_A \parallel x_B^* \parallel x_B) & \end{array}$$

By composing on the left with projections $\mathcal{M}_{\mathcal{S} \parallel \mathcal{B}}(x_S^* \parallel x_B) \xrightarrow{\pi_1} \mathcal{M}_\mathcal{S}(x_S^*)$ and

$\mathcal{M}_{\mathcal{A}\|\mathcal{B}\|\mathcal{B}}(x_A \parallel x_B^* \parallel x_B) \xrightarrow{\mathcal{M}_{\mathcal{A}}(x_A) \times \pi_1} \mathcal{M}_{\mathcal{A}\|\mathcal{B}}(x_A \parallel x_B^*)$, we derive that

$$\begin{array}{ccc}
& \mathcal{M}_{\mathbb{C}_B \odot \mathcal{S}}((x_B^* \parallel x_B) \odot x_S^*) & \\
\pi_1 \circ \text{nt}^{\Pi_1} \swarrow & & \searrow \text{nt}^{\Pi_2} \\
\mathcal{M}_{\mathcal{S}}(x_S^*) & & \mathcal{M}_{\mathcal{A}\|\mathbb{C}_B}(x_A \parallel x_B^* \parallel x_B) \\
\text{nt}^\sigma \searrow & & \swarrow \mathcal{M}_{\mathcal{A}}(x_A) \times (\pi_1 \circ \text{nt}^{\mathbb{C}_B}) \\
& \mathcal{M}_{\mathcal{A}\|\mathcal{B}}(x_A \parallel x_B^*) &
\end{array}$$

is a pullback. Combining this with the remarks in the previous paragraph, we get that

$$\begin{array}{ccc}
& \mathcal{M}_{\mathbb{C}_B \odot \mathcal{S}}((x_B^* \parallel x_B) \odot x_S^*) & \\
\pi_1 \circ \text{nt}^{\Pi_1} \swarrow & & \searrow \text{nt}^{\Pi_2} \\
\mathcal{M}_{\mathcal{S}}(x_S^*) & & \mathcal{M}_{\mathcal{A}\|\mathbb{C}_B}(x_A \parallel x_B^* \parallel x_B) \xrightarrow{\cong} \mathcal{M}_{\mathcal{A}\|\mathcal{B}}(x_A \parallel x_B) \\
\text{nt}^\sigma \searrow & & \swarrow \mathcal{M}_{\mathcal{A}}(x_A) \times (\pi_1 \circ \text{nt}^{\mathbb{C}_B}) \\
& \mathcal{M}_{\mathcal{A}\|\mathcal{B}}(x_A \parallel x_B^*) & \xleftarrow{\mathcal{M}_{\mathcal{A}\|\mathcal{B}}(x_A \parallel x_B^* \hookrightarrow x_A \parallel x_B)}
\end{array}$$

commutes, and hence (rearranging) that

$$\begin{array}{ccc}
\mathcal{M}_{\mathbb{C}_B \odot \mathcal{S}}((x_B^* \parallel x_B) \odot x_S^*) & \longrightarrow & \mathcal{M}_{\mathcal{S}}(x_S^*) \\
\downarrow & & \downarrow \text{nt}^\sigma \\
\mathcal{M}_{\mathcal{A}\|\mathcal{B}}(x_A \parallel x_B) & \xrightarrow{\mathcal{M}_{\mathcal{A}\|\mathcal{B}}(x_A \parallel x_B^* \hookrightarrow x_A \parallel x_B)} & \mathcal{M}_{\mathcal{A}\|\mathcal{B}}(x_A \parallel x_B^*)
\end{array}$$

is a pullback. But by measurable receptivity of $\underline{\sigma}$,

$$\begin{array}{ccc}
\mathcal{M}_{\mathcal{S}}(x_S) & \xrightarrow{\mathcal{M}_{\mathcal{S}}(x_S^* \subseteq x_S)} & \mathcal{M}_{\mathcal{S}}(x_S^*) \\
\text{nt}^\sigma \downarrow & & \downarrow \text{nt}^\sigma \\
\mathcal{M}_{\mathcal{A}\|\mathcal{B}}(x_A \parallel x_B) & \xrightarrow{\mathcal{M}_{\mathcal{A}\|\mathcal{B}}(x_A \parallel x_B^* \hookrightarrow x_A \parallel x_B)} & \mathcal{M}_{\mathcal{A}\|\mathcal{B}}(x_A \parallel x_B^*)
\end{array}$$

is a pullback, and pullbacks are unique up to isomorphism: we get a mediating isomorphism $\mathcal{M}_{\mathbb{C}_B \odot \mathcal{S}}((x_B^* \parallel x_B) \odot x_S^*) \cong \mathcal{M}_{\mathcal{S}}(x_S)$.

The proof that this is a natural transformation is a straightforward diagram chase using naturality of nt^σ . The proof that the induced $\underline{\lambda}_\sigma$ is natural in $\underline{\sigma}$ is very similar, and makes use of the naturality of λ_σ and the additional axiom in the definition of symmetry for maps of measurable essps. The details are easily recovered. \square

Theorem 6.27. *There is a pseudo-double category $\mathcal{M}\mathcal{G}$ having*

- *objects: negative, measurable games;*
- *vertical morphisms: maps of measurable games;*

- *horizontal morphisms: measurable strategies; and*
- *2-cells: maps of measurable strategies.*

The category of objects is written \mathbf{MG}_0 and the category of morphisms \mathbf{MG}_1 .

Proof. We have given all the data. It remains to verify the two coherence axioms – this is done by instantiating the universal property of the appropriate pullback. \square

We proceed to study the categorical structure of measurable games. In passing, we note that the functor $\mathbf{disc} : \mathbf{Ess} \rightarrow \mathbf{Mess}$ defined in 6.1.4 induces a *pseudo-double functor* (the standard notion of morphism between pseudo-double categories) $\mathbf{disc} : \mathcal{G} \rightarrow \mathcal{MG}$.

6.4 Categorical properties

In Chapter 2, we studied in detail the structure of the bicategory \mathbf{G} , and found that, given the appropriate data, it is a symmetric monoidal closed bicategory with finite products and a linear exponential pseudo-comonad. Proceeding in very similar steps, we show that \mathbf{MG} enjoys the same properties.

6.4.1 Symmetric monoidal closed structure

We use Shulman’s theorem (Theorem 2.39). A map $\underline{f} : \underline{\mathcal{A}} \rightarrow \underline{\mathcal{B}}$ which is courteous, receptive, \sim -receptive and measurable receptive can be lifted to a strategy $\widehat{\underline{f}} : \underline{\mathcal{A}} \leftrightarrow \underline{\mathcal{B}}$, and if \underline{f}^\perp satisfies these properties then we get a colifted strategy $\check{\underline{f}} : \underline{\mathcal{B}} \leftrightarrow \underline{\mathcal{A}}$. So \mathcal{MG} is isofibrant.

We have seen (Lemma 6.15) that \mathbf{Mess} is a symmetric monoidal category, and for the same reasons, so is \mathbf{Messp} , the category of measurable essps. This induces a symmetric monoidal structure on both \mathbf{MG}_0 and \mathbf{MG}_1 , with all definitions the same as for \mathcal{G} . (There is one minor subtlety: to ensure that the functor $\mathfrak{c} : \mathbf{MG}_0 \rightarrow \mathbf{MG}_1$ sends the monoidal unit of \mathbf{MG}_0 to that of \mathbf{MG}_1 we must choose the latter to be \mathbb{C}_\emptyset , rather than \emptyset . The two are equal as essps but only isomorphic as measurable essps.)

Lemma 6.28. *The pseudo-double category \mathcal{MG} is symmetric monoidal.*

By Theorem 2.39, its horizontal bicategory $\mathcal{H}(\mathcal{MG})$ is symmetric monoidal. We will focus on a sub-bicategory with finer structure. We apply the same restrictions as for \mathbf{G} and an additional one:

Definition 6.29. A measurable essp \mathcal{E} is **rooted** if $\mathcal{M}_{\mathcal{E}}(\emptyset)$ is a singleton.

Say a measurable game/strategy is rooted when the underlying measurable essp is. Rooted strategies are closed under composition and tensor, and moreover the copycat strategy on a rooted game is itself rooted. So we consider the bicategory \mathbf{MG} having:

- objects: negative, rooted measurable arenas;

- morphisms: negative, well-threaded, rooted measurable strategies;
- 2-cells: maps of measurable strategies.

(A measurable strategy is **well-threaded** and **negative** just when the underlying strategy is.)

The symmetric monoidal structure in \mathbf{MG} , inherited from $\mathcal{H}(\mathcal{MG})$, is *closed*. For (negative) measurable arenas $\underline{\mathcal{B}}, \underline{\mathcal{C}}$ we have $\underline{\mathcal{B}} \multimap \underline{\mathcal{C}} = (\mathcal{B} \multimap \mathcal{C}, \mathcal{M}_{\mathcal{B} \multimap \mathcal{C}})$, where $\mathcal{B} \multimap \mathcal{C}$ is defined as in Chapter 2 and $\mathcal{M}_{\mathcal{B} \multimap \mathcal{C}}(x) = \mathcal{M}_{\mathcal{B}^\perp \parallel \mathcal{C}}(\chi x)$, where $\chi : \mathcal{B} \multimap \mathcal{C} \rightarrow \mathcal{B}^\perp \parallel \mathcal{C}$ is the canonical map. For any $\underline{\mathcal{A}}, \underline{\mathcal{B}}, \underline{\mathcal{C}}$, there is an isomorphism of categories

$$\Phi : \mathbf{MG} [\underline{\mathcal{A}}, \underline{\mathcal{B}} \multimap \underline{\mathcal{C}}] \rightarrow \mathbf{MG} [\underline{\mathcal{A}} \otimes \underline{\mathcal{B}}, \underline{\mathcal{C}}],$$

and to show that \mathbf{MG} is closed it remains to give for $\sigma \in \mathbf{MG} [\underline{\mathcal{A}}, \underline{\mathcal{B}} \multimap \underline{\mathcal{C}}]$ an isomorphism $\text{ev}_{\underline{\mathcal{B}}, \underline{\mathcal{C}}} \odot (\sigma \otimes \alpha_{\underline{\mathcal{B}}}) \cong \Phi(\sigma)$ so that this is natural in σ . We omit the details: the isomorphism of essps is the same as for \mathbf{G} and for the measurable structure the proof is analogous to that for composition with `copycat`.

6.4.2 Products

Additionally \mathbf{MG} has finite products. The terminal object is the same as the monoidal unit, *i.e.* the empty measurable essp.

For binary products, we construct from measurable arenas $\underline{\mathcal{A}}$ and $\underline{\mathcal{B}}$ a measurable arena $\underline{\mathcal{A}} \& \underline{\mathcal{B}} = (\mathcal{A} \& \mathcal{B}, \mathcal{M}_{\mathcal{A} \& \mathcal{B}})$ with $\mathcal{M}_{\mathcal{A} \& \mathcal{B}}(\emptyset) = \{*\}$, and $\mathcal{M}_{\mathcal{A} \& \mathcal{B}}(x)$ defined as $\mathcal{M}_{\mathcal{A}}(x)$ or $\mathcal{M}_{\mathcal{B}}(x)$ according to whether $x \subseteq A$ or $x \subseteq B$. The projections are obtained by co-lifting as for \mathbf{G} , and this makes $\underline{\mathcal{A}} \& \underline{\mathcal{B}}$ into a product.

6.4.3 A pseudo-comonad

The pseudo-comonad $!$ on \mathbf{G} is extended to one on \mathbf{MG} . We first describe the construction on an arbitrary measurable essp. For $\underline{\mathcal{E}} \in \mathbf{Messp}$, let $!\underline{\mathcal{E}} = (!\mathcal{E}, \mathcal{M}_{!\mathcal{E}})$ where for $x = \parallel_{i \in \omega} x_i \in \mathcal{C}(!\mathcal{E})$,

$$\mathcal{M}_{!\mathcal{E}}(x) = \prod_{i \in \omega} \mathcal{M}_{\mathcal{E}}(x_i).$$

The action of $\mathcal{M}_{!\mathcal{E}}$ on morphisms of $\mathcal{C}(!\mathcal{E})$ is well-defined: for inclusions this is clear, and for symmetries, recall that if $y = \parallel_{i \in \omega} y_i \in \mathcal{C}(!\mathcal{E})$ a bijection $\theta : x \cong_{!\mathcal{E}} y$ is determined by a bijection $\pi : \omega \rightarrow \omega$ and symmetries $\theta_i : x_i \cong_E y_{\pi(i)}$. We define $\mathcal{M}_{!\mathcal{E}}(\theta)$ as

$$\prod_{i \in \omega} \mathcal{M}_{\mathcal{E}}(y_i) \xrightarrow{\cong} \prod_{i \in \omega} \mathcal{M}_{\mathcal{E}}(y_{\pi(i)}) \xrightarrow{\prod_i \mathcal{M}_{\mathcal{E}}(\theta_i)} \prod_{i \in \omega} \mathcal{M}_{\mathcal{E}}(x_i).$$

The maps $\underline{\delta}_{\mathcal{E}} : !\mathcal{E} \rightarrow !\mathcal{E}$ and $\underline{\varepsilon}_{\mathcal{E}} : \mathcal{E} \rightarrow !\mathcal{E}$ are turned into maps of measurable essps, with transformations $\mathbf{nt}^{\underline{\delta}}$ and $\mathbf{nt}^{\underline{\varepsilon}}$ obtained via standard product manipulations which we omit. Restricting to positive arenas, we have:

Lemma 6.30. *The triple $(!, \delta_{\mathcal{A}}, \varepsilon_{\mathcal{A}})$ satisfies, up to \sim^+ , the laws for a monad on the subcategory of \mathbf{Messp} having positive arenas as objects.*

Proof. Direct verification. □

We can then lift ! to a pseudo-comonad on \mathbf{MG} . The proof and associated data can be given via the same steps as for \mathbf{G} . Likewise, the “Seely” adjoint equivalence $m_{A,B}$ is obtained exactly as in \mathbf{G} . Because the conditions of Theorem 2.58 are satisfied, we conclude:

Theorem 6.31. *The Kleisli bicategory $\mathbf{MG}_!$ is cartesian closed.*

The cartesian closure of \mathbf{MG} is interesting not least because \mathbf{Meas} itself lacks this property, and there is an embedding of \mathbf{Meas} inside of \mathbf{MG} . The embedding is obtained as follows.

