THE QUALITATIVE COLLAPSE OF CONCURRENT GAMES

PIERRE CLAIRAMBAULT^a

Aix Marseille Univ, CNRS, LIS, Marseille, France *e-mail address*: First.Last@lis-lab.fr

ABSTRACT. In this paper, we construct an interpretation-preserving functor from a category of concurrent games to the category of Scott domains and Scott-continuous functions. We give a concrete description of this functor, extending earlier results on the *relational* collapse of game semantics. The crux is an intricate combinatorial lemma allowing us to synchronize states of strategies which reach the same resources, but with different multiplicity.

Putting this together with the previously established relational collapse, this provides a new proof of the *qualitative-quantitative* correspondence first established by Ehrhard in his celebrated *extensional collapse* theorem. Whereas Ehrhard's proof is indirect and rests on an abstract realizability construction, our result gives a concrete, combinatorial description of the extraction of quantitative information from a qualitative model.

1. INTRODUCTION

The heart of denotational semantics is certainly *domain theory*, where types are interpreted as certain partially ordered sets, and programs as (continuous) functions between those. This idea, originally pioneered by Scott and Strachey [SS71], has spread wide and far, and underlies much of the modern theory of programming languages. In the terminology of this paper, this functional semantics is *qualitative*: it tracks the amount of information about the input need to compute a given part of the output, but not *how many times* that information is needed, or how many times the argument of a function is evaluated.

Another deeply influential discovery, in that field of research, is Girard's invention of *linear logic* [Gir87]. Linear logic is a logic of *resources*; it gives a special status to those functions that are *linear* in the sense that they evaluate their argument *exactly once*. Starting with the interpretation of λ -terms as normal functors [Gir88], linear logic prompted the development of denotational models that are sensitive to resources, in the sense that they also record the multiplicity of resource usage: in the terminology of this paper, they are *quantitative*. Quantitative models have been under active development in the following three decades, with a number of remarkable achievements. For instance, quantitative models (and their type-theoretic presentations as *non-idempotent intersection types*) provide a semantic characterization of execution time [dC18]. Their resource-sensitivity lets them track numerous quantitative aspects of computation [LMMP13], or provide models of properly quantitative computational effects, such as probabilistic choice [EPT18] or quantum effects [PSV14], for which they give fully abstract models [EPT18, CdV20].

The drawback of this quantitative aspect, however, is that they are infinitary. Even for the simply-typed λ -calculus with a finite interpretation for ground types, they give infinitary semantics, because they rely on finite multisets to represent the arrow type. A "proof" that a certain point is in the quantitative semantics of a term (which can often be represented as a derivation in a non-idempotent intersection type system), is really a de-temporised, "static" representation of the full execution. In contrast, the functional models as in domain theory, and their syntactic presentations as idempotent intersection type systems, remain finitary: for instance, they give a finite interpretation to simply-typed programs with finite ground types. In this way, they talk by finitary means of an infinitary object: the execution – this is the key to their algorithmic use in *e.g.* higher-order model-checking [Aeh06, KO09].

There is a fascinating scientific tension between these qualitative and quantitative models. On the one hand, they are remarkably similar: with the right presentation, the only significant difference in their construction is whether the exponential modality should be based on finite sets or finite multisets. On the other hand, the associated proof methods are very different: quantitative models are infinitary, but their connection with the execution is simple logically (though it can still be subtle combinatorially), allowing them to provide useful program approximants [BM20]; qualitative models are finitary, but linking them with execution requires tools with considerable logical complexity, such as logical relations. Surprisingly, this tension has been somewhat little studied, perhaps also because the two families of models correspond to different communities. However, there is one important paper that strikes right at that tension: Ehrhard's result that the linear Scott model of the simply-typed λ -calculus is the *extensional collapse* of its relational model [Ehr12]. Ehrhard's result entails, in particular, that a point a in the qualitative model is in the semantics of a program M iff it has a "quantitativation" a' in the quantitative semantics of M. At the core of this result is the construction of a model that is somewhat hybrid between qualitative and quantitative; of quantitative relations which behave well with respect to a preorder relation rearranging resources. But this hybrid model is obtained by formulating and maintaining an invariant (by *biorthogonality*) implying this quantitativation, it gives us no combinatorial understanding of that process, and no way to compute it in concrete cases.

Here, we provide a combinatorial understanding of this quantitativation process, using game semantics. Game semantics is another quantitative denotational model, originally developped to attack the famous full abstraction problem for PCF [AJM00,HO00]. Game semantics enriches the relational model with time or causality, presenting interactive executions of a program with its runtime environment as plays on a game whose rules are determined by the type. Despite its clear intellectual affiliation with quantitative semantics, the precise relationship between games and relational models has been the center of a longstanding line of research [BDER97, Ehr96, Mel06, Mel05, Bou09]. In the modern dresses of thin concurrent games [CCW19, Cla24], building on concurrent games on event structures [RW11] and in the footsteps of Melliès's insightful work on asynchronous games [Mel05], this relationship now appears as a simple forgetful interpretation-preserving functor to the relational model, erasing the "dynamic" causal dependency coming from the program, keeping only the "static" causal dependency from the type – this is summed up in [Cla24], see also [CCPW18, CdV20] for extensions with quantitative features and [COP23] for a bicategorical version.

In this paper, we complement this "relational collapse" with a related interpretationpreserving functor to the *linear Scott model*, a linear decomposition of a (full subcategory of) Scott domains due independently to Huth [Hut93] and Winskel [Win98]. To construct a Scott domain from a game, we equip the latter with adequate notions of morphisms, cartesian morphisms, which allow the rearrangment (contraction and weakening) of resources. The crux of the issue is then to show that this collapse operation to the linear Scott model preserves composition: this rests on a crucial proposition (Proposition 6.10) showing that if innocent strategies can synchronize up to cartesian morphisms, then one can find adequate expansions of the strategies making them synchronize on the nose. This forms the core of our combinatorial account of Ehrhard's quantitativation result: because our games model has interpretation-preserving functors to both the relational model and the Scott model, we also obtain a precise connection between the two (Theorem 6.16). This is also a contribution to the line of work connecting game semantics to "static" semantics, targetting from the first time a qualitative semantics, spanning accross communities: Scott domains.

After this scientific introduction, we include in the next section a more technical introduction, setting the scene and giving the main intuitions, and exposing the outline.

2. The Qualitative and the Quantitative

2.1. The Relational Model and Quantitative Semantics. Our starting point, in this discussion, will be the *relational model* of linear logic.

2.1.1. The relational model. At its heart, it is the interpretation of the λ -calculus into the category **Rel** of sets and relations, with which we assume that the reader is familiar (see *e.g.* [Ehr12] for a reference). **Rel** is a Seely category [Mel09]: its monoidal product is the cartesian product of sets, its cartesian product is given by the disjoint union, and its exponential modality ! sends a set A to the set $\mathscr{M}_f(A)$ of finite multisets of elements of A. We adopt standard conventions for multisets: we adopt a list notation $[a_1, \ldots, a_n]$ possibly with repetitions, with the empty multiset written []. Multiset union is written with a sum +.

As **Rel** is a Seely category, one can consider the Kleisli category for the exponential comonad !, which is cartesian closed. One can then interpret the simply-typed λ -calculus following the standard lines of its interpretation into a cartesian closed category: this sends any type A to a set [A], following the straightforward inductive definition¹

$$\begin{bmatrix} o \end{bmatrix} = \{\star\} \\ \begin{bmatrix} A \to B \end{bmatrix} = \mathcal{M}_f(\llbracket A \rrbracket) \times \llbracket B \rrbracket,$$

this set $\llbracket A \rrbracket$ is often referred to as the **web** of A. Likewise, any well-typed term $\vdash M : A$ is sent to $\llbracket M \rrbracket \subseteq \llbracket A \rrbracket$ a subset of the web. It is a reasonable intuition to think of elements of $\llbracket A \rrbracket$ as sort of *detemporalized execution traces*, and indeed it is central in this paper that they do correspond to plays in the game semantics sense where time has been suppressed.

Let us illustrate this with an example. We have

$$([([\star],\star),([\star,\star],\star)],([\star,\star],\star)) \in [\lambda fx. f(fx):(o \to o) \to o \to o]$$

$$(2.1)$$

which may be interpreted as an execution of the term (the Church integer for 2) which calls f twice. For one of these calls, f calls its argument once; the other time, twice – so the term ends up using x twice. We will revisit this example later armed with better notations.

¹The interpretation is parametrised by the choice of an interpretation for the base type. For now we use a singleton type, which is restrictive but simplifies the relationship with game semantics. In Section 4.4, we will see how to extend that if the base type is interpreted with an arbitrary set.

2.1.2. Non-idempotent intersection types. The relational model is a denotational semantics – in fact, it is the core of numerous quantitative denotational models including coherence spaces, probabilistic coherence spaces, and many others. But one of its striking features is that it can be presented purely syntactically, via an intersection type system known as non-idempotent intersection types, or multi-types.

Non-idempotent intersection types come in two shapes: on the one hand, single types are of the form \star or $\bar{\alpha} \multimap \beta$, where β is a single type and $\bar{\alpha}$ is a multi-type. On the other hand, a multi-type $\bar{\alpha}$ is a finite multiset $[\alpha_1, \ldots, \alpha_n]$ of single types. Because we refer here to the relational semantics of the simply-typed λ -calculus, we shall only consider those non-idempotent intersection types that refine -i.e., follow the structure of -a simple type:

$$\frac{\bar{\alpha} \lhd A \qquad \beta \lhd B}{\bar{\alpha} \multimap \beta \lhd A \rightarrow B} \qquad \qquad \frac{\forall i \in I \qquad \alpha_i \lhd A}{\lceil \alpha_i \mid i \in I \rceil \lhd A}$$

It looks like there is no intersection in this syntax for non-idempotent intersection types, but a finite multiset $[\alpha_1, \ldots, \alpha_n]$ can be read as a formal intersection

$$\alpha_1 \wedge \ldots \wedge \alpha_n$$

with \wedge an associative, commutative operation – but crucially, not idempotent.

It is direct to see by induction on types that refinements of a simple type A are in one-to-one correspondance with elements of $\llbracket A \rrbracket$, provided we interpret the base type as $\llbracket o \rrbracket = \{\star\}$. Accordingly, we shall from now on identify the two, and see non-idempotent intersection types as a syntax for elements of the relational model. Though we shall not rely on it in this paper, this also extends to typing rules for terms: the relational interpretation of $\vdash M : A$ is precisely the set of $\alpha \triangleleft A$ such that $\vdash M : \alpha$ is derivable – for instance,

$$\lambda f x. f (f x) : [[\star] \multimap \star, [\star, \star] \multimap \star] \multimap [\star, \star] \multimap \star$$

is the typing judgment corresponding to (2.1).

2.1.3. *Plays and refinement types.* These objects, points of the web in the relational model or non-idempotent intersection types, are at the heart of many quantitative models. Importantly for this paper, this includes (some presentations of) game semantics.

Game semantics present computation as an exchange of moves between two players: Player (+), who plays for the program under scrutiny, and Opponent (-), who plays for the execution environment. In this setting, an execution is traditionally represented as a play, a chronological sequence of moves linked with so-called *pointers* indicating their hierarchical relationships. As an example, we show in Figure 1 a play in (the strategy for) $\lambda f x. f(f x)$. It is read from top to bottom, and each move is placed under the corresponding type component. Opponent starts computation, which prompts the evaluation of f with q^+ . Then f calls its argument, which prompts the evaluation of the second occurrence of f. After that, f calls its argument twice, and x gets evaluated twice. Moves are linked with so-called justification pointers, carrying the hierarchical relationships between variable calls.