Any measurable space X can be represented as the negative arena $G(X) = \{\ominus \rightarrow \oplus\}$ with $\mathcal{M}_{G(X)}(\{\ominus, \oplus\}) = X$. Then, it is easy to check that for any measurable function $f : X \rightarrow Y$ there is a map of games $\underline{G}(f) : G(X) \rightarrow G(Y)$, defined as the identity map of essps and such that $\text{nt}_{\{\ominus, \oplus\}}^{\underline{G}(f)} = f$. This can be lifted to a strategy $G(f) := \widehat{\underline{G}(f)} : G(X) \leftrightarrow G(Y)$. From here G is easily turned into a faithful, injective on objects pseudo-functor.

Chapter 7

Probability in measurable games

We finally come to the construction of a concurrent games model allowing for non-discrete probabilistic behaviour. Games and strategies now carry additional structure and can be enriched with probability measures in a natural way, when given access to the tools of probability theory. In this chapter, we describe how the compositional machinery of strategies, and the associated bicategorical structure, can be adapted to this new setting.

In moving from discrete to continuous probability, we seek to replace valuations on strategies with the more general notion of *measures* on a measurable space. It turns out that the notion of valuation ($v : \mathcal{C}(S) \rightarrow [0, 1]$) used in Chapters 4 and 5 does not generalise well. The alternative notion of *conditional* valuation discussed in Chapter 3 (see Definition 3.1) is a more appropriate choice.

It is the polarised nature of strategies which makes valuations unsuitable, and as this plays an important part in the technical development, we devote a section (7.2) to a comparison of the two approaches in the context of event structures *without* polarity. Sections 7.3 and 7.4 are concerned with the construction of the games model. We discuss in particular how this generalises the model of Chapter 3.

7.1 Probability theory

7.1.1 Measures and kernels

A **sub-probability measure** on a measurable space (X, Σ_X) is a map $\mu : \Sigma_X \rightarrow [0, 1]$ such that $\mu(\emptyset) = 0$ and such that for any countable family $\{U_i\}_{i \in I} \subseteq \Sigma_X$ with $U_i \cap U_j = \emptyset$ for every $i \neq j$, we have $\mu(\biguplus_i U_i) = \sum_i \mu(U_i)$. For $x \in X$, the **Dirac measure** δ_x is defined as $\delta_x(U) = 1$ if $x \in U$, and 0 otherwise. Finally, given a sub-probability measure μ on X and a non-negative measurable function $g : X \rightarrow \mathbb{R}$, the integral $\int_{x \in X} g(x) \mu(dx)$ is a well-defined element of $[0, \infty)$.

A **sub-probability kernel** [Gir82] from (X, Σ_X) to (Y, Σ_Y) is a map

$$k : X \times \Sigma_Y \rightarrow [0, 1]$$

such that for every $x \in X$ the map $k(x, -)$ is a sub-probability measure, and for every $U \in \Sigma_Y$ the map $k(-, U)$ is measurable with respect to $\Sigma_{[0,1]}$, the subspace

σ -algebra of $\Sigma_{\mathbb{R}}$. Such a map provides a notion of sub-probability measure on the space Y parametrised by elements of X ; we write $k : X \rightsquigarrow Y$ when k is a kernel from X to Y .

Importantly, kernels can be composed: given $k : X \rightsquigarrow Y$ and $h : Y \rightsquigarrow Z$, their composition is the map $h \circ k : X \times \Sigma_Z \rightarrow [0, 1]$ defined as

$$(x, U) \mapsto \int_{y \in Y} h(y, U) k(x, dy).$$

This is still a sub-probability kernel, and the **Dirac kernel** $\delta : X \rightsquigarrow X$ (defined so that for every $x \in X$ the measure $\delta(x, -)$ on X is the Dirac measure at x) is the identity for composition of kernels.

7.1.2 Products

If X and Y are measurable spaces, recall from the previous chapter that the *product space* $X \times Y$ has σ -algebra generated by the “rectangles” $U \times V$, with $U \in \Sigma_X$ and $V \in \Sigma_Y$. If X and Y are equipped with sub-probability measures μ_X and μ_Y , respectively, the **product measure** $\mu_X \otimes \mu_Y$ on $X \times Y$ is uniquely determined by its value on rectangles:

$$(\mu_X \otimes \mu_Y)(U \times V) = \mu_X(U) \times \mu_Y(V).$$

Using this we can also define the product of kernels: if $k : X \rightsquigarrow Z$ and $h : Y \rightsquigarrow W$, then the kernel $k \otimes h : X \times Y \rightsquigarrow Z \times W$ is defined by

$$(k \otimes h)(x, y) = k(x, -) \otimes h(y, -)$$

for any $x \in X$, $y \in Y$.

7.1.3 Standard Borel spaces

To define our model of probabilistic strategies, we must ensure all measurable spaces are *standard Borel*.

Definition 7.1 ([Kal06]). A measurable space X is **standard Borel** if it is countable and discrete or measurably isomorphic to \mathbb{R}^n for some $n \in \mathbb{N}$.

The restriction is common in probability theory, and many standard theorems hold only for this restricted class. But it is sufficient for our purposes, and moreover the class of standard Borel spaces is closed under all the constructions of the model. We will need the following result:

Lemma 7.2. *Let X, Y, Z be standard Borel spaces, and let $f : Z \rightarrow X$ and $r : Y \rightarrow X$ be measurable functions. Consider the pullback*

$$\begin{array}{ccc} W & \xrightarrow{\Pi_2} & Z \\ \Pi_1 \downarrow & & \downarrow f \\ Y & \xrightarrow{X} & X \end{array}$$

of r along f . Then:

- For every $y \in Y$, $z \in Z$, and $U \in \Sigma_W$, the sections $U_y = \{z \in Z \mid (y, z) \in U\}$ and $U_z = \{y \in Y \mid (y, z) \in U\}$ are in Σ_Z and Σ_Y , respectively.
- If $k : X \times \Sigma_Y \rightarrow [0, 1]$ is a sub-probability kernel satisfying $k(x, Y \setminus r^{-1}\{x\}) = 0$ for every x , then the map $k^\# : Z \times \Sigma_W \rightarrow [0, 1]$ defined by $k^\#(z, U) = k(f(z), U_z)$ is a sub-probability kernel such that $k^\#(z, W \setminus (\Pi_2)^{-1}\{z\}) = 0$ for all $z \in Z$.

Proof. The proof of the first statement is standard. Let $k : X \times \Sigma_Y \rightarrow [0, 1]$ be a stochastic kernel, and let $z \in Z$. Then $k^\#(z, -)$ is a sub-probability measure, because $k(z, -)$ is countably additive and $(-)_z$ commutes with countable disjoint union. Now, for each $U \in \Sigma_W$, we must show that $k^\#(-, U) : W \rightarrow [0, 1]$ is measurable. For any U of the form $E_Y \times E_Z$, and for any $V \in \Sigma_{[0,1]}$, we have $k^\#(-, U)^{-1}V = \{z \in Z \mid k(f(z), U_z) \in V\} = \{z \in E_z \mid k(f(z), E_Y \cap r^{-1}\{f(z)\}) \in V\} \cup \{z \in Z \setminus E_z \mid k(f(z), \emptyset) \in V\}$ but by assumption $k(f(z), E_Y \cap r^{-1}\{f(z)\}) = k(f(z), E_Y)$ for any z , so we get $f^{-1}(k(-, E_Y)^{-1}V \cap E_Z) \cup f^{-1}(k(-, \emptyset)^{-1}V \setminus E_Z)$, a measurable set. So the set \mathcal{D} of $U \in \Sigma_W$ such that $k^\#(-, U)$ is measurable contains all generating elements. To show $\mathcal{D} = \Sigma_W$, by the λ - π theorem [Bil08] it is enough to show that \mathcal{D} is closed under complements and countable disjoint unions. This is easily checked using standard measure-theoretic arguments. \square

From now on we assume all measurable essps are standard Borel.

7.2 Probability in measurable event structures with symmetry

Throughout this section consider a fixed (standard Borel) measurable event structure with symmetry, $(\mathcal{E}, \mathcal{M})$, without polarity. The absence of polarity suggests that a single agent is responsible for all events in the process, and our goal is to make stochastic the behaviour of this agent. Indeed an alternative approach might be to make all events positive; this will be reflected by the various notions of valuations we consider, in which plain extensions \subseteq play the role of positive extensions.

7.2.1 Conditional valuations

Let us restate Definition 3.1 in this context:

Definition 7.3. A **(discrete) conditional valuation** on an event structure with symmetry \mathcal{E} is a family $(v(y \mid x))_{x \subseteq y}$ of coefficients in $[0, 1]$ satisfying

- (1) $v(x \mid x) = 1$ for all $x \in \mathcal{C}(E)$;
- (2) if $x \subseteq y \subseteq z$ then $v(z \mid x) = v(y \mid x)v(z \mid y)$;
- (3) if $\theta : x \cong_E y$ and $\theta \subseteq \theta' : x' \cong_E y'$, then $v(x' \mid x) = v(y' \mid y)$.

(4) if $x \subseteq y_1, \dots, y_n$, then

$$\sum_I (-1)^{|I|+1} v(\cup_{i \in I} y_i \mid x) \leq 1$$

with the sum ranging over the $I \subseteq \{1, \dots, n\}$ with $\cup_{i \in I} y_i \in \mathcal{C}(E)$.

For the measurable ess $(\mathcal{E}, \mathcal{M})$, we generalise the above by considering a family of kernels $k_{x,y} : \mathcal{M}(x) \rightsquigarrow \mathcal{M}(y)$ (labelled $k_{x,y}^\mathcal{E}$ when there is a risk of confusion), indexed by extensions $x \subseteq y$ in $\mathcal{C}(E)$.

The formal definition is as follows:

Definition 7.4. A **conditional valuation** on a measurable ess $(\mathcal{E}, \mathcal{M})$ consists of a sub-probability measure μ_\emptyset on $\mathcal{M}(\emptyset)$, and a family $\mathcal{K} = (k_{x,y})_{x \subseteq y \in \mathcal{C}(E)}$ of sub-probability kernels $k_{x,y} : \mathcal{M}(x) \rightsquigarrow \mathcal{M}(y)$ satisfying the following conditions:

- (Identity) $k_{x,x}(u, -) = \delta_u$ for every $u \in \mathcal{M}(x)$;
- (Composition) if $x \subseteq y \subseteq z$, then $k_{x,z} = k_{y,z} \circ k_{x,y}$;
- (Drop) if $x \subseteq y_1, \dots, y_n$ and $u \in \mathcal{M}(x)$, then

$$\sum_I (-1)^{|I|+1} k_{x, \cup_{i \in I} y_i}(u, \mathcal{M}(\cup_{i \in I} y_i)) \leq 1,$$

where I ranges over nonempty subsets of $\{1, \dots, n\}$ such that $\cup_{i \in I} y_i$ is consistent;

- (Concentration) for all $u \in \mathcal{M}(x)$, $k_{x,y}(u, \mathcal{M}(y) \setminus \mathcal{M}(x \subseteq y)^{-1}\{u\}) = 0$;
- (Symmetry) if $\theta : x \cong_S y$ and $\theta \subseteq \theta' : x' \cong_S y'$, then the kernels $k_{x,x'}$ and $k_{y,y'}$ are equal modulo the isos $\mathcal{M}(x) \cong \mathcal{M}(y)$ and $\mathcal{M}(x') \cong \mathcal{M}(y')$.

Consider for instance the (representable) measurable ess drawn as a measurable fibration below (previously used in Chapter 6):

$$\begin{array}{ccc} \begin{array}{c} \xleftarrow{-\infty} \xrightarrow{\infty} \\ \xleftarrow{-\infty} \xrightarrow{\infty} \end{array} & \xrightarrow{\quad} & \begin{array}{c} a_1 \\ \downarrow \\ a_2 \end{array} \end{array}$$

$r_1 \in \mathbb{R}$ (top interval), $r_2 \in \mathbb{R}$ (bottom interval)

Here the base event structure E has trivial symmetry, and the fibres are defined as $\mathcal{M}(\emptyset) = \{*\}$, $\mathcal{M}(\{a_1\}) = \mathbb{R}$, and $\mathcal{M}(\{a_1, a_2\}) = \mathbb{R} \times \mathbb{R}$, with the restriction map $\mathcal{M}(\{a_1\}) \hookrightarrow \mathcal{M}(\{a_1, a_2\})$ acting as the first projection. Then we could for instance define $k_{\emptyset, \{a_1\}}$ to be the uniform measure on $[0, 1]$ (extended to the reals by assigning measure 0 outside of the interval), and $k_{\{a_1\}, \{a_1, a_2\}}(r_1, -)$ to be a uniform distribution on $[r_1, r_1 + 1]$.

Given a conditional valuation on $(\mathcal{E}, \mathcal{M})$, we can define a measure μ_x on $\mathcal{M}(x)$ for each $x \in \mathcal{C}(E)$ as $\mu_x(U) = \int k_{\emptyset, x}(-, U) d\mu_\emptyset$.

An interesting question is the following: what properties are needed of a family $(\mu_x)_{x \in \mathcal{C}(E)}$ of measures for it to be induced by a conditional valuation in this way? An answer is given by the disintegration theorem [Kal06] of probability theory, but it is outside the scope of this thesis.

7.3 Probability in measurable strategies

7.3.1 Measurable race-freeness

In the discrete model, to make strategies probabilistic we required that the games be *race-free*, meaning that there should be no immediate conflict between moves of opposite polarity.

In measurable games we generalise this condition in a natural way using a pullback condition:

Definition 7.5. A measurable essp $(\mathcal{E}, \mathcal{M})$ is **race-free** if for every $x \in \mathcal{C}(E)$, if $x \subseteq^+ y$ and $x \subseteq^- z$, then $z \cup y \in \mathcal{C}(E)$ (*i.e.* \mathcal{E} is race-free in the usual sense), and moreover the diagram

$$\begin{array}{ccc} \mathcal{M}(y \cup z) & \xrightarrow{\mathcal{M}(y \leftarrow y \cup z)} & \mathcal{M}(y) \\ \mathcal{M}(z \leftarrow y \cup z) \downarrow & & \downarrow \mathcal{M}(x \leftarrow y) \\ \mathcal{M}(z) & \xrightarrow{\mathcal{M}(x \leftarrow z)} & \mathcal{M}(x) \end{array}$$

is a pullback in **Meas**.

It is the case also in measurable games that a strategy on a race-free game is necessarily race-free:

Lemma 7.6. *If $\underline{\mathcal{A}}$ is race-free and $\underline{\sigma} : \underline{\mathcal{S}} \rightarrow \underline{\mathcal{A}}$ is a measurable strategy, then $\underline{\mathcal{S}}$ is race-free.*

Proof. For $x \in \mathcal{C}(S)$, if $x \subseteq^+ y$ and $x \subseteq^- z$, then $z \cup y \in \mathcal{C}(S)$ by Lemma 3.15, and the cube below commutes (horizontal arrows are restrictions maps, and vertical maps are instances of nt^σ):

$$\begin{array}{ccccc} & & \mathcal{M}_{\mathcal{S}}(y \cup z) & \longrightarrow & \mathcal{M}_{\mathcal{S}}(y) \\ & & \downarrow & & \downarrow \\ & & \mathcal{M}_{\mathcal{A}}(\sigma y \cup \sigma z) & \longrightarrow & \mathcal{M}_{\mathcal{A}}(\sigma y) \\ & \swarrow & & \swarrow & \\ \mathcal{M}_{\mathcal{S}}(z) & \longrightarrow & \mathcal{M}_{\mathcal{S}}(x) & & \\ \downarrow & & \downarrow & & \downarrow \\ \mathcal{M}_{\mathcal{A}}(\sigma z) & \longrightarrow & \mathcal{M}_{\mathcal{A}}(\sigma x) & & \end{array}$$

The front and back faces are pullbacks by receptivity of $\underline{\sigma}$, and the bottom face is a pullback by race-freeness of $\underline{\mathcal{A}}$. So the top face is a pullback, *i.e.* $\underline{\mathcal{S}}$ is race-free. \square

7.3.2 Conditional valuations on measurable essps

Definition 7.7. A **conditional valuation** on a race-free measurable essp $(\mathcal{S}, \mathcal{M})$ consists of a measure μ_{\emptyset} on $\mathcal{M}(\emptyset)$, and a family $\mathcal{K} = (k_{x,y})_{x \subseteq^+ y \in \mathcal{C}(E)}$ of stochastic kernels $k_{x,y} : \mathcal{M}(x) \rightsquigarrow \mathcal{M}(y)$ satisfying the following conditions:

- (Identity) $k_{x,x}(u, -) = \delta_u$ for every $u \in \mathcal{M}(x)$;
- (Composition) if $x \sqsubseteq^+ y \sqsubseteq^+ z$, then $k_{x,z} = k_{y,z} \circ k_{x,y}$;
- (Drop) if $x \sqsubseteq^+ y_1, \dots, y_n$ and $u \in \mathcal{M}(x)$, then

$$\sum_I (-1)^{|I|+1} k_{x, \bigcup_{i \in I} y_i}(u, \mathcal{M}(\bigcup_{i \in I} y_i)) \leq 1,$$

where I ranges over nonempty subsets of $\{1, \dots, n\}$ such that $\bigcup_{i \in I} y_i$ is consistent;

- (Concentration) for all $u \in \mathcal{M}(x)$, $k_{x,y}(u, \mathcal{M}(y) \setminus \mathcal{M}(x \sqsubseteq y)^{-1}\{u\}) = 0$;
- (Symmetry) if $\theta : x \cong_S y$ and $\theta \sqsubseteq \theta' : x' \cong_S y'$, then $k_{x,x'} = k_{y,y'}$.
- (+/--Independence) if $x \sqsubseteq^+ y$, $x \sqsubseteq^- z$, then $k_{z,y \cup z}$ is the pullback-lifting of $k_{x,y}$ (with respect to the race-freeness pullback).

Definition 7.8. A **probabilistic** strategy on a measurable game $\underline{\mathcal{A}}$ is a measurable strategy $\underline{\sigma} : \underline{\mathcal{S}} \rightarrow \underline{\mathcal{A}}$, together with a conditional valuation on $\underline{\mathcal{S}}$.

We proceed to discuss composition of probabilistic measurable strategies. Because of the “conditional” approach, this is more involved than in the discrete case. Let us briefly discuss the technical steps.

Let $\underline{\sigma} : \underline{\mathcal{S}} \rightarrow \underline{\mathcal{A}} \perp \underline{\mathcal{B}}$ and $\underline{\tau} : \underline{\mathcal{T}} \rightarrow \underline{\mathcal{B}} \perp \underline{\mathcal{C}}$ probabilistic strategies with conditional valuations $\mathcal{K}^{\mathcal{S}}$ and $\mathcal{K}^{\mathcal{T}}$, respectively. A positive extension $x \sqsubseteq^+ y$ in $\mathcal{C}(T \odot S)$ is always induced from a *positive/internal* extension $[x] \sqsubseteq^{+,0} [y]$ in $\mathcal{C}(T \otimes S)$, meaning that all events in $[y] \setminus [x]$ are either internal or positive. (This is because τ and σ are courteous.) Such an extension can always be decomposed as a chain of extensions

$$[x] \sqsubseteq^{\sigma} u_0 \sqsubseteq^{\tau} u_1 \sqsubseteq^{\sigma} \dots \sqsubseteq^{\tau} [y]$$

where \sqsubseteq^{σ} (resp. \sqsubseteq^{τ}) means all added moves are σ -actions (resp. τ -actions). For each step in the chain, a kernel can be lifted appropriately from either $\mathcal{K}^{\mathcal{S}}$ or $\mathcal{K}^{\mathcal{T}}$. Using kernel composition, we obtain a kernel $\mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}([x]) \rightsquigarrow \mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}([y])$, *i.e.* $\mathcal{M}_{\mathcal{T} \odot \mathcal{S}}(x) \rightsquigarrow \mathcal{M}_{\mathcal{T} \odot \mathcal{S}}(y)$.

One difficulty is to show that this does not depend on the particular choice of chain $[x] \sqsubseteq \dots \sqsubseteq [y]$, and this is the aim of the discussion that follows. We will then see why the obtained family of kernels satisfies the axioms for a conditional valuation.

7.3.3 Interaction of probabilistic strategies

Consider the interaction $\underline{\mathcal{T}} \otimes \underline{\mathcal{S}}$ of $\underline{\sigma}$ and $\underline{\tau}$ as measurable strategies. We define an initial measure $\mu_{\emptyset}^{\mathcal{T} \otimes \mathcal{S}}$, and a family of kernels $k_{x,y}^{\mathcal{T} \otimes \mathcal{S}}$ indexed by positive/internal extensions $x \sqsubseteq^{+,0} y$ in $\mathcal{C}(T \otimes S)$.

The initial measure. For this we argue that $\mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(\emptyset)$ can be seen as a subspace of $\mathcal{M}_{\mathcal{S}}(\emptyset) \times \mathcal{M}_{\mathcal{T}}(\emptyset)$ via the canonical map $h : \mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(\emptyset) \rightarrow \mathcal{M}_{\mathcal{S}}(\emptyset) \times \mathcal{M}_{\mathcal{T}}(\emptyset)$. By construction, the maps $\text{nt}^{\Pi_1} : \mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(\emptyset) \rightarrow \mathcal{M}_{\mathcal{S}}(\emptyset) \times \mathcal{M}_{\mathcal{C}}(\emptyset)$ and $\text{nt}_{\emptyset}^{\Pi_2} : \mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(\emptyset) \rightarrow \mathcal{M}_{\mathcal{A}}(\emptyset) \times \mathcal{M}_{\mathcal{T}}(\emptyset)$ generate the σ -algebra on $\mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(\emptyset)$, so in other words the following set provides a basis for it:

$$\{(\text{nt}_{\emptyset}^{\Pi_1})^{-1}(U_S \times U_C) \cap (\text{nt}_{\emptyset}^{\Pi_2})^{-1}(U_A \times U_T) \mid U_S \in \Sigma_{\mathcal{M}_{\mathcal{S}}(\emptyset)}, \text{ etc. } \}.$$

And indeed every element of the above set is equal to $h^{-1}(U_S \times U_T)$.

Accordingly, a measure $\mu_{\emptyset}^{T \otimes S}$ on $\mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(\emptyset)$ is obtained using the standard *subspace measure* construction on the product measure $\mu_{\emptyset}^S \otimes \mu_{\emptyset}^T$:

$$\mu_{\emptyset}^{T \otimes S}(U) := \inf\{(\mu_{\emptyset}^S \otimes \mu_{\emptyset}^T)(U') \mid U' \in \Sigma_{\mathcal{M}_{\mathcal{S}}(\emptyset) \times \mathcal{M}_{\mathcal{T}}(\emptyset)} \text{ and } h(U) \subseteq U'\}.$$

This is a sub-probability measure because μ_{\emptyset}^S and μ_{\emptyset}^T are. But note that when μ_{\emptyset}^S and μ_{\emptyset}^T are strict probability measures, $\mu_{\emptyset}^{T \otimes S}$ may not be: the strategies $\underline{\sigma}$ and $\underline{\tau}$ are not guaranteed to agree on the choice of initial state. (This is the case only when h above is an isomorphism.)

We have defined a measure on $\mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(\emptyset)$, the “space of initial states”. In what follows we deal with the subsequent steps.

Lifting kernels from \mathcal{K}^S and \mathcal{K}^T . We are only concerned with positive or neutral extensions, since negative ones do not carry any probabilistic information.

First, consider an extension by σ -actions. If $x \subseteq^{\sigma} y$ in $\mathcal{C}(T \circledast S)$, recall from Lemma 6.19 that the following is a pullback

$$\begin{array}{ccc} \mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(y) & \xrightarrow{\mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(x \hookrightarrow y)} & \mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(x) \\ \text{nt}_y^{\Pi_1} \downarrow & & \downarrow \text{nt}_x^{\Pi_1} \\ \mathcal{M}_{\mathcal{S} \parallel \mathcal{C}}(y_S \parallel y_C) & \longrightarrow & \mathcal{M}_{\mathcal{S} \parallel \mathcal{C}}(x_S \parallel x_C) \end{array}$$

and because $y_C = x_C$, the diagram below is also a pullback:

$$\begin{array}{ccc} \mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(y) & \xrightarrow{\mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(x \hookrightarrow y)} & \mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(x) \\ \pi_1 \circ \text{nt}_y^{\Pi_1} \downarrow & & \downarrow \pi_1 \circ \text{nt}_x^{\Pi_1} \\ \mathcal{M}_{\mathcal{S}}(y_S) & \xrightarrow{\mathcal{M}(x_S \hookrightarrow y_S)} & \mathcal{M}_{\mathcal{S}}(x_S) \end{array}$$

Using Lemma 7.2, the kernel $k_{x_S, y_S}^S : \mathcal{M}_{\mathcal{S}}(x_S) \rightsquigarrow \mathcal{M}_{\mathcal{S}}(y_S)$ can be lifted through to a kernel $k_{x, y}^{T \otimes S} : \mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(x) \rightsquigarrow \mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(y)$ satisfying the concentration property.