Time is critical in traditional presentations of game semantics. But it is nonetheless sensible to forget time, retaining only the tree structure induced by moves and justification pointers – as pictured in Figure 2, where the correspondence between moves and atoms in the type is conveyed via subscripts rather than the horizontal position of moves. In this time-forgetting operation lies the main intuition behind the link between game semantics



Figure 1: Example of a play



and relational semantics: Figure 2 (ignoring solid arrows \rightarrow) turns out to be an alternative representation of the non-idempotent intersection refinement

$$[[\star] \multimap \star, [\star, \star] \multimap \star] \multimap [\star, \star] \multimap \star \quad \lhd \quad (o \to o) \to o \to o$$

encountered earlier – carrying the same information as in the play in Figure 1 about the distinct variable calls and their hierarchical dependencies². Modern presentations of game semantics [Mel05, RW11, CCRW17] reject time; instead, *positions* as pictured in Figure 2 are primitive. In *concurrent games*, positions are enriched instead with *causal* wiring conveying the causal dependencies from the term, pictured with solid arrows \rightarrow in Figure 2.

2.1.4. *Rigidity and symmetries.* There is a seemingly small, but actually fundamental subtlety in the explanation above. *Positions*, those trees matching points in the relational model, are *unordered*: children of a same node which correspond to the same type component can be permuted at will – this corresponds to the fact that elements of !A in the relational model are *multisets* rather than merely lists. They are quotiented structures.

In thin concurrent games, strategies play not on these quotiented structures, but rather on some choice of concrete representatives; this is in particular required so that positions can be sensibly enriched with causal information as in Figure 2. In thin concurrent games in particular, concrete representatives of positions are called *configurations*. In these objects, distinct copies of moves are kept separate by attributing each an identifier, an integer called its *copy index*. Copy indices are not unique to a move, but two moves sharing the same justifier and corresponding to the same type component cannot have the same copy index. As an example, we draw in Figure 3 concrete representatives for the position of Figure 2. Copy indices appear in grey, to distinguish them from the subscript for the type component.

This feature of working with concrete representatives of positions is not unique to thin concurrent games; it is in fact common in *categorifications* of the relational model, such as generalized species of structure [FGHW08]. There, types are interpreted not as sets

 $^{^{2}}$ The formal link between plays in Hyland-Ong games and points of the web in the relational model was first established by Boudes [Bou09] – at least under this specific form: the correspondence between static and dynamic denotational models was first explored in [BDER97, Ehr96]. This tension between static and dynamic models underlies Melliès's work on asynchronous games [Mel05].



Figure 3: Concrete configurations with copy indices

comprising quotiented structures, but as *groupoids*: their objects are concrete representatives of non-idempotent intersection types obtained as for the relational models but with *lists*

 $\langle \alpha_1, \ldots, \alpha_n \rangle$

instead of finite multisets. In quantitative semantics, these concrete representatives of quotiented objects are often referred to as *rigid*. In this categorified situation, the quotient is replaced with explicit morphisms generated by permutations between elements of these lists.

Just like generalized species, thin concurrent games are *rigid*; and concrete configurations are related by so-called *symmetries*, tree isomorphisms which can only change copy indices. In Figure 3 we show two symmetries, which are the tree isomorphisms respecting the topological position of nodes. In thin concurrent games, the polarity of moves lets us set apart sub-groupoids of *polarized symmetries*³: some symmetries, dubbed *positive*, only change the copy index of positive moves; while others, dubbed *negative*, only reindex *negative* moves. Not every symmetry is negative or positive: the composite symmetry in Figure 3 is neither. But every symmetry factors as a composition of the two, as illustrated in Figure 3. We leave these structures on the side from now on, but they will play a crucial role later on.

Note in passing that these similarities between thin concurrent games and generalized species of structure are not superficial; the two models are connected at the bicategorical level by an interpretation-preserving (cartesian closed) pseudofunctor [COP23].

2.2. The Qualitative. Quantitative models record the multiplicity of resource usage. This is their defining feature, and a significant part of their remarkable ability to handle quantitative effects such as probabilistic choice [DE11], quantum primitives [PSV14] and others [LMMP13]. But most denotational models – starting with Scott domains [Plo77] – do not record such quantitative information: they are qualitative; they record the presence of resource calls but not their multiplicity. They correspond to a different notion of approximation according to which a function is greater than another if it produces more using less information about the input, regardless of how many times the input is evaluated.

2.2.1. Idempotent intersection types. In terms of intersection types, being qualitative means

$$\alpha \wedge \alpha = \alpha$$
,

i.e. that \wedge is *idempotent*: an expression $\alpha_1 \wedge \ldots \wedge \alpha_n$ no longer corresponds to the finite multiset $[\alpha_1, \ldots, \alpha_n]$, but to the set $\{\alpha_1, \ldots, \alpha_n\}$. But brutally enforcing $\alpha \wedge \alpha = \alpha$ in this

³This fundamental distinction was identified for the first time in Melliès's *orbital games* [Mel03]; it is also at the core of the more recent setting of *thin spans of groupoids* [CF23].

way fails: the finite powerset endofunctor on **Rel** fails to be a comonad, as the candidate co-unit fails naturality. For instance, with $A = \{a_1, a_2\}, B = \{b\}$, with $R = \{(a_1, b), (a_2, b)\}$,

$$(\{a_1, a_2\}, b) \in \operatorname{der}_B \circ ! R, \qquad (\{a_1, a_2\}, b) \notin R \circ \operatorname{der}_A.$$

Intuitively, the issue is that seeing b may correspond to one occurrence, or to many. If it is really just one, then we get it through either $\{a_1\}$ or $\{a_2\}$ but not through $\{a_1, a_2\}$. However, if there are actually several "merged" occurrences of b, then some may arise through a_1 and others through a_2 , so we really do need $\{a_1, a_2\}$. To sort out this confusion, we must relax our intuition about resources: while in the relational model, $(a, b) \in \llbracket M \rrbracket$ means that there is an execution where M consumes *exactly* a to produce b, here we must allow M to produce b or more using a or less. In idempotent intersection types, this means that it is unavoidable to consider a *partial order* on types, *i.e.* a *subtyping relation*. By insisting that R is adequately down-closed for this subtyping relation, one reinstates naturality.

It is then fairly natural to consider "idempotent" intersection types as simply the addition of a preorder on non-idempotent intersection types, defined by the rules below:

$$\frac{\bar{\alpha}_{2} \leq \bar{\alpha}_{1}}{\star \leq \star} \qquad \frac{\bar{\alpha}_{2} \leq \bar{\alpha}_{1}}{\bar{\alpha}_{1} \multimap \beta_{1} \leq \bar{\alpha}_{2} \multimap \beta_{2}} \qquad \frac{\forall i \in I \; \exists j \in J \; \alpha_{i} \leq \beta_{j}}{[\alpha_{i} \mid i \in I] \leq [\beta_{j} \mid j \in J]}$$
(2.2)

noting that the order is contravariant on the left hand side for the arrow. Those may not look idempotent as they are still based on multisets; but it follows that for all $\bar{\alpha}$,

$$\bar{\alpha} \wedge \bar{\alpha} \le \bar{\alpha} \,, \qquad \qquad [] \le \bar{\alpha}$$

which reminds us of the logical laws of contraction and weakening. In particular, $\bar{\alpha} \leq \bar{\alpha} \wedge [] \leq \bar{\alpha} \wedge \bar{\alpha}$. Hence, though idempotence is not enforced primitively, it is derived as the equivalence relation generated by the subtyping preorder⁴. This view on idempotent intersection types is implicit in the *linear Scott model*, the linear decomposition of Scott domains discovered by Huth [Hut93] and Winskel [Win98]; and in particular in Winskel's presentation⁵.

The heart of Ehrhard's celebrated *extensional collapse* theorem [Ehr12] is then that in this language, the interpretation of a simply-typed λ -term in the linear Scott model is simply the *down-closure* of its relational interpretation, linking the qualitative and the quantitative.

2.2.2. Cartesian maps. Now, as we discussed earlier, the set of non-idempotent intersection types may be categorified into a groupoid of rigid intersection types, as formalized in the cartesian closed bicategory of generalized species of structure; or into a groupoid of configurations in the sense of thin concurrent games. It is natural that the preorder \leq should then be refined into a category: this is achieved, for instance, through the cartesian closed bicategory of cartesian distributors [Oli21]. There, given a category A a morphism in !A

$$\langle \alpha_i \mid i \in I \rangle \ \to \ \langle \beta_j \mid j \in J \rangle$$

consists in a function $h: I \to J$, together with a family $(f_i)_{i \in I}$ where $f_i: \alpha_i \to \beta_{f(i)}$ in A– we call this a *cartesian* morphism, as it can *contract* several copies together, and it can weaken copies on the right hand side by not giving them a pre-image. A morphism from $\alpha_1 \to \beta_1$ to $\alpha_2 \to \beta_2$ consists in morphisms $f: \alpha_2 \to \alpha_1$ and $g: \beta_1 \to \beta_2$, reflecting (2.2), contravariantly on the left hand side.

⁴One can equivalently work with non-idempotent intersection types equipped with the subtyping preorder, or with idempotent intersection types equipped with the induced partial order.

 $^{^{5}}$ The preordered set of intersection types is not the order of Scott domains, which may be obtained by considering the complete lattice of down-closed sets – see Section 5.1.3.

In this paper, we achieve an analogous categorification in thin concurrent games, turning the groupoid of configurations in thin concurrent games into a *category* of configurations. So, what should be the adequate morphisms between configurations? Drawing inspiration from symmetries and the contraction maps above, a natural guess is that they should simply be forest morphisms which preserve the type component. For instance, we could have



contracting all copies down to copy index 0. But this is not right, because it does not take into account the contravariance on the left hand side of arrows. The missing ingredient is to account for *polarity* – *negative* contraction maps can only contract and weaken negative moves, while *positive* contraction maps can only contract and weaken positive moves:



and *cartesian morphisms* between configurations are obtained as relational compositions

 $\langle \overline{a} \overline{b} \rangle = \overline{b} \langle \overline{a}, \overline{c} \rangle \langle \overline{a}, \overline{c} \rangle \langle \overline{$

which are therefore no longer forest morphisms, but *do* induce the adequate preorder on configurations to match the linear Scott model, as we shall demonstrate in this paper.

2.2.3. Cartesian matching problems. The main contribution of this paper is to extend this into a structure-preserving functor from a category of thin concurrent games into the linear Scott model. This builds on earlier results on the relational collapse of thin concurrent games: informally, a strategy $\sigma : A \vdash B$ from A to B is an aggregate of states x^{σ} , with

$$x_A^{\sigma} \quad \leftrightarrow \quad x^{\sigma} \quad \mapsto \quad x_B^{\sigma}$$

projections obtained – ignoring the issue of taking symmetry classes – by simply forgetting the causal arrows \rightarrow displayed *e.g.* in Figure 2. The *relational collapse* of σ then gathers all pairs $(x_A^{\sigma}, x_B^{\sigma})$. But to reach the linear Scott model, we need to build a relation that is *down-closed*! We shall achieve this by sending σ to all pairs (y_A, y_B) such that we have

$$y_A \quad \stackrel{+-}{\longleftrightarrow} \quad x_A^{\sigma} \quad \leftarrow \quad x^{\sigma} \quad \mapsto \quad x_B^{\sigma} \quad \stackrel{-+}{\longleftrightarrow} \quad y_B \,,$$

i.e. simply the down-closure with respect to $\stackrel{+}{\longleftrightarrow}$. This does yield a valid morphism in linear Scott models. But this leaves us with the demanding task to show that this down-closure remains functorial. And this means that for $\sigma : A \vdash B$ and $\tau : B \vdash C$, given

$$x^{\sigma} \mapsto x^{\sigma}_B \quad \stackrel{\leftarrow}{\leftrightarrow} \quad x_B \quad \stackrel{\leftarrow}{\leftrightarrow} \quad x^{\tau}_B \quad \leftarrow \quad x^{\tau} ,$$

i.e. $(x^{\sigma}, x_B^{\sigma} \stackrel{\leftarrow}{\longrightarrow} x_B^{\tau}, x^{\tau})$, we must find a synchronization $z^{\tau \odot \sigma}$ in $\tau \odot \sigma$ whose (down-closure of the) projection on A, C is the same. An analogous property is necessary with respect to

symmetries to construct thin concurrent games [Cla24, Proposition 7.4.4]. But here, the situation is significantly more complex: both x^{σ} and x^{τ} are trying to duplicate and erase each other, and we must find a satisfactory state where all these duplications and erasures are satisfied – we call this a *cartesian matching problem*.

In game semantics we approach the question concretely, and provide a combinatorial argument to resolve such matchings. This is the crux of this paper; once this is solved it is not hard to provide an interpretation-preserving functor from to the linear Scott model.

2.3. Outline. This sums up the main thrust of the paper; of course there are various technical subtleties that come blur the waters. First, we must introduce thin concurrent games, along with their relational collapse. This already comes with a significant technical set-up on top of thin concurrent games: typically, those configurations that match points in the relational model must be identified via a *payoff* mechanism. One must also introduce the slightly unorthodox concept of a *relative Seely* (\sim -)*categories*, a weakening of Seely categories, as the structure of plain Seely categories is not preserved by the relational collapse. This content is well-covered in other sources [Cla24, COP23], which our presentation follows.

We must then give concrete definitions for cartesian morphisms. Unfortunately, we can only do that relying on a fairly concrete description of the games considered, referring explicitly to copy indices. For this we must import from [Cla24] the rather clunky notion of *mixed board*. Though this paper focuses on the semantic structures, we will apply those to obtain results on the simply-typed λ -calculus. We consider the λ -calculus with one base type o – our results of course apply in the presence of many base types, but for this paper we estimate that the additional notational burden is not worth the extra generality.

Concretely, the paper is organized as follows. In Section 3, we recall the main definitions of thin concurrent games. In Section 4 we describe the *relational collapse* of thin concurrent games – the material is mainly taken from [Cla24], with the extra development required in order to allow interpreting the atom as an arbitrary set. They main thrust of the contributions start in Section 5: there, we refine our games to allow the collapse to the linear Scott model. In particular, we introduce the notion of *cartesian morphism* on a mixed board, and develop some of their basic properties. Finally, in Section 6 we show how to solve cartesian matching problems, and derive our main results.

3. Thin Concurrent Games

The goal of this section is to give an introduction to thin concurrent games, geared towards its relational collapse: we wish to present the situation in which this collapse is the most natural, which we may then specialize in later parts of this paper.

Though those are well-understood notions, setting up all the required infrastructure for the collapse is demanding – the reader can find a more detailed introduction in [Cla24].

3.1. **Basic Concurrent Games.** The framework of *concurrent games* [MM07, FP09, RW11] is not merely a game semantics for concurrency – though it can serve that purpose – but a deep reworking of the basic mechanisms of game semantics using causal "truly concurrent" structures from concurrency theory [NPW79], which we must first introduce.

3.1.1. *Event structures.* Concurrent games and strategies are based on event structures. An event structure represents the behaviour of a system as a set of possible computational events equipped with dependency and incompatibility constraints.

Definition 3.1. An event structure (es) is $E = (|E|, \leq_E, \#_E)$, where |E| is a (countable) set of events, \leq_E is a partial order called **causal dependency** and $\#_E$ is an irreflexive symmetric binary relation on |E| called **conflict**, satisfying:

finite causes: $\forall e \in |E|, \ [e]_E = \{e' \in |E| \mid e' \leq_E e\}$ is finite, vendetta: $\forall e_1 \ \#_E \ e_2, \ \forall e_2 \leq_E e'_2, \ e_1 \ \#_E \ e'_2.$

Operationally, an event can occur if all its dependencies are met, and no conflicting events have occurred. A finite set $x \subseteq_f |E|$ down-closed for \leq_E and comprising no conflicting pair is called a **configuration** – we write $\mathscr{C}(E)$ for the set of configurations on E, naturally ordered by inclusion. If $x \in \mathscr{C}(E)$ and $e \in |E|$ is such that $e \notin x$ but $x \cup \{e\} \in \mathscr{C}(E)$, we say that e is **enabled** by x and write $x \vdash_E e$. For $e_1, e_2 \in |E|$ we write $e_1 \rightarrow_E e_2$ for the **immediate causal dependency**, *i.e.* $e_1 <_E e_2$ with no event strictly in between. Finally, two events $e_1, e_2 \in |E|$ are in **immediate conflict**, written $e_1 \sim_E e_2$, if $e_1 \#_E e_2$, and this conflict is not inherited: if $e'_1 < e_1$ then $\neg(e_1 \#_E e_2)$, and likewise on the other side.

There is an accompanying notion of map: a **map of event structures** from E to F is a function $f : |E| \to |F|$ such that: (1) for all $x \in \mathscr{C}(E)$, the direct image $fx \in \mathscr{C}(F)$; and (2) for all $x \in \mathscr{C}(E)$ and $e, e' \in x$, if fe = fe' then e = e'. There is a category **ES** of event structures and maps.

3.1.2. Games and strategies. Throughout this paper, we will gradually refine our notion of game. For now, a **plain game** is simply an event structure A together with a **polarity** function $\text{pol}_A : |A| \to \{-, +\}$ which specifies, for each event $a \in A$, whether it is **positive** (*i.e.* due to Player / the program) or **negative** (*i.e.* due to Opponent / the environment). Events are often called **moves**, and annotated with their polarity.

A strategy is an event structure with a projection map to A:

Definition 3.2. Consider A a plain game. A strategy on A, written $\sigma : A$, is an event structure σ together with a map $\partial_{\sigma} : \sigma \to A$ called the **display map**, satisfying:

- (1) for all $x \in \mathscr{C}(\sigma)$ and $\partial_{\sigma} x \vdash_A a^-$, there is a unique $x \vdash_{\sigma} s$ such that $\partial_{\sigma} s = a$.
- (2) for all $s_1 \twoheadrightarrow_{\sigma} s_2$, if $\operatorname{pol}_A(\partial_{\sigma}(s_1)) = +$ or $\operatorname{pol}_A(\partial_{\sigma}(s_2)) = -$, then $\partial_{\sigma}(s_1) \twoheadrightarrow_A \partial_{\sigma}(s_2)$.

There two conditions (called *receptivity* and *courtesy*) ensure that the strategy does not constrain the behaviour of Opponent any more than the game does. They are essential for the compositional structure we describe below, but they do not play a major role in this paper (their use is encapsulated in technical lemmas and propositions proved elsewhere). Note also that though a strategy does not come with a polarity function for the moves in σ , they do inherit a polarity through ∂_{σ} . This is used implicitly from now on.

As a simple example, the usual game \mathbb{B} for booleans in call-by-name is



drawn from top to bottom (Player moves are blue, and Opponent moves are red): Opponent initiates computation with the first move \mathbf{q} , to which Player can react with either \mathbf{tt} or \mathbf{ff} .

Strategies give a "proof-relevant" account of execution, in the sense that moves and configurations of the game can have multiple witnesses in the strategy. For example, on the left below, b and c are both mapped to the same move **tt**:



We denote immediate causality by \rightarrow in strategies, and by dotted lines for games – this lets us represent the strategy in a single diagram, as on the right above. Similar diagrams may represent not entire games and strategies but *configurations* of games and strategies, which implicitly inherit a partial order from the ambiant event structure.

3.1.3. Morphisms between strategies. For σ and τ two strategies on A, a morphism from σ to τ , written $f : \sigma \Rightarrow \tau$, is a map of event structures $f : \sigma \to \tau$ preserving the dependency relation \leq (we say it is **rigid**) and such that $\partial_{\tau} \circ f = \partial_{\sigma}$.

3.1.4. *+-covered configurations.* We now describe a useful technical tool: it turns out that a strategy is completely characterized by a subset of its configurations, called *+-*covered.

For a strategy σ on a game A, a configuration $x \in \mathscr{C}(\sigma)$ is +-covered if all its maximal events are positive, so every Opponent move in x has at least one Player successor. We write $\mathscr{C}^+(\sigma)$ for the partially ordered set (by inclusion) of +-covered configurations of σ .

Lemma 3.3. Consider a plain game A, and strategies $\sigma, \tau : A$.

If $f : \mathscr{C}^+(\sigma) \cong \mathscr{C}^+(\tau)$ is an order-isomorphism such that $\partial_{\tau} \circ f = \partial_{\sigma}$, then there is a unique isomorphism of strategies $\hat{f} : \sigma \cong \tau$ such that for all $x \in \mathscr{C}^+(\sigma)$, $\hat{f}(x) = f(x)$.

Proof. Immediate consequence of [Cla24, Lemma 6.3.4].

This is the first hint of a methodology that is central to this paper: in concurrent games, we rarely reason at the level of individual events, preferring whenever possible to reason at the level of configurations, especially when linking with relational-like models.

3.2. A \sim -category of concurrent games and strategies. We now show how games and strategies are organized into a \sim -category – that is, a bicategory where 2-cells are degenerated so that each hom-set forms a setoid, a set with an equivalence relation.

3.2.1. Strategies between games. If A is a plain game, its **dual** A^{\perp} has the same components as A except for the reversed polarity. In particular $\mathscr{C}(A) = \mathscr{C}(A^{\perp})$. The **parallel composition** $A \parallel B$ of A and B is simply A and B side by side, with no interaction – its events are the tagged disjoint union $|A \parallel B| = |A| + |B| = \{1\} \times |A| \uplus \{2\} \times |B|$, and other components are inherited. Likewise, the **hom** $A \vdash B$ is simply defined as $A^{\perp} \parallel B$. We write $x_A \parallel x_B$

for the configuration of $A \otimes B$ that has $x_A \in \mathscr{C}(A)$ on the left and $x_B \in \mathscr{C}(B)$ on the right, and likewise for $x_A \vdash x_B \in \mathscr{C}(A \vdash B)$, informing order-isomorphisms⁶

$$-\parallel - : \mathscr{C}(A) \times \mathscr{C}(B) \cong \mathscr{C}(A \parallel B),$$
 (3.1)

$$- \vdash - : \mathscr{C}(A) \times \mathscr{C}(B) \cong \mathscr{C}(A \vdash B).$$
 (3.2)

A strategy from A to B is a strategy on the game $A \vdash B$. Note that if $\sigma : A \vdash B$ and $x^{\sigma} \in \mathscr{C}(\sigma)$, by convention we write $\partial_{\sigma}(x^{\sigma}) = x_A^{\sigma} \vdash x_B^{\sigma} \in \mathscr{C}(A \vdash B)$.