This can be done similarly whenever $x \subseteq^{\tau} y$. More generally, for an extension $x \subseteq^{+,0} y$, events of $y \setminus x$ are either σ -actions or τ -actions, so every covering chain from x to y is of the form $x \text{---} \lambda_1 u_1 \text{---} \lambda_2 \dots \text{---} \lambda_n u_n \text{---} \lambda_{n+1} y$, with $\lambda_i \in \{\sigma, \tau\}$ for each i . With respect to this chain we can define a kernel $k_{p, u_1, \dots, u_n, q}^{T \otimes S} = k_{u_n, y}^{T \otimes S} \circ \dots \circ k_{x, u_1}^{T \otimes S}$. If the extension is trivial ($x = y$) we assign it the identity (Dirac) kernel.

In fact, we will see that the properties of the valuations \mathcal{K}^S and \mathcal{K}^T ensure that this kernel does not depend on the choice of covering chain. We first prove two auxiliary lemmas.

Lemma 7.9. *Let $\lambda \in \{\sigma, \tau\}$ and suppose $x \prec^\lambda u \prec^\lambda y$ in $\mathcal{C}(T \otimes S)$. Then $k_{x,y}^{T \otimes S} = k_{u,y}^{T \otimes S} \circ k_{x,u}^{T \otimes S}$.*

Proof. It is enough to show that the lifting $(-)^{\#}$ of Lemma 7.2 preserves composition. It is a straightforward verification. \square

Lemma 7.10. *If $x \prec^\sigma y$ and $x \prec^\tau z$ in $\mathcal{C}(T \otimes S)$, then $z \cup y \in \mathcal{C}(T \otimes S)$ and $k_{y,z \cup y}^{T \otimes S} \circ k_{x,y}^{T \otimes S} = k_{z,z \cup y}^{T \otimes S} \circ k_{x,z}^{T \otimes S}$.*

Proof. Using that \mathcal{S} and \mathcal{T} are race-free,

$$\begin{array}{ccc} \mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(y \cup z) & \longrightarrow & \mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(y) \\ \downarrow & & \downarrow \\ \mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(z) & \longrightarrow & \mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(x) \end{array}$$

is a pullback in **Meas**. Then, for $u \in \mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(x)$, we check that the sub-probability measures $k_{y,z \cup y}^{T \otimes S} \circ k_{x,y}^{T \otimes S}(u, -)$ and $k_{z,z \cup y}^{T \otimes S} \circ k_{x,z}^{T \otimes S}(u, -)$ on $\mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(z \cup y)$ are the same. By the concentration property, it is sufficient that they agree on measurable sets $U \in \Sigma_{\mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(y \cup z)}$ such that $U \subseteq \mathcal{M}(x \hookrightarrow y \cup z)^{-1}\{u\}$, and an inspection of the pullback above shows that $\mathcal{M}(x \hookrightarrow y \cup z)^{-1}\{u\}$, viewed as a subspace of $\mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(z \cup y)$, is isomorphic to the product $\mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(x \hookrightarrow z)^{-1}\{u\} \times \mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(x \hookrightarrow y)^{-1}\{u\}$.

Moreover, using that the diagrams

$$\begin{array}{ccc} \mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(z) & \longrightarrow & \mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(x) & & \mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(y) & \longrightarrow & \mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(x) \\ \downarrow & & \downarrow & & \downarrow & & \downarrow \\ \mathcal{M}_{\mathcal{T}}(z_T) & \longrightarrow & \mathcal{M}_{\mathcal{T}}(x_T) & & \mathcal{M}_{\mathcal{S}}(y_S) & \longrightarrow & \mathcal{M}_{\mathcal{S}}(x_S) \end{array}$$

are pullbacks, we see that

$$\mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(x \hookrightarrow z)^{-1}\{u\} \cong \mathcal{M}_{\mathcal{T}}(x_T \hookrightarrow z_T)^{-1}\{u_T\}$$

and

$$\mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(x \hookrightarrow y)^{-1}\{u\} \cong \mathcal{M}_{\mathcal{S}}(x_S \hookrightarrow z_S)^{-1}\{u_S\}.$$

So we let V be of the form $E \times E'$ for measurable sets $E \subseteq \mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(x \hookrightarrow y)^{-1}\{u\}$ and $E' \subseteq \mathcal{M}_{\mathcal{S}}(x_S \hookrightarrow z_S)^{-1}\{u_S\}$. Then, we have

$$\begin{aligned} & k_{y,z \cup y}^{T \otimes S} \circ k_{x,y}^{T \otimes S}(u, V) \\ &= \int_{u' \in \mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(y)} k_{y,z \cup y}^{T \otimes S}(u', V) k_{x,y}^{T \otimes S}(u, du') \\ &= \int_{u' \in \mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(x \hookrightarrow y)^{-1}\{u\}} k_{y_T, (z \cup y)_T}^T(u'_T, V_y) k_{x_S, y_S}^S(u_S, du'_S) \\ &= \int_{w \in E'} k_{x_T, z_T}^T(\mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(x \hookrightarrow y)(w), E) k_{x_S, y_S}^S(u_S, dw) \\ &= \int_{w \in E'} k_{x_T, z_T}^T(u_T, E) k_{x_S, y_S}^S(u_S, dw) \\ &= k_{x_S, y_S}^S(u_S, E') \times k_{x_T, z_T}^T(Z_T, E') \end{aligned}$$

and a symmetric calculation shows that $k_{z,z \cup y}^{T \otimes S} \circ k_{x,z}^{T \otimes S}(u, V)$ has the same value. \square

We are now in a position to show that any two parallel chains in $\mathcal{T} \circledast \mathcal{S}$ yield the same composite kernel:

Lemma 7.11. *If $x \subseteq^{+,0} y \in \mathcal{C}(T \circledast S)$ and we have two chains*

$$\begin{aligned} x \text{---} \text{c}^{\lambda_1} z_1 \text{---} \text{c}^{\lambda_2} \dots \text{---} \text{c}^{\lambda_{n-1}} z_{n-1} \text{---} \text{c}^{\lambda_n} y \\ x \text{---} \text{c}^{\rho_1} z'_1 \text{---} \text{c}^{\rho_2} \dots \text{---} \text{c}^{\rho_{n-1}} z'_{n-1} \text{---} \text{c}^{\rho_n} y \end{aligned}$$

where $\lambda_i, \rho_i \in \{\sigma, \tau\}$ for each i , then $k_{x, z_1, \dots, z_n, y}^{T \circledast S} = k_{x, z'_1, \dots, z'_n, y}^{T \circledast S}$. Thus we may write $k_{x, y}^{T \circledast S}$ for the kernel obtained via any chain.

Proof. By induction on n . If $n = 0$, the result holds directly since there is only one possible chain from x to y . If $n > 0$, consider $w = z_1 \cup z'_1$. By the induction hypothesis, any chain from z_1 to y yields the same kernel, so in particular

$$k_{z_1, z_2, \dots, z_n, y}^{T \circledast S} = k_{w, y}^{T \circledast S} \circ k_{z_1, w}^{T \circledast S}$$

and similarly we have

$$k_{z'_1, z'_2, \dots, z'_n, y}^{T \circledast S} = k_{w, y}^{T \circledast S} \circ k_{z'_1, w}^{T \circledast S}.$$

Next, observe that $x \text{---} \text{c}^{\lambda_1} z_1 \text{---} \text{c}^{\rho_1} w$ and $x \text{---} \text{c}^{\rho_1} z'_1 \text{---} \text{c}^{\lambda_1} w$. If $\lambda_1 = \rho_1$, then it follows from Lemma 7.9 that $k_{z_1, w}^{T \circledast S} \circ k_{x, z_1}^{T \circledast S} = k_{z'_1, w}^{T \circledast S} \circ k_{x, z'_1}^{T \circledast S}$, since both are equal to $k_{x, w}^{T \circledast S}$. If instead $\lambda_1 \neq \rho_1$, then Lemma 7.10 shows that the same equality holds. Thus

$$\begin{aligned} k_{x, z_1, \dots, z_n, y}^{T \circledast S} &= k_{z_1, \dots, z_n, y}^{T \circledast S} \circ k_{x, z_1}^{T \circledast S} \\ &= k_{w, y}^{T \circledast S} \circ k_{z_1, w}^{T \circledast S} \circ k_{x, z_1}^{T \circledast S} \\ &= k_{w, y}^{T \circledast S} \circ k_{z'_1, w}^{T \circledast S} \circ k_{x, z'_1}^{T \circledast S} \\ &= k_{w, y}^{T \circledast S} \circ k_{z'_1, w}^{T \circledast S} \circ k_{x, z'_1}^{T \circledast S} \\ &= k_{z'_1, \dots, z'_n, y}^{T \circledast S} \circ k_{x, z'_1}^{T \circledast S} \\ &= k_{x, z'_1, \dots, z'_n, y}^{T \circledast S}. \end{aligned}$$

□

Following the above process, we obtain a family $\mathcal{K}^{T \circledast S} = (k_{x, y}^{T \circledast S})_{x \subseteq^{+,0} y}$ of kernels, indexed by the positive/neutral extensions in $\mathcal{C}(T \circledast S)$. We proceed to show that this satisfies a number of properties, ensuring that the *hiding* step turns $\mathcal{K}^{T \circledast S}$ into a valid conditional valuation.

Symmetry. Conditional valuations must assign the same kernel to *symmetric* extensions. It is straightforward to see why this property holds in the interaction.

Lemma 7.12. *If $\theta : x \cong_{T \circledast S} y$ and $\theta \subseteq \theta' : x' \cong_{T \circledast S} y'$, then $k_{x, x'}^{T \circledast S} = k_{y, y'}^{T \circledast S}$.*

Proof. We first show the result holds for a one-step extension $\theta \text{---} \text{c}^\sigma \theta'$. By definition, the kernel $k_{x, x'}^{T \circledast S}$ is lifted from $k_{x_S, x'_S}^S : \mathcal{M}_S(x_S) \rightsquigarrow \mathcal{M}_S(x'_S)$, and similarly $k_{y, y'}^{T \circledast S}$ is lifted from k_{y_S, y'_S}^S . But since the extension projects down to an extension $\theta_S \text{---} \text{c}^+ \theta'_S$

For the purposes of the upcoming development it will be convenient to extend the set $\mathcal{C}(T \otimes S)$ to a lattice $\mathcal{C}(T \otimes S)^\top$, obtained by adding to $\mathcal{C}(T \otimes S)$ a top element \top . That is, whenever $x, y \in \mathcal{C}(T \otimes S)$ and $x \cup y \notin \mathcal{C}(T \otimes S)$, we set $x \vee y = \top$. (If x and y are compatible we have $x \vee y = x \cup y$.) The drop function is also extended: assume $x \neq \top$ and $x \subseteq y_1, \dots, y_n \in \mathcal{C}(T \otimes S)^\top$. Suppose $y_1, \dots, y_k \neq \top$ and $y_{k+1}, \dots, y_n = \top$ (without loss of generality, since the definition of $\text{dr}[x; y_1, \dots, y_n]$ is insensitive to a permutation of the y_i). Then we define

$$\text{dr}[x; y_1, \dots, y_n] := \text{dr}[x; y_1, \dots, y_k].$$

We show two technical lemmas, before carrying out the first step in the plan above. (The integrals are well-defined by Lemma 7.13.)