Our first example of a strategy between games is **copycat** $\mathbf{c}_A : A \vdash A$, the identity morphism on A in our \sim -category. Concretely, copycat on A has the same events as $A \vdash A$, but adds immediate causal links between copies of the same move across components. By Lemma 3.3, the following characterizes copycat up to isomorphism.

Proposition 3.4. If A is a game, there is an order-isomorphism

$$\mathbf{c}_{(-)} \quad : \quad \mathscr{C}(A) \cong \mathscr{C}^+(\mathbf{c}_A)$$

such that for all $x \in \mathscr{C}(A)$, $\partial_{\mathbf{c}_A}(\mathbf{c}_x) = x \vdash x$.

Proof. Follows from [Cla24, Lemma 6.4.4].

This shows that the copycat strategy is essentially the diagonal relation, which is the first hint of the connection between concurrent games and the relational model.

3.2.2. Composition. Consider $\sigma : A \vdash B$ and $\tau : B \vdash C$. We define their composition $\tau \odot \sigma : A \vdash C$. Concurrent games are a dynamic model, and to successfully synchronize, σ and τ must agree to play the same events *in the same order*; this is defined in two steps.

We say that configurations $x^{\sigma} \in \mathscr{C}(\sigma)$ and $x^{\tau} \in \mathscr{C}(\tau)$ are **matching** if they reach the same configuration on B, *i.e.* $x_B^{\sigma} = x_B^{\tau} = x_B$. If that is the case, it induces a synchronization (and we may then ask if that synchronization induces a deadlock). If all events of x^{σ} and x^{τ} were in B, this synchronization would take the form of a bijection $x^{\sigma} \simeq x^{\tau}$. But some moves of x^{σ} are in A and some moves of x^{τ} are in C, so instead we form the bijection

$$\varphi[x^{\sigma}, x^{\tau}] : x^{\sigma} \parallel x_{C}^{\tau} \stackrel{\partial_{\sigma} \parallel x_{C}^{\tau}}{\simeq} x_{A}^{\sigma} \parallel x_{B} \parallel x_{C}^{\tau} \stackrel{x_{A}^{\sigma} \parallel \partial_{\tau}^{-1}}{\simeq} x_{A}^{\sigma} \parallel x_{C}^{\tau}$$

where $x \parallel y$ is the tagged disjoint union. This uses the fact that from the conditions on maps of event structures, $\partial_{\sigma} : x^{\sigma} \simeq x_A^{\sigma} \vdash x_B^{\sigma}$ is a bijection and likewise for ∂^{τ} .

Next, we import the causal constraints of σ and τ to (the graph of) $\varphi[x^{\sigma}, x^{\tau}]$, via:

$$\begin{array}{lll} (m,n) \triangleleft_{\sigma} (m',n') & \Leftrightarrow & m <_{\sigma \parallel C} m \\ (m,n) \triangleleft_{\tau} (m',n') & \Leftrightarrow & n <_{A \parallel \tau} n' \end{array}$$

letting us finally say that matching x^{σ} and x^{τ} are **causally compatible** if $\triangleleft = \triangleleft_{\sigma} \cup \triangleleft_{\tau}$ on (the graph of) $\varphi[x^{\sigma}, x^{\tau}]$ is acyclic. In particular, x^{σ} and x^{τ} in Figure 4 are *not* causally compatible, the synchronization induces a *deadlock*.

The **composition** of σ and τ is the unique (up to iso) strategy whose +-covered configurations are essentially causally compatible pairs of +-covered configurations. Writing $\mathbf{CC}(\sigma, \tau)$ for causally compatible pairs $(x^{\sigma}, x^{\tau}) \in \mathscr{C}^+(\sigma) \times \mathscr{C}^+(\tau)$ (ordered componentwise):

⁶Throughout this paper, we write \simeq for mere bijections, and \cong for isomorphisms also preserving structure.



Figure 4: An example of matching but causally incompatible configurations, in the composition of $\sigma : \mathbb{U} \to \mathbb{U}$ and $\tau : \mathbb{U} \to \mathbb{U} \vdash \mathbb{N}$. The underlying games are left undefined, but can be recovered by removing the arrows \rightarrow . The configurations are matching on $\mathbb{U} \to \mathbb{U}$, but the arrows \rightarrow impose incompatible orders (i.e. a cycle) between the two occurrences of \checkmark .

Proposition 3.5. Consider strategies $\sigma : A \vdash B$ and $\tau : B \vdash C$. There is a strategy $\tau \odot \sigma : A \vdash C$, unique up to isomorphism, with an order-isomorphism

$$-\odot - : \mathbf{CC}(\sigma, \tau) \cong \mathscr{C}^+(\tau \odot \sigma)$$

s.t. for all $x^{\sigma} \in \mathscr{C}^+(\sigma)$ and $x^{\tau} \in \mathscr{C}^+(\tau)$ causally compatible, $\partial_{\tau \odot \sigma}(x^{\tau} \odot x^{\sigma}) = x^{\sigma}_A \vdash x^{\tau}_C$.

Proof. See [Cla24, Proposition 6.2.1].

This description of composition emphasizes the conceptual difference between a static model, in which composition is based merely on matching pairs, and a dynamic model, based on causal compatibility and sensitive to deadlocks. We get [Cla24, Theorem 6.4.11]:

Theorem 3.6. There is a \sim -category CG with: (1) objects, plain games; (2) morphisms from A to B, strategies $\sigma : A \vdash B$; and equivalence relation, isomorphism of strategies.

3.3. Adding Symmetry. The ambiant ~-category in which this paper takes place is not quite CG, but a refinement sensitive to symmetry – this is necessary so that the model supports an exponential modality. We now go from CG to TCG by replacing the set of configurations $\mathscr{C}(A)$ with a groupoid of configurations $\mathscr{S}(A)$ whose morphisms are chosen bijections called symmetries, that behave well with respect to the causal order.

3.3.1. Event structures with symmetry. Our starting point is to replace event structures with event structures with symmetry, due to Winskel [Win07]:

Definition 3.7. An isomorphism family on es E is a groupoid $\mathscr{S}(E)$ having as objects all configurations, and as morphisms certain bijections between configurations, satisfying:

 $\begin{array}{ll} \textit{restriction:} & \text{for all } \theta : x \simeq y \in \mathscr{S}(E) \text{ and } x \supseteq x' \in \mathscr{C}(E), \\ & \text{there is } \theta \supseteq \theta' \in \mathscr{S}(E) \text{ such that } \theta' : x' \simeq y'. \\ \textit{extension:} & \text{for all } \theta : x \simeq y \in \mathscr{S}(E), \ x \subseteq x' \in \mathscr{C}(E), \\ & \text{there is } \theta \subseteq \theta' \in \mathscr{S}(E) \text{ such that } \theta' : x' \simeq y'. \end{array}$

We call $(E, \mathscr{S}(E))$ an event structure with symmetry (ess).

We refer to morphisms in $\mathscr{S}(E)$ as symmetries, and write $\theta : x \cong_E y$ if $\theta : x \simeq y$ with $\theta \in \mathscr{S}(E)$. The **domain** dom(θ) of $\theta : x \cong_E y$ is x, and likewise its **codomain** cod(θ) is y. A map of ess $E \to F$ is a map of event structures that preserves symmetry: the bijection

 $f\theta \stackrel{\text{def}}{=} fx \stackrel{f^{-1}}{\simeq} x \stackrel{\theta}{\simeq} y \stackrel{f}{\simeq} fy,$

is in $\mathscr{S}(F)$ for every $\theta: x \cong_E y$ (recall that f restricted to any configuration is bijective). This exactly amounts to making $f : \mathscr{S}(E) \to \mathscr{S}(F)$ a functor of groupoids.

We can define a 2-category **ESS** of ess, maps of ess, and natural transformations between the induced functors. For $f, g: E \to F$ such a natural transformation is necessarily unique [Win07], and corresponds to the fact that for every $x \in \mathscr{C}(E)$ the composite bijection

$$fx \stackrel{f^{-1}}{\simeq} x \stackrel{g}{\simeq} gx$$

via local injectivity of f and q, is in $\mathscr{S}(F)$. So this is an equivalence, denoted $f \sim q$.

3.3.2. Thin games. We define games with symmetry. To match the polarized structure, a game is an ess with two sub-symmetries, one for each player (see e.g. [Mel03, CCW19, Paq22]).

Definition 3.8. A thin concurrent game (tcg) is a game A with isomorphism families $\mathscr{S}(A), \mathscr{S}_+(A), \mathscr{S}_-(A)$ s.t. $\mathscr{S}_+(A), \mathscr{S}_-(A) \subseteq \mathscr{S}(A)$, symmetries preserve polarity, and

- (1) if $\theta \in \mathscr{S}_+(A) \cap \mathscr{S}_-(A)$, then $\theta = \mathrm{id}_x$ for $x \in \mathscr{C}(A)$,
- (2) if $\theta \in \mathscr{S}_{-}(A), \ \theta \subseteq^{-} \theta' \in \mathscr{S}(A)$, then $\theta' \in \mathscr{S}_{-}(A)$, (3) if $\theta \in \mathscr{S}_{+}(A), \ \theta \subseteq^{+} \theta' \in \mathscr{S}(A)$, then $\theta' \in \mathscr{S}_{+}(A)$,

where $\theta \subset^p \theta'$ is $\theta \subset \theta'$ with (pairs of) events of polarity p.

Elements of $\mathscr{S}_+(A)$ (resp. $\mathscr{S}_-(A)$) are called **positive** (resp. **negative**); they intuitively correspond to symmetries carried by positive (resp. negative) moves, introduced by Player (resp. Opponent). We write $\theta : x \cong_A^- y$ (resp. $\theta : x \cong_A^+ y$) if $\theta \in \mathscr{S}_-(A)$ (resp. $\theta \in \mathscr{S}_+(A)$).

Each symmetry has a unique positive-negative factorization [Cla24, Lemma 7.1.18]:

Lemma 3.9. Consider A a tcg and $\theta : x \cong_A z$ a symmetry. Then, there are unique $y \in \mathscr{C}(A)$, $\theta_- : x \cong_A^- y$ and $\theta_+ : y \cong_A^+ z$ s.t. $\theta = \theta_+ \circ \theta_-$.

We extend with symmetry the basic constructions on games: the **dual** A^{\perp} has the same symmetries as A, but $\mathscr{S}_+(A^{\perp}) = \mathscr{S}_-(A)$ and $\mathscr{S}_-(A^{\perp}) = \mathscr{S}_+(A)$; the **parallel composition** $A_1 \parallel A_2$ has symmetries those $\theta_1 \parallel \theta_2 : x_1 \parallel x_2 \cong_{A_1 \parallel A_2} y_1 \parallel y_2$, where each $\theta_i : x_i \cong_{A_i} y_i$, and similarly for positive and negative symmetries; the hom $A \vdash B$ is $A^{\perp} \parallel B$.