Lemma 7.14. *For every $x \subseteq^{+,0} y_1, \dots, y_n \in \mathcal{C}(T \otimes S)^\top$,*

$$\begin{aligned} \text{dr}[x; y_1, \dots, y_n](u) &= \text{dr}[x; y_1, \dots, y_{n-1}](u) \\ &\quad - \int_{u_n \in \mathcal{M}(y_n)} \text{dr}[y_n; y_1 \vee y_n, \dots, y_{n-1} \vee y_n](u_n) k_{x, y_n}(u, du_n). \end{aligned}$$

Proof. The integral in the RHS can be rewritten as

$$\int_{u_n \in \mathcal{M}(y_n)} \left[1 - \sum_I (-1)^{|I|+1} k_{y_n, y_n \cup y_I}(u_n, \mathcal{M}(y_n \cup y_I)) \right] k_{x, y_n}(u, du_n) \quad (7.1)$$

where we write y_I for $\bigcup_{i \in I} y_i$, and I ranges over nonempty subsets of $\{1, \dots, n-1\}$ such that $y_I \cup y_n$ is consistent. Since $k_{x, y_I \cup y_n} = k_{y_n, y_I \cup y_n} \circ k_{x, y_n}$, by definition of kernel composition we have

$$\int_{u_n \in \mathcal{M}(y_n)} k_{y_n, y_n \cup y_I}(u_n, \mathcal{M}(y_n \cup y_I)) k_{x, y_n}(u, du_n) = k_{x, y_I \cup y_n}(u, \mathcal{M}(y_n \cup y_I)).$$

So 7.1 above is equal to

$$k_{x, y_n}(u, \mathcal{M}(y_n)) - \sum_I (-1)^{|I|+1} k_{x, y_I \cup y_n}(u, \mathcal{M}(y_n \cup y_I)).$$

Looking back at the main statement, we have

$$\begin{aligned} \text{RHS} &= \text{dr}[x; y_1, \dots, y_{n-1}](u) \\ &\quad - k_{x, y_n}(u, \mathcal{M}(y_n)) + \sum_I (-1)^{|I|+1} k_{x, y_I \cup y_n}(u, \mathcal{M}(y_n \cup y_I)) \end{aligned}$$

and by substituting the first term, and combining the sums, we recover the expression for $\text{dr}[x; y_1, \dots, y_n](u)$. \square

From this we derive another technical result:

Lemma 7.15. For every $x \subseteq^{+,0} y_1, \dots, y_n \in \mathcal{C}(T \otimes S)^\top$, $y_n \subseteq^{+,0} y'_n$, and $u \in \mathcal{M}(x)$

$$\begin{aligned} \text{dr}[x; y_1, \dots, y_{n-1}, y'_n](u) &= \text{dr}[x; y_1, \dots, y_n](u) \\ &+ \int_{u_n \in \mathcal{M}(y_n)} \text{dr}[y_n; y_1 \vee y_n, \dots, y_{n-1} \vee y_n, y'_n](u_n) \text{dk}_{x, y_n}(u, u_n). \end{aligned}$$

Proof. Applying Lemma 7.14 to both terms and cancelling out, we obtain:

$$\begin{aligned} \text{RHS} &= \text{dr}[x; y_1, \dots, y_{n-1}](u) \\ &- \int_{u_n \in \mathcal{M}(y_n)} \int_{u'_n \in \mathcal{M}(y'_n)} \text{dr}[y'_n; y_1 \vee y'_n, \dots, y_{n-1} \vee y'_n](u'_n) \text{dk}_{y_n, y'_n}(u_n, u'_n) \text{dk}_{x, y_n}(u, u_n). \end{aligned}$$

The double integral can be simplified using the definition of kernel composition. By another application of Lemma 7.14 we then that this is equal to the LHS. \square

Using the above, we show that for the drop condition to hold in $\mathcal{T} \otimes \mathcal{S}$, it suffices to check that the property holds for *one-step* extensions:

Proposition 7.16. *The following are equivalent:*

- (1) For $x \subseteq^{+,0} y_1, \dots, y_n \in \mathcal{C}(T \otimes S)$ and $u \in \mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(x)$, $\text{dr}[x; y_1, \dots, y_n](u) \geq 0$.
- (2) For $x \subset^{+,0} y_1, \dots, y_n \in \mathcal{C}(T \otimes S)$ and $u \in \mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(x)$, $\text{dr}[x; y_1, \dots, y_n](u) \geq 0$.

Proof. The (1) \Rightarrow (2) direction is clear. Assume (2) holds. We show the property holds for every extension $x \subseteq^{+,0} y_1, \dots, y_n \in \mathcal{C}(T \otimes S)$ and $u \in \mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(x)$, by induction on the *weight* of the extension, defined as $\prod_{i=1}^n |y_i \setminus x|$. If the weight is 0, then $x = y_i$ for some i . Without loss of generality, assume $i = n$. Then, the expression for $\text{dr}[x; y_1, \dots, y_n](u)$ given by Lemma 7.14 simplifies to 0 via a routine manipulation which we omit.

In the general case, either $x \subset y_i$ for each i , in which case (2) applies directly, or there is $x \subsetneq y'_i \subsetneq y_i$ for some i and we are in the situation of Lemma 7.15: $\text{dr}[x; y_1, \dots, y_n](u)$ is equal to

$$\text{dr}[x; y_1, \dots, y'_i, \dots, y_n](u) + \int_{u'_i \in \mathcal{M}(y'_i)} \text{dr}[y'_i; y_1 \vee y'_i, \dots, y_{n-1} \vee y'_i, y_i](u'_i) \text{dk}_{x, y'_i}(u, u'_i).$$

Observe that all extensions in the above expression have weight strictly lower than that of $x \subseteq y_1, \dots, y_n$. By the induction hypothesis, these are all non-negative, so that $\text{dr}[x; y_1, \dots, y_n] \geq 0$. \square

This completes the first step in our proof that the “drop condition” holds in the interaction $\mathcal{T} \otimes \mathcal{S}$. By Proposition 7.16, it is enough to consider one-step extensions of the form $x \subset^{+,0} y_1, \dots, y_n$. Recall that positive and neutral events in $T \otimes S$ are either “ σ -actions” or “ τ -actions”, so the y_i can be partitioned into two groups, accordingly. We will see that the drop $\text{dr}[x; y_1, \dots, y_n]$ reduces to a product of a drop in σ and one in τ . For $u \in \mathcal{M}(x)$, we write u_S for its image under the map

$$\mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(x) \xrightarrow{\text{nt}\pi_1} \mathcal{M}_{S \parallel C}(x_S \parallel x_C) \xrightarrow{\pi_1} \mathcal{M}_S(x_S),$$

and we define $u_T \in \mathcal{M}_{\mathcal{T}}(x_T)$ analogously.

Lemma 7.17. *Let $x \prec^\sigma y_1, \dots, y_k$ and $x \prec^\tau y_{k+1}, \dots, y_n$ in $\mathcal{C}(T \otimes S)$. Then, for all $u \in \mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(x)$,*

$$\text{dr}[x; y_1, \dots, y_n](u) = \text{dr}[x_S; (y_1)_S, \dots, (y_k)_S](u_S) \times \text{dr}[x_T; (y_{k+1})_T, \dots, (y_n)_T](u_T).$$

Proof. Let $I \subseteq \{1, \dots, n\}$ be such that $y_I = \bigcup_{i \in I} y_i$ is consistent. Write $I_S = \{i \in I \mid i \leq k\}$ and $I_T = \{i \in I \mid i > k\}$. Then $x \preceq^\sigma \cup_{i \in I_S} y_i$ and $x \preceq^\tau \cup_{i \in I_T} y_i$. As a shorthand, y_{I_S} and y_{I_T} denote $\cup_{i \in I_S} (y_i)_S \in \mathcal{C}(S)$ and $\cup_{i \in I_T} (y_i)_T \in \mathcal{C}(T)$, respectively.

By the same argument as in the proof of Lemma 7.10 we derive that

$$k_{x, y_I}^{\mathcal{T} \otimes \mathcal{S}}(u, \mathcal{M}(y_I)) = k_{x_S, y_{I_S}}^{\mathcal{S}}(u_S, \mathcal{M}_{\mathcal{S}}(y_{I_S})) \times k_{x_T, y_{I_T}}^{\mathcal{T}}(u_T, \mathcal{M}_{\mathcal{T}}(y_{I_T})). \quad (7.2)$$

Then, calculating from the RHS, we have

$$\begin{aligned} & \text{dr}[x_S; (y_1)_S, \dots, (y_k)_S](u_S) \times \text{dr}[x_T; (y_{k+1})_T, \dots, (y_n)_T](u_T) \\ &= \left(\sum_{I_S} (-1)^{|I_S|} k_{x_S, y_{I_S}}^{\mathcal{S}}(u_S, \mathcal{M}_{\mathcal{S}}(y_{I_S})) \right) \times \left(\sum_{I_T} (-1)^{|I_T|} k_{x_T, y_{I_T}}^{\mathcal{T}}(u_T, \mathcal{M}_{\mathcal{T}}(y_{I_T})) \right) \end{aligned}$$

where I_S ranges over (possibly empty) subsets of $\{1, \dots, k\}$ such that $\cup_{i \in I_S} y_i$ is consistent, and I_T similarly over subsets of $\{k+1, \dots, n\}$, and we take $y_\emptyset = x$. Now, pairs (I_S, I_T) correspond to subsets $I \subseteq \{1, \dots, n\}$ such that $\bigcup_{i \in I} y_i$ is consistent, since by race-freeness there are no conflicts between σ - and τ -actions. Thus, by 7.2, the expression above is equal to $\text{dr}[x; y_1, \dots, y_n]$. \square

We arrive at our main result:

Corollary 7.18. *Let $x \prec^{+,0} y_1, \dots, y_n$ in $\mathcal{C}(T \otimes S)$. Then, for all $u \in \mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}(x)$,*

$$\text{dr}[x; y_1, \dots, y_n](u) \geq 0.$$

Proof. Direct consequence of the previous lemma and the drop condition for σ and τ . \square

7.3.5 Composition of probabilistic strategies

We have studied in detail the probabilistic interaction of two measurable strategies $\underline{\sigma} : \underline{\mathcal{S}} \rightarrow \underline{\mathcal{A}}^\perp \parallel \underline{\mathcal{B}}$ and $\underline{\tau} : \underline{\mathcal{T}} \rightarrow \underline{\mathcal{B}}^\perp \parallel \underline{\mathcal{C}}$, equipped with conditional valuations $\mathcal{K}^{\mathcal{S}}$ and $\mathcal{K}^{\mathcal{T}}$, respectively. We have defined a family of kernels $\mathcal{K}^{\mathcal{T} \otimes \mathcal{S}} = (k_{x,y}^{\mathcal{T} \otimes \mathcal{S}})_{x \preceq^{+,0} y}$ indexed by positive/neutral extensions in $\mathcal{C}(T \otimes S)$.

We turn to the *composition* of $\underline{\sigma}$ and $\underline{\tau}$, the measurable strategy $\underline{\tau} \odot \underline{\sigma} : \underline{\mathcal{T}} \odot \underline{\mathcal{S}} \rightarrow \underline{\mathcal{A}}^\perp \parallel \underline{\mathcal{C}}$, and equip it with a conditional valuation $\mathcal{K}^{\mathcal{T} \odot \mathcal{S}}$.

It is key to observe (as we have already in 7.3.2) that whenever $x, y \in \mathcal{C}(T \odot S)$ satisfy $x \preceq^+ y$, their interaction witnesses satisfy $[x] \preceq^{+,0} [y]$. We define $k_{x,y}^{\mathcal{T} \odot \mathcal{S}} = k_{[x],[y]}^{\mathcal{T} \otimes \mathcal{S}}$, since by definition $\mathcal{M}_{\mathcal{T} \odot \mathcal{S}}(x) = \mathcal{M}_{\mathcal{T} \otimes \mathcal{S}}([x])$ (and the same for y). The results we have about $\mathcal{K}^{\mathcal{T} \otimes \mathcal{S}}$ suffice to show that this defines a conditional valuation.

Theorem 7.19. *The family $\mathcal{K}^{\mathcal{T} \odot \mathcal{S}} = (k_{x,y}^{\mathcal{T} \odot \mathcal{S}})_{x \preceq^+ y}$ is a conditional valuation, making $\underline{\tau} \odot \underline{\sigma}$ a probabilistic strategy.*

Proof. The (Identity) and (Composition) axioms hold immediately. The (Drop) axiom is a consequence of Corollary 7.18. The (Concentration) axiom holds because the concentration property is preserved by kernel composition and lifting (Lemma 7.2). The (Symmetry) axiom holds by Lemma 7.12, and by definition on the symmetry in $\mathcal{T} \odot \mathcal{S}$. \square

7.3.6 Probabilistic copycat

The identity strategy in this setting is a probabilistic version of the “measurable copycat” introduced in 6.3.1. This was defined for every measurable game \underline{A} as $\mathbb{C}_{\underline{A}} \rightarrow \mathbb{C}_{\underline{A}} \rightarrow \underline{A}^\perp \parallel \underline{A}$, and as we will see, the valuation $\mathcal{K}^{\mathbb{C}_{\underline{A}}}$ is obtained in a canonical way, because (whenever \underline{A} is race-free) measurable copycat satisfies a kind of *determinism* property.

In the probabilistic model of Chapter 3, the race-freeness property for arenas ensures that copycat is deterministic, meaning that Player never has to make a choice between two moves in conflict. Consequently all configurations can be assigned probability 1. The *measurable* copycat strategy satisfies an additional property:

Lemma 7.20. *Let \underline{A} be a race-free measurable game. Then, if $w \sqsubseteq^+ w' \in \mathcal{C}(\mathbb{C}_{\underline{A}})$, the map $\mathcal{M}_{\mathbb{C}_{\underline{A}}}(w \hookrightarrow w') : \mathcal{M}_{\mathbb{C}_{\underline{A}}}(w') \rightarrow \mathcal{M}_{\mathbb{C}_{\underline{A}}}(w)$ is an isomorphism.*

Proof. Recall that configurations of $\mathbb{C}_{\underline{A}}$ are of the form $x \parallel y$, where $x, y \in \mathcal{C}(\underline{A})$ and $x \sqsubseteq_A y$, i.e. $x \sqsupseteq_A^- x \cap y \sqsubseteq_A^+ y$.

Consider a one-step positive extension in $\mathcal{C}(\mathbb{C}_{\underline{A}})$. (It suffices to show the result in the one-step case, since isomorphisms compose.) Without loss of generality we can take this to be of the form $x \parallel y \dashv\vdash^+ x' \parallel y$, because the case $x \parallel y \dashv\vdash^+ x \parallel y'$ is symmetric.

By definition, the map $\mathcal{M}_{\mathbb{C}_{\underline{A}}}(x \parallel y \dashv\vdash^+ x' \parallel y)$ is the dashed mediating map below:

$$\begin{array}{ccccc}
 \mathcal{M}_{\mathbb{C}_{\underline{A}}}(x' \parallel y) & \xrightarrow{\quad} & \mathcal{M}_{\underline{A}}(y) & & \\
 \downarrow & \dashrightarrow & \downarrow & \searrow & \\
 \mathcal{M}_{\underline{A}}(x') & \xrightarrow{\quad} & \mathcal{M}_{\underline{A}}(y \cap x') & & \\
 & \searrow & \downarrow & \searrow & \\
 & & \mathcal{M}_{\mathbb{C}_{\underline{A}}}(x \parallel y) & \xrightarrow{\quad} & \mathcal{M}_{\underline{A}}(y) \\
 & \searrow & \downarrow & \searrow & \downarrow \\
 & & \mathcal{M}_{\underline{A}}(x) & \xrightarrow{\quad} & \mathcal{M}_{\underline{A}}(y \cap x)
 \end{array}$$

The bottom face is a pullback by race-freeness of \underline{A} , since $y \cup x \sqsubseteq^+ y$ and $y \cap x \sqsubseteq^- y \cap x'$ (because $x \sqsubseteq_{\underline{A}^\perp}^+ x'$). The front and back faces are pullbacks by definition, so the top face is a pullback.

Since pullbacks of identity maps are isomorphisms, we are done. \square

This means that if $w \sqsubseteq^+ w'$ in $\mathcal{C}(\mathbb{C}_{\underline{A}})$, each $u \in \mathcal{M}_{\mathbb{C}_{\underline{A}}}(w)$ has a unique *extension* in $\mathcal{M}_{\mathbb{C}_{\underline{A}}}(w')$. We define the kernel

$$k^{\mathbb{C}_{\underline{A}}} : \mathcal{M}_{\mathbb{C}_{\underline{A}}}(w) \rightsquigarrow \mathcal{M}_{\mathbb{C}_{\underline{A}}}(w')$$

as the identity kernel, modulo the iso.

7.4 The bicategory of probabilistic measurable strategies

We give a proof that measurable games and probabilistic strategies form a cartesian closed bicategory **PMG**, and we show that the bicategory **PG** of Chapter 3 arises as a sub-bicategory. To do this we follow the same principles as in Chapters 2 and 6, and first construct a pseudo-double category.

7.4.1 The pseudo-double category \mathcal{PMG}

We start by defining the 2-cells.

Definition 7.21. A map of probabilistic measurable strategies from $\underline{\sigma} : \underline{\mathcal{S}} \rightarrow \underline{\mathcal{A}}^\perp \parallel \underline{\mathcal{B}}$ to $\underline{\tau} : \underline{\mathcal{T}} \rightarrow \underline{\mathcal{C}}^\perp \parallel \underline{\mathcal{D}}$ is a map $(\underline{f}, \underline{g}, \underline{h}) : \underline{\sigma} \Rightarrow \underline{\tau}$ in \mathcal{MG} , such that:

- (1) for every $x \in \mathcal{C}(S)$, the map $\text{nt}_x^f : \mathcal{M}_S(x) \rightarrow \mathcal{M}_T(fx)$ is an isomorphism;
- (2) for every $x \subseteq y \in \mathcal{C}(S)$, we have $k_{x,y}^S \leq k_{fx,fy}^T$ modulo the isos, that is, for every $u \in \mathcal{M}_S(x)$ and $U \in \Sigma_{\mathcal{M}_S(y)}$,

$$k_{x,y}^S(u, U) \leq k_{fx,fy}^T(\text{nt}_x^f(u), \text{nt}_y^f U).$$

Remark. This definition can be given in greater generality, removing the requirement (1) and rephrasing (2) as follows: for every $u \in \mathcal{M}_S(x)$ and $V \in \Sigma_{\mathcal{M}_T(fy)}$,

$$k_{x,y}^S(u, (\text{nt}_y^f)^{-1}V) \leq k_{fx,fy}^T(\text{nt}_x^f(u), V).$$

We conjecture that the forthcoming development would support the extra generality; however the proofs would likely be more technical. The definition as stated is sufficient for our purposes: it generalises the 2-cells in the discrete setting of Chapter 3, and the resulting bicategory is cartesian closed.

Lemma 7.22. *Maps of probabilistic strategies are stable under vertical and horizontal composition.*

Proof. Vertical composition of maps poses no problem, as axioms (1) and (2) are directly seen to be stable under composition.

For horizontal composition, consider maps of strategies $\underline{f} : \underline{\mathcal{S}} \rightarrow \underline{\mathcal{S}}'$ and $\underline{g} : \underline{\mathcal{T}} \rightarrow \underline{\mathcal{T}}'$ and the induced map $\underline{g} \circledast \underline{f} : \underline{\mathcal{T}} \circledast \underline{\mathcal{S}} \rightarrow \underline{\mathcal{T}}' \circledast \underline{\mathcal{S}}'$. Note that (1) holds directly by Lemma 6.19, since pullbacks of isomorphisms are isomorphisms. If $x \subseteq^\sigma y \in \mathcal{C}(T \circledast S)$, then the kernel $k_{x,y}^{T \circledast S}$ is lifted from k_{x_T, y_T}^T . By assumption we have $k_{x_T, y_T}^T \leq k_{gx_T, gy_T}^{T'}$ modulo the iso, and since $k_{gx_T, gy_T}^{T'}$ induces $k_{(g \circledast f)x, (g \circledast f)y}^{T' \circledast S'}$ (via lifting and across the symmetry isomorphism), we conclude that $k_{x,y}^{T \circledast S} \leq k_{(g \circledast f)x, (g \circledast f)y}^{T' \circledast S'}$, again modulo the iso. The same argument goes for σ -extensions, and so the map $\underline{g} \circledast \underline{f}$ satisfies axiom (2).

The conditions are not affected by hiding, so the composition $\underline{g} \circledcirc \underline{f}$ is a map of probabilistic strategies. \square

Unitors and associators. We now identify the structural 2-cells, starting with the unitors.

Lemma 7.23. *For every probabilistic strategy $\sigma : \underline{\mathcal{S}} \rightarrow \underline{\mathcal{A}}^\perp \parallel \underline{\mathcal{B}}$, there is an isomorphism $\lambda_\sigma : \alpha_{\underline{\mathcal{B}}} \odot \sigma \cong \sigma$ of probabilistic strategies.*

Proof. The iso is the same as in **MG**. All we need to show is the kernel preservation property. We show the existence of λ_σ^{-1} .

Let $x_S \text{--} \text{c}^e y_S$ with e positive. We know that $\lambda_\sigma^{-1} x_S = (x_B^* \parallel x_B) \odot x_S^*$ and $\lambda_\sigma^{-1} y_S = (y_B^* \parallel y_B) \odot y_S^*$, where x_S^* (resp. y_S^*) is x_S (resp. y_S) with its maximal negative B -moves taken out.

Because λ_σ^{-1} is an isomorphism of essps, we have $(x_B^* \parallel x_B) \odot x_S^* \text{--} \text{c}^+ (y_B^* \parallel y_B) \odot y_S^*$, and we investigate the shape of covering chains for the extension $(x_B^* \parallel x_B) \otimes x_S^* \subseteq (y_B^* \parallel y_B) \otimes y_S^*$. It turns out these are of two kinds: either e is an A move with all immediate predecessors in A , in which case $y_B = x_B$ and so we have

$$(x_B^* \parallel x_B) \otimes x_S^* \text{--} \text{c}^+ (y_B^* \parallel y_B) \otimes y_S^*,$$

or, e is an B -move or some of its immediate predecessors are, in which case the necessary moves must be played by copycat in the internal copy of B , and:

$$(x_B^* \parallel x_B) \otimes x_S^* \text{--} \text{c}^+ (y_B^* \parallel x_B) \otimes z_S \text{--} \text{c}^+ (y_B^* \parallel y_B) \otimes y_S^*$$

for some z_S . In the first case, the kernel we are interested in is directly lifted from $\underline{\mathcal{S}}$, whereas in the second case it is lifted from $\underline{\mathcal{S}}$ and composed with an identity kernel, by definition of copycat, so we are done. \square

For associativity, as discussed in Chapter 2, it suffices to consider the associator at the level of interactions $(\mathcal{R} \otimes \mathcal{T}) \otimes \mathcal{S} \rightarrow \mathcal{R} \otimes (\mathcal{T} \otimes \mathcal{S})$ as discussed. This is an isomorphism of essps, so it clearly preserves covering chains and associated kernels, and therefore the kernel preservation property holds after hiding.

The coherence laws hold because they hold in \mathcal{MG} , so that:

Theorem 7.24. *There is a pseudo-double category \mathcal{PMG} where*

- PMG_0 consists of race-free, standard Borel measurable games and maps between them, and
- PMG_1 consists of probabilistic standard Borel strategies and maps between them.

7.4.2 Monoidal structure

To show that probabilistic strategies support a tensor construction, we first investigate the parallel composition $\underline{\mathcal{S}} \parallel \underline{\mathcal{T}}$ of two measurable essps $\underline{\mathcal{S}}$ and $\underline{\mathcal{T}}$ with conditional valuations $\mathcal{K}^{\mathcal{S}}$ and $\mathcal{K}^{\mathcal{T}}$.

To an extension $x = x_S \parallel x_T \subseteq^+ y = y_S \parallel y_T$ in $\mathcal{C}(S \parallel T)$ we assign the appropriate product kernel:

$$k_{x,y}^{S \parallel T} = k_{x_S, y_S}^{\mathcal{S}} \otimes k_{x_T, y_T}^{\mathcal{T}}.$$

We must show that this satisfies the conditions for a valuation. We start with the drop condition; other conditions will be straightforward.

Lemma 7.25. *Let $y \subseteq^+ x_1, \dots, x_n \in \mathcal{C}(S \parallel T)$, and let $u \in \mathcal{M}_y$. Then*

$$\text{dr}[y; x_1, \dots, x_n](u) \geq 0.$$

Proof. Let $y \subseteq^+ x_1, \dots, x_n \in \mathcal{C}(S \parallel T)$, writing $y = y^S \parallel y^T$ and $x_i = x_i^S \parallel x_i^T$ for each i . The proof that $\text{dr}[y; x_1, \dots, x_n](u) \geq 0$ uses two facts:

$$\text{dr}[y; x_1, \dots, x_n](u) \geq \text{dr}[y; y^S \parallel x_1^T, x_1^S \parallel y^T, \dots, y^S \parallel x_n^T, x_n^S \parallel y^T](u) \quad (7.3)$$

and

$$\begin{aligned} \text{dr}[y; y^S \parallel x_1^T, x_1^S \parallel y^T, \dots, y^S \parallel x_n^T, x_n^S \parallel y^T](u) \\ = \text{dr}[y^S; x_1^S, \dots, x_n^S](u_S) \text{dr}[y^T; x_1^T, \dots, x_n^T](u_T). \end{aligned} \quad (7.4)$$

Combining (1) and (2) and applying the drop condition for \mathcal{K}^S and \mathcal{K}^T gives the desired result.

Proof of (1). By Lemma 7.15, we have $\text{dr}[y; x_1, \dots, x_n](u) = \text{dr}[y; x_1, \dots, x_{n-1}](u) - \int_{u_n \in \mathcal{M}(x_n)} \text{dr}[x_n; x_1 \vee x_n, \dots, x_{n-1} \vee x_n](u_n) k(u, \text{d}u_n)$. Using this definition, one can show by a straightforward induction on n that if $x_i \subseteq x'_i$ for all i , then $\text{dr}[y; x_1, \dots, x_n](u) \leq \text{d}_v[y; x'_1, \dots, x'_n]$. Fact (1) then holds because $\text{d}_v[y; x_1, \dots, x_n] = \text{d}_v[y; x_1, x_1, \dots, x_n, x_n]$, and $y^S \parallel x_i^T$ and $x_i^S \parallel y^T$ are subsets of x_i for all i .