3.3.3. Thin strategies. We now define strategies on thin concurrent games:

Definition 3.10. Consider A a tcg.

A strategy on A, written $\sigma: A$, is an ess σ equipped with a morphism of ess $\partial_{\sigma}: \sigma \to A$ forming a strategy in the sense of Definition 3.2, and such that:

- (1) if $\theta \in \mathscr{S}(\sigma), \partial_{\sigma}\theta \vdash_A (a^-, b^-)$, there are unique $\theta \vdash_{\sigma} (s, t)$ s.t. $\partial_{\sigma}s = a$ and $\partial_{\sigma}t = b$.
- (2) if $\theta : x \cong_{\sigma} y$ is such that $\partial_{\sigma} \theta \in \mathscr{S}_{+}(A)$, then x = y and $\theta = \mathrm{id}_{x}$.

As before, a strategy from A to B is a strategy on $\sigma : A \vdash B$.

The first condition forces σ to acknowledge Opponent symmetries in A; the notation $\theta \vdash_A (a, b)$ means $(a, b) \notin \theta$ and $\theta \cup \{(a, b)\} \in \mathscr{S}(A)$. The second condition is **thinness**: it means that any non-identity symmetry in the strategy must originate from Opponent.

3.3.4. $A \sim$ -category. The composition of thin strategies $\sigma : A \vdash B$ and $\tau : B \vdash C$ is obtained by equipping $\tau \odot \sigma$ (Proposition 3.5) with an adequate isomorphism family.

If $\mathscr{S}^+(\sigma)$ is the restriction of $\mathscr{S}(\sigma)$ to +-covered configurations, then we can write $\mathbf{CC}(\mathscr{S}^+(\sigma), \mathscr{S}^+(\tau))$ for the pairs $(\varphi^{\sigma}, \varphi^{\tau})$ of symmetries which are matching, i.e. $\varphi^{\sigma}_B = \varphi^{\tau}_B$ and whose domain (or equivalently, codomain) are causally compatible.

Proposition 3.11. Consider $\sigma : A \vdash B$ and $\tau : B \vdash C$ thin strategies.

There is a unique symmetry on $\tau \odot \sigma$ with a bijection commuting with dom and cod

$$(-\odot -): \mathbf{CC}(\mathscr{S}^+(\sigma), \mathscr{S}^+(\tau)) \simeq \mathscr{S}^+(\tau \odot \sigma)$$

and compatible with display maps, i.e. $(\varphi^{\tau} \odot \varphi^{\sigma})_A = \varphi^{\sigma}_A$ and $(\varphi^{\tau} \odot \varphi^{\sigma})_C = \varphi^{\tau}_C$.

Proof. This follows from [Cla24, Proposition 7.3.1].

It can be checked that this makes $\tau \odot \sigma : A \vdash C$ a thin strategy as required. In order to form a \sim -category, it is necessary to give the adequate equivalence relation between thin strategies. For this, recall first the 2-dimensional structure in **ESS**, given by the equivalence relation \sim on morphisms (Section 3.3.1). For two maps $f, g : E \to A$ into a tcg, we write $f \sim^+ g$ if $f \sim g$ and for every $x \in \mathscr{C}(E)$ the symmetry obtained as the composition

$$f x \stackrel{f^{-1}}{\simeq} x \stackrel{g}{\simeq} g x$$

witnessing $f \sim g$ for x, is positive. This lets us give the next definition:

Definition 3.12. Let $\sigma, \tau : A \vdash B$ be thin strategies. A **positive morphism of strategies** from σ to τ is a rigid map of ess $f : \sigma \to \tau$ such that $\partial_{\tau} \circ f \sim^+ \partial_{\sigma}$. We write $f : \sigma \Rightarrow \tau$ to mean that f is a positime morphism from σ to τ .

A positive isomorphism $f : \sigma \cong \tau$ is an invertible (on the nose) positive morphism.

As a convention, if f is a 2-cell as above, for $x^{\sigma} \in \mathscr{C}(\sigma)$ we write

$$f[x^{\sigma}]$$
 : $\partial_{\sigma} x^{\sigma} \cong^{+}_{A \vdash B} \partial_{\tau}(f x^{\sigma})$

for the positive symmetry witnessing this, which may be decomposed into two symmetries on the two sides, $f[x^{\sigma}]_A : x_A^{\sigma} \cong_A^- (f x^{\sigma})_A$ and $f[x^{\sigma}]_B : x_B^{\sigma} \cong_B^+ (f x^{\sigma})_B$.

Positive isomorphism will provide the equivalence relation for the \sim -categorical structure of thin concurrent games. A crucial challenge in constructing this \sim -category is then to ensure that positive isomorphism is preserved under composition. This demands in particular, given $f: \sigma \Rightarrow \sigma': A \vdash B$ and $g: \tau \Rightarrow \tau': B \vdash C$, to form a *horizontal composition*

$$g \odot f$$
 : $\tau \odot \sigma \Rightarrow \tau' \odot \sigma'$: $A \vdash C$

which requires us to transport $x^{\tau} \odot x^{\sigma} \in \mathscr{C}^+(\tau \odot \sigma)$ to $\mathscr{C}^+(\tau' \odot \sigma')$ via f and g. However, the issue is that $f(x^{\sigma})$ and $g(x^{\tau})$ may not be matching: the hypotheses at hand only yield

$$g[x^{\tau}]_B \circ f[x^{\sigma}]_B^{-1} \quad : \quad f(x^{\sigma})_B \cong_B g(x^{\tau})_B$$

a mediating symmetry – hence to achieve our goals, we use [Cla24, Proposition 7.4.4]:

Proposition 3.13. Consider $x^{\sigma} \in \mathscr{C}^+(\sigma), \theta_B : x_B^{\sigma} \cong_B x_B^{\tau}, x^{\tau} \in \mathscr{C}^+(\tau)$ causally compatible, i.e. the relation \triangleleft induced on the graph of the composite bijection

 $x^{\sigma} \parallel x_{C}^{\tau} \quad \stackrel{\partial_{\sigma} \parallel x_{C}^{\tau}}{\simeq} \quad x_{A}^{\sigma} \parallel x_{B}^{\sigma} \parallel x_{C}^{\tau} \quad \stackrel{x_{A}^{\sigma} \parallel \theta \parallel x_{C}^{\tau}}{\simeq} \quad x_{A}^{\sigma} \parallel x_{B}^{\tau} \parallel x_{C}^{\tau} \quad \stackrel{x_{A}^{\sigma} \parallel \partial_{\tau}^{-1}}{\simeq} \quad x_{A}^{\sigma} \parallel x_{T}^{\tau}$

by $<_{\sigma \parallel C}$ and $<_{A \parallel \tau}$ as in §3.2.2, is acyclic – we also say the composite bijection is **secured**. Then, there are unique $y^{\tau} \odot y^{\sigma} \in \mathscr{C}^+(\tau \odot \sigma)$ with symmetries $\varphi^{\sigma} : x^{\sigma} \cong_{\sigma} y^{\sigma}$ and $\varphi^{\tau} : x^{\tau} \cong_{\tau} y^{\tau}$, such that $\varphi^{\sigma}_A \in \mathscr{S}_-(A)$ and $\varphi^{\tau}_C \in \mathscr{S}_+(C)$, and $\varphi^{\tau}_B \circ \theta = \varphi^{\sigma}_B$.

Altogether, this allows us to construct the desired \sim -category:

Theorem 3.14. There is a \sim -category **TCG** with: (1) objects, thin concurrent games; (2) morphisms, strategies $\sigma : A \vdash B$; (3) equivalence, positive isomorphism.

3.4. **Boards.** In this line of work connecting concurrent games with relational-like models, a difficulty is that *points* in the sense of the relational model are not *all* configurations, but only some of them. This means that following the approach first outlined by Melliès [Mel05] and adapted to concurrent games in earlier work (see *e.g.* [Cla24]), we must enrich tcgs with structure allowing us to identify those configurations that are *stopping*, in the sense that they correspond to points in the relational model.

We now introduce boards along with useful constructions on them.

Definition 3.15. A board is a tcg A along with $\kappa_A : \mathscr{C}(A) \to \{-1, 0, +1\}$ a payoff function, such that this data satisfies the following conditions:

invariant: for all $\theta : x \cong_A y$, we have $\kappa_A(x) = \kappa_A(y)$, *race-free:* for all $a \sim A a'$, we have $\operatorname{pol}_A(a) = \operatorname{pol}_A(a')$. *forestial:* for all $a_1, a_2, a \in A$, if $a_1, a_2 \leq_A a$, then $a_1 \leq_A a_2$ or $a_2 \leq_A a_1$, *alternating:* for all $a_1, a_2 \in A$, if $a_1 \Rightarrow_A a_2$, then $\operatorname{pol}_A(a_1) \neq \operatorname{pol}_A(a_2)$,

A –-board must also satisfy the following two additional conditions:

negative: for all a minimal in A, $\text{pol}_A(a) = -$, initialized: $\kappa_A(\emptyset) \ge 0$.

Finally, a –-board A is strict if $\kappa_A(\emptyset) = 1$ and all its initial moves are in pairwise conflict. It is well-opened if it is strict with exactly one initial move.

The payoff function κ_A assigns a value to each configuration. Configurations x with payoff 0 are called **complete**, written $x \in \mathscr{C}^0(A)$: those correspond to points in the relational model. Otherwise, κ_A assigns a responsibility for why a configuration is non-complete: if $\kappa_A(x) = -1$ then Player is responsible, otherwise it is Opponent.

We recall a few constructions on boards. The objects of our forthcoming category will be --boards. The first basic --boards are the units. In the presence of the payoff function the empty tcg \emptyset splits into two units, reflecting the units of multiplicative and additive conjunctions in linear logic: the **top** \top has $\kappa_{\top}(\emptyset) = 1$, while the **one 1** has $\kappa_1(\emptyset) = 0$. To interpret the base type we shall use a strict board, also written o, with only one move \mathbf{q} , which is negative. Its payoff function is given by $\kappa_o(\emptyset) = 1$ and $\kappa_o(\{\mathbf{q}\}) = 0$.

\otimes	-1	0	+1	23	-1	0	+1
-1	-1	-1	-1	-1	-1	-1	+1
0	-1	0	+1	0	-1	0	+1
+1	-1	+1	+1	+1	+1	+1	+1

Figure 5: Payoff for \otimes and \Re

3.4.1. Dual, tensor and par. First, the **dual** extends with payoff via $\kappa_{A^{\perp}}(x) = -\kappa_A(x)$ as expected. Of course, the dual does not preserve –-boards. Parallel composition splits into:

Definition 3.16. Consider A and B two boards.

Their **tensor** $A \otimes B$ and their **par** $A \Im B$ are $A \parallel B$ enriched with:

$$\kappa_{A\otimes B}(x_A \parallel x_B) = \kappa_A(x_A) \otimes \kappa_B(x_B), \qquad \kappa_{A\mathfrak{B}}(x_A \parallel x_B) = \kappa_A(x_A) \,\mathfrak{F} \, \kappa_B(x_B)$$

with the operations \otimes and \Re defined on $\{-1, 0, +1\}$ in Figure 5.

The tensor of two –-boards is still a –-board, though tensor does not preserve *strict* –-boards. The par also preserves –-boards, but we shall not use it on –-boards: if A and B are –-boards, let us use $A \vdash B$ to denote the board $A^{\perp} \mathfrak{N} B$ used to define the strategies from A to B. Observe that even if A and B are –-boards, $A \vdash B$ is not.