Proof of (2). In what follows, I ranges over nonempty subsets of $\{1, \dots, n\}$ such that $\bigcup_{i \in I} x_i^S \in \mathcal{C}(S)$, J over nonempty subsets of $\{1, \dots, n\}$ such that $\bigcup_{j \in J} x_j^T \in \mathcal{C}(T)$. We write z_1, \dots, z_{2n} for $y^S \parallel x_1^T, x_1^S \parallel y^T, \dots, y^S \parallel x_n^T, x_n^S \parallel y^T$, and K ranges over nonempty subsets of $\{1, \dots, 2n\}$ such that $\bigcup_{k \in K} z_k \in \mathcal{C}(S \parallel T)$. To alleviate notation we write x_I^S for $\bigcup_{i \in I} x_i^S$, and so on. Then we compute:

$$\begin{aligned} & \text{dr}[y^S; x_1^S, \dots, x_n^S](u_S) \text{dr}[y^T; x_1^T, \dots, x_n^T](u_T) \\ &= \left[1 - \sum_I (-1)^{|I|+1} k_{y^S, x_I^S}(u_S, \mathcal{M}(x_I^S)) \right] \left[1 - \sum_J (-1)^{|J|+1} k_{y^T, x_J^T}(u_T, \mathcal{M}(x_J^T)) \right] \\ &= 1 - \sum_I (-1)^{|I|+1} k_{y^S, x_I^S}(u_S, \mathcal{M}(x_I^S)) - \sum_J (-1)^{|J|+1} k_{y^T, x_J^T}(u_T, \mathcal{M}(x_J^T)) \\ & \quad + \sum_I \sum_J (-1)^{|I|+|J|} k_{y^S, x_I^S}(u_S, \mathcal{M}(x_I^S)) k_{y^T, x_J^T}(u_T, \mathcal{M}(x_J^T)) \\ &= 1 - \sum_I (-1)^{|I|+1} k_{y, x_I^S \parallel y^T}(u, \mathcal{M}(x_I^S \parallel y^T)) - \sum_J (-1)^{|J|+1} k_{y, y^S \parallel x_J^T}(u, \mathcal{M}(y^S \parallel x_J^T)) \\ & \quad + \sum_I \sum_J (-1)^{|I|+|J|} k_{y, x_I^S \parallel x_J^T}(u, \mathcal{M}(x_I^S \parallel x_J^T)) \\ &= 1 - \sum_K (-1)^{|K|+1} k_{y, z_K}(u, \mathcal{M}(z_K)) \\ &= \text{dr}[y; z_1, \dots, z_{2n}]. \end{aligned}$$

□

From this we define the tensor product of two strategies.

Lemma 7.26. *For two probabilistic strategies $\underline{\sigma} : \underline{\mathcal{S}} \rightarrow \underline{\mathcal{A}}^\perp \parallel \underline{\mathcal{A}}'$ and $\underline{\tau} : \underline{\mathcal{T}} \rightarrow \underline{\mathcal{B}}^\perp \parallel \underline{\mathcal{B}}'$, the family $\mathcal{K}^{\mathcal{S} \parallel \mathcal{T}}$ defines a conditional valuation and thus there is a probabilistic strategy*

$$\underline{\sigma} \otimes \underline{\tau} : \underline{\mathcal{S}} \parallel \underline{\mathcal{T}} \rightarrow (\underline{\mathcal{A}} \parallel \underline{\mathcal{B}})^\perp \parallel (\underline{\mathcal{A}}' \parallel \underline{\mathcal{B}}').$$

Proof. Routine verification. □

The tensor construction for strategies induces a symmetric monoidal structure on \mathcal{PMG} . Observe that \mathcal{PMG} is isofibrant, since the lifting and co-lifting constructions are copycat-like and can always be equipped with the canonical *deterministic* conditional valuation. Therefore:

Lemma 7.27. *The bicategory $\mathcal{H}(\mathcal{PMG})$ is symmetric monoidal.*

7.4.3 The cartesian closed bicategory

We continue on the usual path. We consider the sub-bicategory \mathbf{PMG} of $\mathcal{H}(\mathcal{PMG})$ having:

- objects: negative, rooted, race-free, standard Borel measurable arenas;
- morphisms: negative, well-threaded, rooted, standard Borel probabilistic strategies;
- 2-cells: maps of probabilistic strategies.

The constructions performed in \mathbf{MG} to obtain cartesian closure extend to valuations, and all steps are unsurprising. Here is a brief account:

- The monoidal structure is closed: since $\text{Cur}(\underline{\sigma})$ and $\underline{\sigma}$ have the same internal event structure $\underline{\mathcal{S}}$, leaving the valuation $\mathcal{K}^{\mathcal{S}}$ unchanged makes Cur an isomorphism also in this context.
- In the binary product $\underline{\mathcal{S}} \& \underline{\mathcal{T}}$ the two components are completely inconsistent: $\mathcal{K}^{\mathcal{S} \& \mathcal{T}}$ is the obvious construction and all axioms are immediate.
- The pseudo-comonad $!$ on \mathbf{MG} becomes one on \mathbf{PMG} : we only ever consider finite configurations of $!\mathcal{S}$, and so whenever necessary we may restrict to an appropriate finitary tensor $\bigotimes_{i \in I} \mathcal{S}_i$, and apply results about the monoidal structure. This gives a valuation $\mathcal{K}^{!\mathcal{S}}$.

Theorem 7.28. *The bicategory $\mathbf{PMG}_!$ is cartesian closed.*

7.4.4 The discrete sub-bicategory

We conclude by observing that for a measurable strategy of the form $\text{disc}(\sigma)$, for $\sigma : \mathcal{S} \rightarrow \mathcal{A}$ a strategy in \mathbf{G} , to give a conditional valuation (*i.e.* a family of kernels) on $\text{disc}(\sigma)$ is to give a conditional valuation (*i.e.* a family of coefficients) on σ . This way we obtain a faithful (and full and faithful on 2-cells) pseudo-functor $\mathbf{PG} \rightarrow \mathbf{PMG}$, and since discrete strategies are closed under all constructions, the structure of \mathbf{PMG} (resp. $\mathbf{PMG}_!$) is already present in \mathbf{PG} (resp. $\mathbf{PG}_!$).

Chapter 8

Conclusion

We have presented several models for concurrent games semantics of programs, building on existing work by Winskel, Clairambault, Castellan and others. For each model we give a notion of game, a notion of strategy on a game, and a notion of maps between strategies, and investigate how these objects interact. Each time we describe in detail the development of a bicategory with structure.

The thesis contains four models:

- Concurrent games with symmetry, known as *thin concurrent games* in [CCW19]. This model is not new, but some progress is made in understanding the structure underlying the theory. This model is the main basis for the next three.
- Probabilistic concurrent games, enriching the latter with probability.
- Measurable concurrent games, in which the games and strategies of the first model are further equipped with measure-theoretic structure, so as to refine the modelling of computation with continuous data types.
- Probabilistic measurable concurrent games, in which we use measure-theoretic probability theory to enrich the latter with quantitative information.

We additionally discuss two applications of the probabilistic concurrent games model, *i.e.* the second model in the list above. The two applications are to an untyped and a typed language, respectively, and we have described connection with related work.

Throughout the thesis we have paid special attention to the *bicategorical* structure of the models above, *i.e.* the various coherence laws obeyed by the 2-cells. Furthermore, in the construction of cartesian closed bicategories, efforts have been made to follow a principled approach, so that, when building new models, much of the verification work can be avoided. As should be clear from the thesis, we have taken full advantage of this. But this also applies to potential further work.

Remarkably, in the development of measurable concurrent games (Chapter 6), no aspect of the development appeals to specific properties of measurable spaces. Our construction only requires a category with finite products and pullbacks, so in principle **Meas** could be replaced with another category of spaces, should one require structure of a different kind.

We speculate further. It is well-known that the *kernels* used in Chapter 7 to make probabilistic the measurable strategies of Chapter 6 are Kleisli maps $X \rightarrow GY$ for G a probability monad on \mathbf{Meas} [Gir82]. Then, kernel composition is Kleisli composition, and the identity kernel is the Kleisli identity. So the question must be asked whether one could generalise the model of probabilistic measurable games to a model for a Moggi-style computational λ -calculus [Mog91] parametrised by a monad on an arbitrary category (with finite products and pullbacks). This is less clear, as some of the results in Chapter 7 (Lemma 7.2, for example) are inherently based on the nature of kernels. We leave this for further work.

A more concrete direction for future work consists in applying the framework of measurable games to a versions of PCF and Probabilistic PCF with (say) a type of real numbers. Such a language is studied for instance in [EPT17]. The questions of innocence and definability in this framework remain to be answered, although recent advances in *quantum* concurrent game semantics may indeed be relevant [CdVW19].

Finally, in yet another direction, we aim to develop intensional semantics for “statistical” probabilistic programming as discussed briefly in the introduction. This direction is promising: an initial step was recently made in this direction in joint work with Simon Castellan [CP19].

Bibliography

- [A⁺61] Robert J Aumann et al. Borel structures for function spaces. *Illinois Journal of Mathematics*, 5(4):614–630, 1961.
- [AB06] Samy Abbes and Albert Benveniste. True-concurrency probabilistic models: Branching cells and distributed probabilities for event structures. *Information and Computation*, 204(2):231–274, 2006.
- [AHM98] Samson Abramsky, Kohei Honda, and Guy McCusker. A fully abstract game semantics for general references. In *Proceedings. Thirteenth Annual IEEE Symposium on Logic in Computer Science (Cat. No. 98CB36226)*, pages 334–344. IEEE, 1998.
- [AJ94] Samson Abramsky and Radha Jagadeesan. Games and full completeness for multiplicative linear logic. *The Journal of Symbolic Logic*, 59(2):543–574, 1994.
- [AJM00] Samson Abramsky, Radha Jagadeesan, and Pasquale Malacaria. Full abstraction for PCF. *Inf. Comput.*, 163(2):409–470, 2000.
- [AM96] Samson Abramsky and Guy McCusker. Linearity, sharing and state: a fully abstract game semantics for idealized algol with active expressions. *Electr. Notes Theor. Comput. Sci.*, 3:2–14, 1996.
- [AM97] Samson Abramsky and Guy McCusker. Call-by-value games. In *International Workshop on Computer Science Logic*, pages 1–17. Springer, 1997.
- [Bar84] Hendrik P Barendregt. *The lambda calculus*, volume 3. North-Holland Amsterdam, 1984.
- [BDLGS16] Johannes Borgström, Ugo Dal Lago, Andrew D Gordon, and Marcin Szymczak. A lambda-calculus foundation for universal probabilistic programming. In *ACM SIGPLAN Notices*, volume 51, pages 33–46. ACM, 2016.
- [Bil08] Patrick Billingsley. *Probability and measure*. John Wiley & Sons, 2008.
- [Bla92] Andreas Blass. A game semantics for linear logic. *Annals of Pure and Applied logic*, 56(1-3):183–220, 1992.

- [Cas17] Simon Castellan. *Concurrent structures in game semantics*. PhD thesis, Université de Lyon, 2017.
- [CC16] Simon Castellan and Pierre Clairambault. Causality vs. interleavings in concurrent game semantics. In *27th International Conference on Concurrency Theory, CONCUR 2016, August 23-26, 2016, Québec City, Canada*, pages 32:1–32:14, 2016.
- [CCHW18] Simon Castellan, Pierre Clairambault, Jonathan Hayman, and Glynn Winskel. Non-angelic concurrent game semantics. In *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*, pages 3–19, 2018.
- [CCPW18] Simon Castellan, Pierre Clairambault, Hugo Paquet, and Glynn Winskel. The concurrent game semantics of probabilistic PCF. In *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science*, pages 215–224. ACM, 2018.
- [CCRW17] Simon Castellan, Pierre Clairambault, Silvain Rideau, and Glynn Winskel. Games and strategies as event structures. *LMCS*, 13(3), 2017.
- [CCW14] Simon Castellan, Pierre Clairambault, and Glynn Winskel. Symmetry in concurrent games. In *CSL-LICS '14, Vienna*, pages 28:1–28:10, 2014.
- [CCW15] Simon Castellan, Pierre Clairambault, and Glynn Winskel. The parallel intensionally fully abstract games model of PCF. In *Proceedings, LICS 2015, Kyoto, Japan, July 6-10, 2015*, pages 232–243, 2015.
- [CCW17] Simon Castellan, Pierre Clairambault, and Glynn Winskel. Observably deterministic concurrent strategies and intensional full abstraction for parallel-or. In *2nd International Conference on Formal Structures for Computation and Deduction, FSCD 2017, September 3-9, 2017, Oxford, UK*, pages 12:1–12:16, 2017.
- [CCW19] Simon Castellan, Pierre Clairambault, and Glynn Winskel. Thin games with symmetry and concurrent Hyland-Ong games. *Logical Methods in Computer Science*, 15(1), 2019.
- [CdVW19] Pierre Clairambault, Marc de Visme, and Glynn Winskel. Game semantics for quantum programming. *PACMPL*, 3(POPL):32:1–32:29, 2019.
- [CP18] Pierre Clairambault and Hugo Paquet. Fully abstract models of the probabilistic lambda-calculus. In *27th EACSL Annual Conference on Computer Science Logic, CSL 2018, September 4-7, 2018, Birmingham, UK*, pages 16:1–16:17, 2018.