By definition of payoff, the order-isomorphisms of (3.1) and (3.2) refines to bijections:

$$-\otimes - : \mathscr{C}^{0}(A) \times \mathscr{C}^{0}(B) \cong \mathscr{C}^{0}(A \otimes B), \qquad (3.3)$$

$$-\mathfrak{V} - : \mathscr{C}^{0}(A) \times \mathscr{C}^{0}(B) \cong \mathscr{C}^{0}(A \mathfrak{V} B).$$

$$(3.4)$$

3.4.2. *The with.* We only consider the additive conjunction of linear logic: the with. But in order to define it, we must first define a new operation on ess and tcgs.

Definition 3.17. Let A_1 and A_2 be two tcgs.

Then, we define their sum $A_1 + A_2$ as comprising the components:

events:	$ A_1 \parallel A_2 $	=	$ A_1 + A_2 $
causality:	$(i,a) \leq_{A_1 \parallel A_2} (j,a')$	\Leftrightarrow	$i = j \& a \leq_{A_i} a'$
conflict:	$(i,a) #_{A_1 \parallel A_2} (j,a')$	\Leftrightarrow	$i \neq j \lor a \#_{E_i} a',$
symmetry:	$\theta \in \mathscr{S}(A_1 \parallel A_2)$	\Leftrightarrow	$\exists \theta_i \in \mathscr{S}(A_i), \theta = \theta_1 \parallel \theta_2,$
positive symmetries:	$\theta_1 \parallel \theta_2 \in \mathscr{S}_+(A_1 + A_2)$	\Leftrightarrow	$\theta_1 \in \mathscr{S}_+(A_1) \& \theta_2 \in \mathscr{S}_+(A_2)$
negative symmetries:	$\theta_1 \parallel \theta_2 \in \mathscr{S}(A_1 + A_2)$	\Leftrightarrow	$\theta_1 \in \mathscr{S}(A_1) \& \theta_2 \in \mathscr{S}(A_2)$

where, necessarily, one of θ_1 or θ_2 must be empty.

Ignoring the positive and negative symmetries, this also yields an operation + on plain event structures with symmetry that we shall use later on.

If A, B are tcgs and $x_A \in \mathscr{C}(A)$, we write $(1, x_A) \in \mathscr{C}(A+B)$ as a shorthand for $\{1\} \times x_A$ and likewise for $(2, x_B) = \{2\} \times x_B \in \mathscr{C}(A+B)$ for $x_B \in \mathscr{C}(B)$. Note that all configurations of A + B have the form $(1, x_A)$ for $x_A \in \mathscr{C}(A)$ or $(2, x_B)$ for $x_B \in \mathscr{C}(B)$. For non-empty configurations, this decomposition is *unique*. We shall also use the corresponding notations for symmetries, with *e.g.* $(1, \theta_A) : (1, x_A) \cong_{A+B} (1, y_A)$ for $\theta_A : x_A \cong_A y_A$ comprising all ((1, a), (1, a')) for $(a, a') \in \theta_A$. This sum operation yields the *with* operation on strict boards: **Definition 3.18.** Consider S and T two strict --boards.

Then, their with S & T is the strict –-board with tcg the sum S + T and

$$\kappa_{S\&T}(1, x_S) = \kappa_S(x_S), \qquad \kappa_{S\&T}(2, x_T) = \kappa_T(x_T),$$

for non-empty configurations and $\kappa_{S\&T}(\emptyset) = 1$.

As we will see, this construction will give a cartesian product in the subcategory of strict –-boards. It can also be applied to non-strict –-boards, but then it is not a product: if in $A_1 \& A_2$, one of the A_i is not strict, then the corresponding projection does not respect payoff (in the sense of Definition 3.26), because we have set $\kappa_{A_1\&A_2}(\emptyset) = 1$. On the other hand, setting $\kappa_{A_1\&A_2}(\emptyset) = 0$ breaks the correspondence with the relational model, since the empty configuration does not correspond in a canonical way to one of the components.

In the sequel, we shall use the obvious *n*-ary generalization of the product. Observe that any strict –-board *S* decomposes uniquely (up to forest isomorphism) as $S \cong \&_{i \in I} S_i$, where each S_i is well-opened. We need notations for configurations of this board. Writing $\mathscr{C}^{\neq \emptyset}(E)$ (resp. $\mathscr{S}^{\neq \emptyset}(E)$) for the non-empty configurations (resp. symmetries) of *E*, we observe:

Lemma 3.19. Consider $(S_i)_{i \in I}$ a family of well-opened --boards. Then there are

order-isos commuting with dom and cod.

3.4.3. *Linear closure*. First, we define it in the case the rhs board is well-opened:

Definition 3.20. Consider A a --board and S a well-opened --board.

Then, $A \multimap S$ has tcg all components set as $A \vdash S$ except for:

causality:
$$\leq_{A \multimap S} = \leq_{A \vdash S} \uplus \{((2, s_0), (1, a)) \mid a \in A\}$$

writing $\min(S) = \{s_0\}$, yielding a well-opened tcg. Its payoff function is:

 $\kappa_{A\multimap S}(x_A \parallel x_S) = \kappa_{A\vdash S}(x_A \parallel x_S) = \kappa_{A^{\perp}}(x_A) \ \mathfrak{V} \ \kappa_S(x_S) \,.$

This corrects the non-negativity of $A \vdash S$, by forcing the missing dependency. In the sequel, we shall need $A \multimap S$ not only when S is well-opened (which has no particular status in the definition of relative Seely categories), but when it is strict. In that case, $A \multimap S$ may be defined directly via the decomposition into strict boards, as done in:

Definition 3.21. Consider A a –-board, and S a strict –-board, with $S \cong \&_{i \in I} S_i$. Then, we define $A \multimap S = \&_{i \in I} (A \multimap S_i)$.

The following lemma then follows from Lemma 3.19:

Lemma 3.22. Consider A, S –-boards with S strict. Then, there are:

10

order-isos commuting with dom and cod.

This anticipates on the link with the relational model, where the linear arrow is obtained with a cartesian product. Following this, we adopt the convention that for each $x_A \in \mathscr{C}(A)$ and $x_S \in \mathscr{C}^{\neq \emptyset}(S), x_A \multimap x_S \in \mathscr{C}^{\neq \emptyset}(A \multimap S)$ denotes the corresponding configuration. 3.4.4. *Exponential*. We start by defining the bang as a construction on mere ess:

Definition 3.23. Consider E an ess. Then, we define the **bang** !E with:

In $(i, e) \in !E$, we refer to *i* as a **copy index**. This is in this definition of games that symmetries really come into play: they are used to express that these copy indices can be reindexed at will. It will be convenient to characterize the shape of configurations on !E:

Lemma 3.24. For E an ess, there is an order-isomorphism

$$\lfloor - \rfloor$$
 : Fam $\left(\mathscr{C}^{\neq \emptyset}(E) \right) \simeq \mathscr{C}(!E)$ (3.5)

with Fam(X) the set of families of elements of X indexed by finite subsets of \mathbb{N} .

Proof. This associates to $(x_i)_{i \in I}$ the set $||_{i \in I} x_i = \bigcup_{i \in I} \{i\} \times x_i$.

To clarify the notation above: we mean that for any family $(x_i)_{i \in I} \in \mathsf{Fam}(\mathscr{C}^{\neq \emptyset}(E))$, we write $\lfloor x_i \mid i \in I \rfloor \in \mathscr{C}(!E)$ for the corresponding configuration of !E. In addition, it is clear that if X is a set, then $\mathsf{Fam}(X)$ quotiented by permutations of indices is in bijection with $\mathscr{M}_f(X)$. Hence, the order-isomorphism of Lemma 3.24 immediately yields a bijection

$$\mathscr{C}(!E)/\cong_{!E} \simeq \mathscr{M}_f(\mathscr{C}^{\neq\emptyset}(E)/\cong_E)$$
(3.6)

which again suggests the forthcoming connection with the relational model.

Now, we extend the bang construction to boards:

Definition 3.25. Consider S a strict board. Then, we define the **bang** !S additionally has:

with payoff given by $\kappa_{!S}(\lfloor x_i \mid i \in I \rfloor) = \bigotimes_{i \in I} \kappa_S(x_i)$, and $\kappa_{!S}(\lfloor \rfloor) = 0$.

If S is a strict board, then !S is still a --board, but no longer strict. Since the minimal events of S are negative, an exchange in the copy indices arising from this definition is viewed as *negative*. Hence, positive symmetries can not affect them. Also, because S is strict, its complete configurations are non-empty; hence (3.5) and (3.6) refine to:

$$\lfloor - \rfloor : \quad \mathsf{Fam}(\mathscr{C}^0(S)) \cong \mathscr{C}^0(!S) \tag{3.7}$$

$$[-]: \mathscr{M}_f(\mathscr{C}^0(S)/\cong_S) \simeq \mathscr{C}^0(!S)/\cong_{!S}$$
(3.8)

which again suggests the forthcoming relational collapse. Note that this only holds if S is a strict board. The bang !S does work in more generality [CC21, Cla24], but not in a way that is compatible with the relational model. Accordingly, we shall focus not on the Seely category structure where ! is a comonad, but in a variation called *relative Seely category* where ! is a *relative comonad*; we will come back to that point later on. 3.5. The Relative Seely Category of Sequential Innocence. As explained before, we will form a category whose objects are –-boards. The morphisms will be certain strategies in the sense of Definition 3.10. But to ensure the existence of a functorial collapse to the relational model, we must impose additional conditions on strategies.

3.5.1. Deterministic sequential innocence. We first introduce our notion of strategies:

Definition 3.26. Consider A a board, and $\sigma : A$ a strategy. We define conditions:

 $\begin{array}{ll} negative: & \text{for all } s \in \sigma \text{ minimal for } \leq_{\sigma}, \text{ we have } \mathrm{pol}_{A}(\partial_{\sigma} \, s) = -, \\ winning: & \text{for all } x^{\sigma} \in \mathscr{C}^{+}(\sigma), \, \kappa_{A}(\partial_{\sigma} \, x^{\sigma}) \geq 0, \\ forestial: & \leq_{\sigma} \text{ is a forest.} \\ deterministic: & \text{if } s^{-} \twoheadrightarrow_{\sigma} s_{1}^{+}, s_{2}^{+} \text{ then } s_{1}^{+} = s_{2}^{+}. \end{array}$

We say σ is deterministic sequential innocent (dsinn) if it satisfies all four.

Winning ensures that strategies are well-behaved with respect to payoff: in particular, a closed interaction between winning strategies will always result in a complete position, which is essential for the relational collapse. Forestial and deterministic make σ a negatively branching forest, which mimics a syntactic tree⁷; this ensures that composition is deadlockfree and therefore matches relational composition.

Copycat strategies on –-boards are automatically deterministic sequential innocent. Deterministic sequential innocence is also stable under composition, which ensures [Cla24]:

Theorem 3.27. There is a \sim -category **DSInn** with: (1) objects, --boards; (2) morphisms, dsinn strategies; (3) equivalence, positive isomorphism.