- [CP19] Simon Castellan and Hugo Paquet. Probabilistic programming inference via intensional semantics. In *European Symposium on Programming*, pages 322–349. Springer, 2019.
- [DE11] Vincent Danos and Thomas Ehrhard. Probabilistic coherence spaces as a model of higher-order probabilistic computation. *Inf. Comput.*, 209(6):966–991, 2011.
- [DH02] Vincent Danos and Russell Harmer. Probabilistic game semantics. *ACM Trans. Comput. Log.*, 3(3):359–382, 2002.
- [dV19] Marc de Visme. Event structures for mixed choice. In *30th International Conference on Concurrency Theory, CONCUR 2019, August 27-30, 2019, Amsterdam, the Netherlands.*, pages 11:1–11:16, 2019.
- [EH19] Clovis Eberhart and Tom Hirschowitz. Fibred Pseudo Double Categories for Game Semantics. *Theory and Applications of Categories*, 34:514 – 572, 2019.
- [EPT11] Thomas Ehrhard, Michele Pagani, and Christine Tasson. The computational meaning of probabilistic coherence spaces. In *2011 IEEE 26th Annual Symposium on Logic in Computer Science*, pages 87–96. IEEE, 2011.
- [EPT17] Thomas Ehrhard, Michele Pagani, and Christine Tasson. Measurable cones and stable, measurable functions: a model for probabilistic higher-order programming. *Proceedings of the ACM on Programming Languages*, 2(POPL):59, 2017.
- [ETP14] Thomas Ehrhard, Christine Tasson, and Michele Pagani. Probabilistic coherence spaces are fully abstract for probabilistic PCF. In *Proceedings, POPL '14, San Diego, USA, January 20-21*, pages 309–320, 2014.
- [FGHW08] Marcelo Fiore, Nicola Gambino, Martin Hyland, and Glynn Winskel. The cartesian closed bicategory of generalised species of structures. *Journal of the London Mathematical Society*, 77(1):203–220, 2008.
- [FS19] Marcelo Fiore and Philip Saville. A type theory for cartesian closed bicategories (extended abstract). In *34th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2019, Vancouver, BC, Canada, June 24-27, 2019*, pages 1–13. IEEE, 2019.
- [GHNR14] Andrew D Gordon, Thomas A Henzinger, Aditya V Nori, and Sriram K Rajamani. Probabilistic programming. In *Proceedings of the on Future of Software Engineering*, pages 167–181. ACM, 2014.
- [Gir82] Michele Giry. A categorical approach to probability theory. In *Categorical aspects of topology and analysis*, pages 68–85. Springer, 1982.
- [Gir87] Jean-Yves Girard. Linear logic. *Theoretical computer science*, 50(1):1–101, 1987.

- [Gir88] Jean-Yves Girard. Normal functors, power series and λ -calculus. *Annals of pure and applied logic*, 37(2):129–177, 1988.
- [Gou15] Jean Goubault-Larrecq. Full abstraction for non-deterministic and probabilistic extensions of PCF I: the angelic cases. *J. Log. Algebraic Methods Program.*, 84(1):155–184, 2015.
- [Gou19] Jean Goubault-Larrecq. A probabilistic and non-deterministic call-by-push-value language. In *34th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2019, Vancouver, BC, Canada, June 24–27, 2019*, pages 1–13. IEEE, 2019.
- [GP99] Marco Grandis and Robert Paré. Limits in double categories. *Cahiers de Topologie et Géométrie Différentielle Catégoriques*, 40(3):162–220, 1999.
- [GP04] Marco Grandis and Robert Paré. Adjoint for double categories. *Cahiers de topologie et géométrie différentielle catégoriques*, 45(3):193–240, 2004.
- [Har99] Russell Spencer Harmer. *Games and full abstraction for non-deterministic languages*. PhD thesis, Imperial College London (University of London), 1999.
- [HKSJ17] Chris Heunen, Ohad Kammar, Sam Staton, and Hongseok Yang. A convenient category for higher-order probability theory. In *LICS’17, Reykjavik.*, pages 1–12. IEEE Computer Society, 2017.
- [HM99] Russell Harmer and Guy McCusker. A fully abstract game semantics for finite nondeterminism. In *Proceedings. 14th Symposium on Logic in Computer Science (Cat. No. PR00158)*, pages 422–430. IEEE, 1999.
- [HO00] J. M. E. Hyland and C.-H. Luke Ong. On full abstraction for PCF: i, ii, and III. *Inf. Comput.*, 163(2):285–408, 2000.
- [HY97] Kohei Honda and Nobuko Yoshida. Game theoretic analysis of call-by-value computation. In *International Colloquium on Automata, Languages, and Programming*, pages 225–236. Springer, 1997.
- [Jon90] Claire Jones. *Probabilistic non-determinism*. PhD thesis, University of Edinburgh, 1990.
- [JP89] Claire Jones and Gordon D. Plotkin. A probabilistic powerdomain of evaluations. In *(LICS ’89), Pacific Grove, California, USA.*, pages 186–195. IEEE Computer Society, 1989.
- [JT98] Achim Jung and Regina Tix. The troublesome probabilistic powerdomain. *Electr. Notes Theor. Comput. Sci.*, 13:70–91, 1998.
- [Kal06] Olav Kallenberg. *Foundations of modern probability*. Springer Science & Business Media, 2006.

- [KF09] Daphne Koller and Nir Friedman. *Probabilistic graphical models: principles and techniques*. MIT press, 2009.
- [KNO02] Andrew D Ker, Hanno Nickau, and C-H Luke Ong. Innocent game models of untyped λ -calculus. *Theoretical Computer Science*, 272(1-2):247–292, 2002.
- [Koz79] Dexter Kozen. Semantics of probabilistic programs. In *20th Annual Symposium on Foundations of Computer Science (sfcs 1979)*, pages 101–114. IEEE, 1979.
- [Lac00] Stephen Lack. A coherent approach to pseudomonads. *Advances in Mathematics*, 152(2):179–202, 2000.
- [Lai97] James Laird. Full abstraction for functional languages with control. In *LICS'97, Warsaw, Poland, June 29 - July 2, 1997*, pages 58–67, 1997.
- [Lam92] François Lamarche. Quantitative domains and infinitary algebras. *Theor. Comput. Sci.*, 94(1):37–62, 1992.
- [Lei98] Tom Leinster. Basic bicategories. *arXiv preprint math/9810017*, 1998.
- [Lev16] Thomas Leventis. *Lambdas-théories probabilistes*. PhD thesis, Aix-Marseille, 2016.
- [LMMP13] Jim Laird, Giulio Manzonetto, Guy McCusker, and Michele Pagani. Weighted relational models of typed lambda-calculi. In *2013 28th Annual ACM/IEEE Symposium on Logic in Computer Science*, pages 301–310. IEEE, 2013.
- [LP19] Thomas Leventis and Michele Pagani. Strong adequacy and untyped full-abstraction for probabilistic coherence spaces. In *International Conference on Foundations of Software Science and Computation Structures*, pages 365–381. Springer, 2019.
- [Mel09] Paul-André Melliès. Categorical semantics of linear logic. *Panoramas et synthèses*, 27:15–215, 2009.
- [Mel19] Paul-André Melliès. Categorical combinatorics of scheduling and synchronization in game semantics. *Proceedings of the ACM on Programming Languages*, 3(POPL):23, 2019.
- [ML13] Saunders Mac Lane. *Categories for the working mathematician*, volume 5. Springer Science & Business Media, 2013.
- [MM12] Saunders MacLane and Ieke Moerdijk. *Sheaves in geometry and logic: A first introduction to topos theory*. Springer Science & Business Media, 2012.
- [Mog91] Eugenio Moggi. Notions of computation and monads. *Information and computation*, 93(1):55–92, 1991.

- [NPW79] Mogens Nielsen, Gordon Plotkin, and Glynn Winskel. Petri nets, event structures and domains. In *Semantics of Concurrent Computation*, pages 266–284. Springer, 1979.
- [Pea00] Judea Pearl. *Causality: models, reasoning and inference*, volume 29. Springer, 2000.
- [Pit93] Andrew M Pitts. Computational adequacy via mixed inductive definitions. In *International Conference on Mathematical Foundations of Programming Semantics*, pages 72–82. Springer, 1993.
- [Plo77] Gordon D. Plotkin. LCF considered as a programming language. *Theor. Comput. Sci.*, 5(3):223–255, 1977.
- [PW18] Hugo Paquet and Glynn Winskel. Continuous probability distributions in concurrent games. *Electronic Notes in Theoretical Computer Science*, 341:321–344, 2018.
- [RW11] Silvain Rideau and Glynn Winskel. Concurrent strategies. In *LICS '11, June 21-24, 2011, Toronto, Canada*, pages 409–418, 2011.
- [Sah80] N. Saheb-Djahromi. Cpo’s of measures for nondeterminism. *Theor. Comput. Sci.*, 12:19–37, 1980.
- [Sco69] Dana S Scott. A type-theoretical alternative to ISWIM, CUCH, OWHY. *Theoretical Computer Science*, 121(1-2):411–440, 1969.
- [Sco70] Dana Scott. *Outline of a mathematical theory of computation*. Oxford University Computing Laboratory, Programming Research Group, 1970.
- [Sco80] Dana S Scott. Relating theories of the lambda calculus. *To HB Curry: Essays on combinatory logic, lambda calculus and formalism*, pages 403–450, 1980.
- [SD78] Nasser Saheb-Djahromi. Probabilistic LCF. In *International Symposium on Mathematical Foundations of Computer Science*, pages 442–451. Springer, 1978.
- [See87] Robert AG Seely. Modelling computations: a 2-categorical framework. In *LICS*, pages 65–71, 1987.
- [Shu10] Michael A Shulman. Constructing symmetric monoidal bicategories. *arXiv preprint arXiv:1004.0993*, 2010.
- [Sta16] Michael Stay. Compact closed bicategories. *Theory and applications of Categories*, 31(26):755–798, 2016.
- [Sta17] Sam Staton. Commutative semantics for probabilistic programming. In *European Symposium on Programming*, pages 855–879. Springer, 2017.

- [SW10] Sam Staton and Glynn Winskel. On the expressivity of symmetry in event structures. In *2010 25th Annual IEEE Symposium on Logic in Computer Science*, pages 392–401. IEEE, 2010.
- [SYW⁺16] Sam Staton, Hongseok Yang, Frank D. Wood, Chris Heunen, and Ohad Kammar. Semantics for probabilistic programming: higher-order functions, continuous distributions, and soft constraints. In *Proceedings of LICS '16, New York, NY, USA, July 5-8, 2016*, pages 525–534, 2016.
- [TO14] Takeshi Tsukada and C.-H. Luke Ong. Innocent strategies are sheaves over plays - deterministic, non-deterministic and probabilistic innocence. *CoRR*, abs/1409.2764, 2014.
- [TO15] Takeshi Tsukada and C.-H. Luke Ong. Nondeterminism in game semantics via sheaves. In *30th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2015, Kyoto, Japan, July 6-10, 2015*, pages 220–231, 2015.
- [VKS19] Matthijs Vákár, Ohad Kammar, and Sam Staton. A domain theory for statistical probabilistic programming. *Proceedings of the ACM on Programming Languages*, 3(POPL):36, 2019.
- [VWV04] Daniele Varacca, Hagen Völzer, and Glynn Winskel. Probabilistic event structures and domains. In *International Conference on Concurrency Theory*, pages 481–496. Springer, 2004.
- [Win] Glynn Winskel. Event structures, stable families and concurrent games. Notes associated with ERC grant ECSYM, <https://www.cl.cam.ac.uk/gw104/ecsym-notes.pdf>.
- [Win86] Glynn Winskel. Event structures. In *Advances in Petri Nets*, pages 325–392, 1986.
- [Win99] Glynn Winskel. Event structures as presheaves: two representation theorems. In *International Conference on Concurrency Theory*, pages 541–556. Springer, 1999.
- [Win07] Glynn Winskel. Event structures with symmetry. *Electr. Notes Theor. Comput. Sci.*, 172:611–652, 2007.
- [Win12] Glynn Winskel. Deterministic concurrent strategies. *Formal Asp. Comput.*, 24(4-6):647–660, 2012.
- [Win13a] Glynn Winskel. Distributed probabilistic and quantum strategies. *Electr. Notes Theor. Comput. Sci.*, 298:403–425, 2013.
- [Win13b] Glynn Winskel. Strategies as profunctors. In *Proceedings, FOSSACS '13, 16-24 March 2013, Rome, Italy*, pages 418–433, 2013.

- [Win15] Glynn Winskel. On probabilistic distributed strategies. In *International Colloquium on Theoretical Aspects of Computing*, pages 69–88. Springer, 2015.
- [Zei15] Noam Zeilberger. Linear lambda terms as invariants of rooted trivalent maps. *CoRR*, abs/1512.06751, 2015.