3.5.2. Relative Seely categories. This category (we shall often omit the "~-" from now on) has significant further structure. In particular, it can be organized into a Seely category, a categorical model of intuitionistic linear logic [Mel09]; but unfortunately that structure is not preserved by the relational collapse. For instance, for a game !!A, there is only one empty configuration, whereas in $\mathcal{M}_f(\mathcal{M}_f(\mathcal{C}(A)))$ there are multiple ways to be "empty": [], [[], []], ...; even if the empty configuration on A is not deemed a valid position. In a sense, the relational model counts how many times a program "does nothing", which is meaningless if states in the relational model are to correspond to sets of events.

Fortunately, this mismatch arises outside of the translation of simple types (or even standard linear/non-linear systems, where ! can only occur on the left hand side of an arrow). The categorical structure describing the structure that *is* preserved is called a *relative Seely category* [CP21, Cla24]; we now recall the definition.

Definition 3.28. A relative Seely category is a symmetric monoidal category $(\mathcal{C}, \otimes, 1)$ equipped with a full subcategory \mathcal{C}_s together with the following data and axioms:

- \mathcal{C}_s has finite products $(\&, \top)$ preserved by the inclusion functor $J : \mathcal{C}_s \hookrightarrow \mathcal{C}$.
- For every $B \in \mathcal{C}$ there is a functor $B \multimap -: \mathcal{C}_s \to \mathcal{C}_s$, such that there is

$$\Lambda(-): \mathcal{C}(A \otimes B, S) \simeq \mathcal{C}(A, B \multimap S).$$

a bijection natural in $A \in \mathcal{C}$ and $S \in \mathcal{C}_s$.

⁷The relational collapse is possible for the more general notion of *visible* strategies, see [Cla24]. But the rest of this paper will depend on this specific notion of deterministic sequential strategy.

• There is a *J*-relative comonad $!: \mathcal{C}_s \to \mathcal{C}$. This means that we have, for every $S \in \mathcal{C}_s$, an object $!S \in \mathcal{C}$ and a **dereliction** morphism der_S $: !S \to S$, and for every $\sigma : !S \to T$, a **promotion** $\sigma^! : !S \to !T$, subject to three axioms [ACU10]:

$\operatorname{der}_T \circ \sigma^!$	=	σ	$(\sigma : !S \to T)$
$\operatorname{der}_{S}^{!}$	=	$\mathrm{id}_{!S}$	$(S \in \mathcal{C}_s)$
$(\tau \circ \sigma^!)^!$	=	$ au^! \circ \sigma^!$	$(\sigma: !S \to T, \tau: T \to U) ,$

which make $!: \mathcal{C}_s \to \mathcal{C}$ a functor, via $!\sigma = (\sigma \circ \operatorname{der}_S)^!$ for $\sigma: S \to T$.

• The functor $!: \mathcal{C}_s \to \mathcal{C}$ is symmetric strong monoidal $(\mathcal{C}_s, \&, \top) \to (\mathcal{C}, \otimes, 1)$, so

$$m_0: 1 \to !\top$$
 $m_{S,T}: !S \otimes !T \to !(S \& T)$

are natural isos, additional compatible with promotion: the diagram

commutes for all $\Gamma, S, T \in \mathcal{C}_s, f \in \mathcal{C}(!\Gamma, S), g \in \mathcal{C}(!\Gamma, T).$

Any Seely category is a relative Seely category with $C = C_s$. For any relative Seely category, the Kleisli category associated with !, denoted $C_!$, is cartesian closed: it has objects those of C_s , morphisms $C_!(S,T) = C(!S,T)$, products &, and function space $S \Rightarrow T = !S \multimap T$.

There is an accompanying notion of morphism: a relative Seely functor from C to D is a functor $F : C \to D$ together with isomorphisms

$t_{A,B}^{\otimes}$:	$FA \otimes FB$	\cong	$F(A\otimes B)$	t^1	:	1	\cong	F1
$t_{S,T}^{\&}$:	FS & FT	\cong	F(S & T)	$t^{ op}$:	Т	\cong	$F \top$
$t_{A,S}^{-\circ}$:	$FA \multimap FS$	\cong	$F(A \multimap S)$	$t_S^!$:	!FS	\cong	F!S

satisfying appropriate naturality and coherence conditions [Cla24]. This ensures in particular that F lifts to a cartesian closed functor $F_1 : C_1 \to D_1$.⁸

3.5.3. The relative Seely category **DSInn**. We now spell out the additional structure in **DSInn**. From now on, by *strategy* we mean a morphism in **DSInn**, *i.e.* a sequential deterministic strategy in the sense of Definition 3.26. First, the symmetric monoidal structure is handled by [Cla24, Propositions 8.1.1 and 8.2.16]:

Proposition 3.29. Consider A, B, C, D –-boards, and $\sigma : A \vdash B, \tau : C \vdash D$ strategies. Then, there is a strategy $\sigma \otimes \tau : A \otimes C \vdash B \otimes D$, unique up to iso, s.t. there are

$$\begin{array}{rcl} (-\otimes -) & : & \mathscr{C}^+(\sigma) \times \mathscr{C}^+(\tau) & \simeq & \mathscr{C}^+(\sigma \otimes \tau) \\ (-\otimes -) & : & \mathscr{S}^+(\sigma) \times \mathscr{S}^+(\tau) & \simeq & \mathscr{S}^+(\sigma \otimes \tau) \end{array}$$

order-isos commuting with dom, cod, and s.t. for all $\theta^{\sigma} \in \mathscr{S}^+(\sigma)$ and $\theta^{\tau} \in \mathscr{S}^+(\tau)$,

$$\partial_{\sigma\otimes\tau}(\theta^{\sigma}\otimes\theta^{\tau})=(\theta^{\sigma}_{A}\otimes\theta^{\tau}_{C})\otimes(\theta^{\sigma}_{B}\otimes\theta^{\tau}_{D}).$$

⁸These definitions must of course be taken in \sim -categorical form: as before this means that all operations preserve \sim , and all conditions hold up to \sim ; yielding notions of **relative Seely** \sim -category and **relative Seely** \sim -functor. We omit the straightforward adaptation.

Moreover, $(-\otimes -)$ preserves positive isomorphism.

From this definition, Propositions 3.5, 3.11 and a straightforward symmetry-aware extension of Lemma 3.3, it follows that \otimes is a bifunctor. In addition, there are associativity and unit natural isomorphisms defined as the obvious copycat strategies, satisfying the required conditions for a symmetric monoidal category [Cla24, Proposition 8.2.20].

Next, we move to the product, which is defined via [Cla24, Proposition 8.2.22]:

Proposition 3.30. For –-boards Γ , A, B, with A, B strict, and strategies $\sigma : \Gamma \vdash A, \tau : \Gamma \vdash$ B, there is a strategy $\langle \sigma, \tau \rangle : \Gamma \vdash A \& B$, unique up to iso, such that there are order-isos:

$$\begin{aligned} \mathscr{C}^+(\sigma) + \mathscr{C}^+(\tau) &\simeq & \mathscr{C}^+(\langle \sigma, \tau \rangle) \\ \mathscr{S}^+(\sigma) + \mathscr{S}^+(\tau) &\simeq & \mathscr{S}^+(\langle \sigma, \tau \rangle) \end{aligned}$$

commuting with dom, cod, and such that for all $\theta^{\sigma} \in \mathscr{S}^+(\sigma)$ and $\theta^{\tau} \in \mathscr{S}^+(\tau)$, we have

$$\partial_{\langle \sigma, \tau \rangle}(\mathsf{i}_{\sigma}(\theta^{\sigma})) = \theta^{\sigma}_{\Gamma} \vdash \mathsf{i}_{A}(\theta^{\sigma}_{A}), \qquad \partial_{\langle \sigma, \tau \rangle}(\mathsf{i}_{\tau}(\theta^{\tau})) = \theta^{\tau}_{\Gamma} \vdash \mathsf{i}_{B}(\theta^{\tau}_{B})$$

with $i_{\sigma}: \mathscr{C}^+(\sigma) \to \mathscr{C}^+(\langle \sigma, \tau \rangle)$ and $i_{\tau}: \mathscr{C}^+(\tau) \to \mathscr{C}^+(\langle \sigma, \tau \rangle)$ the induced injections. Moreover, $\langle -, - \rangle$ preserves positive isomorphism.

There are also projections, obtained again as the obvious copycat strategies

 $\pi_A: A \& B \vdash A$, $\pi_B : A \& B \vdash B$

turning A & B into a categorical product [Cla24, Proposition 8.2.24] for A and B strict, and additionally \top is terminal. Finally, for the arrow type, we have [Cla24, Proposition 8.2.25]:

Proposition 3.31. Consider Γ , A, B –-boards with B strict. For $\sigma: \Gamma \otimes A \vdash B$, there is $\Lambda(\sigma): \Gamma \vdash A \multimap B$, unique up to iso, s.t. there are

$$\begin{array}{rcl} \Lambda(-) & : & \mathscr{C}^+(\sigma) & \cong & \mathscr{C}^+(\Lambda(\sigma)) \\ \Lambda(-) & : & \mathscr{S}^+(\sigma) & \cong & \mathscr{S}^+(\Lambda(\sigma)) \end{array}$$

order-isos commuting with dom, cod, and such that for all θ^{σ} non-empty,

$$\partial_{\Lambda(\sigma)}(\Lambda(\theta^{\sigma})) = \theta^{\sigma}_{\Gamma} \vdash \theta^{\sigma}_{A} \multimap \theta^{\sigma}_{B}$$

$$(3.9)$$

whenever $\partial_{\sigma}(\theta^{\sigma}) = \theta^{\sigma}_{\Gamma} \parallel \theta^{\sigma}_{A} \vdash \theta^{\sigma}_{B}$. Moreover, $\Lambda(-)$ preserves positive isomorphism.

This construction, **currying**, is easily shown to be invertible up to isomorphism. As usual, the **evaluation** is defined as $ev_{A,B} = \Lambda^{-1}(\mathbf{c}_{A \to B})$, the uncurrying of the identity, for A, B two –-boards with B strict; altogether this makes $A \multimap B$ an arrow of A and B.

The last outstanding construction is the *promotion*, relative to the exponential:

Proposition 3.32. Consider S, T strict –-boards, and $\sigma : !S \vdash T$ a strategy. Then, the ess σ may be equipped with a display map $\partial_{\sigma^{\dagger}}$ such that

$$\partial_{\sigma^{\dagger}}(\lfloor x^{\sigma,i} \mid i \in I \rfloor) = \lfloor x^{\sigma,i}_{A,j} \mid \langle i,j \rangle \in \Sigma_{i \in I} J_i \rfloor \vdash \lfloor x^{\sigma,i}_B \mid i \in I \rfloor$$

where $\partial_{\sigma}(x^{\sigma,i}) = \lfloor x_{A,j}^{\sigma,i} \mid j \in J_i \rfloor \vdash x_B^{\sigma,i}$, writing $\sum_{i \in I} J_i$ for the set of encodings $\langle i, j \rangle \in \mathbb{N}$ of all pairs of $i \in I$ and $j \in J_i$, where $\langle -, - \rangle : \mathbb{N}^2 \simeq \mathbb{N}$ is an arbitrary bijection.

This makes $\sigma^{\dagger} : !S \vdash !T$ a strategy, and the construction $(-)^{\dagger}$ preserves positive iso.

The **dereliction** strategy der_A : $!A \vdash A$ is defined as a copycat strategy, opening one copy with copy index 0. Finally, the *Seely isomorphisms* (up to positive iso) are

$$\operatorname{see}_{A,B}$$
 : $!A \otimes !B \cong !(A \& B)$,

defined again by the obvious copycat strategy, and the obvious isomorphism $!\top \cong 1$ between empty games. Altogether, these provide all the components for the desired structure:

Theorem 3.33. The components above make **DSInn** a relative Seely category; where the strict full subcategory **DSInn**_s is restricted to strict --boards.

In particular, the Kleisli category **DSInn**_l is cartesian closed.

4. Relational Collapse

Now that we have set up our ambiant game semantics, we are in position to resume the discussion in the introduction now resting on precise definitions. The aim of this section is to recall the relational collapse of thin concurrent games, for deterministic sequential strategies.

This takes the form of a relative Seely functor

$$\mathfrak{R}(-):\mathbf{DSInn}
ightarrow\mathbf{Rel}$$

which we describe in this section. Again, we borrow much of the presentation from [Cla24].

4.1. General Idea. We start by giving the basic definition of the collapse.

4.1.1. *Collapsing games.* As argued before, *boards* come equipped with a notion of **stopping configurations**: namely, those configurations whose payoff is null:

$$(A) = \{ x \in \mathscr{C}(A) \mid \kappa_A(x) = 0 \},\$$

which as argued in the introduction, are designed to match notions of rigid intersection types; or alternatively, the objects of the interpretation of a type in generalized species [COP23].

In turn, points of the web in the sense of the relational model – or non-idempotent intersection types – will correspond to stopping configurations up to symmetry

$$\mathfrak{R}(A) = (A)/\cong_A \tag{4.1}$$

called the **positions** of A. Here, we use symbols x, y, z... to range over symmetry classes of configurations – note the different font than for configurations.

Now, the idea is simply to send a board A to its set of positions. This is well-behaved, in the sense that (by design!) there are relatively straightforward bijections presented in Figure 6 where A, B are any --boards and S, T are strict. For **1** and \top this is clear. For the tensor and hom-game this follows from (3.4) and (3.4). For the with this comes from Lemma 3.19, for the linear arrow from Lemma 3.22, and for the bang, from (3.8).

From the above, as an immediate corollary we get a bijection, for every simply type A:

$$s_A : \llbracket A \rrbracket_{\mathbf{Rel}} \simeq \Re(\llbracket A \rrbracket_{\mathbf{DSInn}})$$
 (4.2)

obtained simply by induction on A – and similarly for a simply-typed context Γ . So as announced, we are able to identify the points in the relational interpretation of a type A with certain (symmetry classes of) configurations of the game semantics interpretation of A.

$$\begin{array}{rclcrcl} \mathbf{s}_{A,B}^{\vdash} & : & \mathfrak{R}(A) \times \mathfrak{R}(B) & \simeq & \mathfrak{R}(A \vdash B) \\ \mathbf{s}_{A,B}^{\otimes} & : & \mathfrak{R}(A) \times \mathfrak{R}(B) & \simeq & \mathfrak{R}(A \otimes B) \\ \mathbf{s}^{\mathbf{1}} & : & \mathbf{1} & \simeq & \mathfrak{R}(\mathbf{1}) \\ \mathbf{s}_{S}^{!} & : & \mathcal{M}_{f}(\mathfrak{R}(S)) & \simeq & \mathfrak{R}(!S) \\ \mathbf{s}^{\top} & : & \emptyset & \simeq & \mathfrak{R}(\top) \\ \mathbf{s}_{S,T}^{\&} & : & \mathfrak{R}(S) + \mathfrak{R}(T) & \simeq & \mathfrak{R}(S \& T) \\ \mathbf{s}_{A,S}^{\&} & : & \mathfrak{R}(A) \times \mathfrak{R}(S) & \simeq & \mathfrak{R}(A \multimap S) \end{array}$$

Figure 0: Structural collapse blie	lections
------------------------------------	----------

$$\begin{split} s^{\mathsf{L}}_{A,B} &: \ \mathfrak{R}^{\mathcal{C}}(A) \times \mathfrak{R}^{\mathcal{C}}(B) &\simeq \ \mathfrak{R}^{\mathcal{C}}(A \vdash B) \\ s^{\otimes}_{A,B} &: \ \mathfrak{R}^{\mathcal{C}}(A) \times \mathfrak{R}^{\mathcal{C}}(B) &\simeq \ \mathfrak{R}^{\mathcal{C}}(A \otimes B) \\ s^{\mathbf{1}} &: \qquad \mathbf{1} &\simeq \ \mathfrak{R}^{\mathcal{C}}(1) \\ s^{\mathsf{I}}_{S} &: \qquad \mathscr{M}_{f}(\mathfrak{R}^{\mathcal{C}}(S)) &\simeq \ \mathfrak{R}^{\mathcal{C}}(!S) \\ s^{\mathsf{T}} &: \qquad \emptyset &\simeq \ \mathfrak{R}^{\mathcal{C}}(\top) \\ s^{\&}_{S,T} &: \ \mathfrak{R}^{\mathcal{C}}(S) + \mathfrak{R}^{\mathcal{C}}(T) &\simeq \ \mathfrak{R}^{\mathcal{C}}(S \& T) \\ s^{\to}_{A,S} &: \ \mathfrak{R}^{\mathcal{C}}(A) \times \mathfrak{R}^{\mathcal{C}}(S) &\simeq \ \mathfrak{R}^{\mathcal{C}}(A \multimap S) \end{split}$$

Figure 7: Colored collapse bijections



Figure 8: The relational collapse : forgetting the dynamic order

4.1.2. Collapsing strategies. Now, we must extend this to strategies. The idea is rather simple: we intend to simply send a strategy to its set of reached positions. In fact, the slightly better behaved definition consists in sending a strategy $\sigma : A \vdash B$ to

$$\mathfrak{R}(\sigma) = \{ (\mathsf{x}_A, \mathsf{x}_B) \in \mathfrak{R}(A) \times \mathfrak{R}(B) \mid \exists x^{\sigma} \in \mathscr{C}^+(\sigma), \ x^{\sigma}_A \in \mathsf{x}_A, \ x^{\sigma}_B \in \mathsf{x}_B \}$$
(4.3)

those positions reached by +-covered configurations only⁹. This is illustrated in Figure 8, with one +-covered configuration arising from the interpretation of λfx . $f(fx) : (o_1 \rightarrow o_2) \rightarrow o_3 \rightarrow o_4$, labelling the occurrences of the base type to match the moves in the diagram. There are two phenomena at play here: firstly, the dynamic causal links \rightarrow are forgotten, leaving only the underlying configuration. Secondly, the copy indices are forgotten by taking the symmetry class. Note that by doing only the first, one may collapse not merely to the relational model, but to its categorification, generalized species of structure [COP23].

4.2. **Composition.** Hopefully, the above conveys the idea of the relational collapse. There are however a few conceptual subtleties that have to do with preservation of composition.

For this section, fix two strategies $\sigma : A \vdash B$ and $\tau : B \vdash C$ in **DSInn**.

4.2.1. Oplax preservation. We start with the easy inclusion, namely:

$$\mathfrak{R}(\tau \odot \sigma) \subseteq \mathfrak{R}(\tau) \circ \mathfrak{R}(\sigma)$$
.

Consider $(\mathsf{x}_A, \mathsf{x}_C) \in \mathfrak{R}(\tau \odot \sigma)$. From the definition, this means that there is some $x^{\tau \odot \sigma} \in \mathscr{C}^+(\tau \odot \sigma)$ such that $x_A^{\tau \odot \sigma} \in \mathsf{x}_A$ and $x_C^{\tau \odot \sigma} \in \mathsf{x}_C$. By Proposition 3.5, +-covered configurations of $\tau \odot \sigma$ are in one-to-one correspondence with pairs (x^{σ}, x^{τ}) of causally

⁹In fact, in the case of deterministic sequential innocent strategies, the configurations that reach positions are always +-covered, under the mild assumption that stopping configurations have as many Opponent as Player events. We stick with the +-covered definition, which works without sequential innocence [Cla24].

compatible $x^{\sigma} \in \mathscr{C}^+(\sigma)$ and $x^{\tau} \in \mathscr{C}^+(\tau)$. Here, causally compatible means that $x_B^{\sigma} = x_B^{\tau}$, and that the induced synchronization of x^{σ} and x^{τ} causes no deadlock – recall that we then write $x^{\tau \odot \sigma} = x^{\tau} \odot x^{\sigma} \in \mathscr{C}^+(\tau \odot \sigma)$.

To show $(\mathsf{x}_A, \mathsf{x}_C) \in \mathfrak{R}(\tau) \circ \mathfrak{R}(\sigma)$, we must exhibit $\mathsf{x}_B \in \mathfrak{R}(B)$ such that $(\mathsf{x}_A, \mathsf{x}_B) \in \mathfrak{R}(\sigma)$ and $(\mathsf{x}_B, \mathsf{x}_C) \in \mathfrak{R}(C)$. This seems simple: we have $x_B^{\sigma} = x_B^{\tau}$, so we may simply take x_B their symmetry class. We must however ensure that $\kappa_A(\mathsf{x}_B) = 0$ – and this follows easily from $\kappa_A(\mathsf{x}_A) = \kappa_C(\mathsf{x}_C) = 0$ using the fact that σ and τ are winning [Cla24, Lemma 10.4.6].

4.2.2. Lax preservation. We now focus on the other inclusion:

$$\mathfrak{R}(\tau) \circ \mathfrak{R}(\sigma) \subseteq \mathfrak{R}(\tau \odot \sigma).$$

$$(4.4)$$

In this direction, the situation is more subtle. Consider $(\mathsf{x}_A, \mathsf{x}_B) \in \mathfrak{R}(\sigma)$ and $(\mathsf{x}_B, \mathsf{x}_C) \in \mathfrak{R}(\tau)$. By definition, those are witnessed by $x^{\sigma} \in \mathscr{C}^+(\sigma)$ and $x^{\tau} \in \mathscr{C}^+(\tau)$ such that

$$x_A^{\sigma} \in \mathsf{x}_A, \qquad x_B^{\sigma} \in \mathsf{x}_B, \qquad x_B^{\tau} \in \mathsf{x}_B, \qquad x_C^{\tau} \in \mathsf{x}_C$$

and ideally, we would like to form $x^{\tau} \odot x^{\sigma} \in \mathscr{C}^+(\tau \odot \sigma)$ using Proposition 3.5. But there are two issues: firstly, we may not have $x_B^{\sigma} = x_B^{\tau}$, in general all we have is $x_B^{\sigma}, x_B^{\tau} \in \mathsf{x}_B$ so that there must exists some unspecified symmetry $\theta : x_B^{\sigma} \cong_B x_B^{\tau}$. Secondly, even if we had $x_B^{\sigma} = x_B^{\tau}$, it is not clear why the induced synchronization would be deadlock-free.

Fortunately, it is a general fact that sequential innocent strategies cannot deadlock:

Lemma 4.1. Consider A, B, C --boards, $\sigma : A \vdash B$ and $\tau : B \vdash C$ deterministic sequential innocent strategies, $x^{\sigma} \in \mathscr{C}(\sigma)$ and $x^{\tau} \in \mathscr{C}(\tau)$ with a symmetry $\theta : x_B^{\sigma} \cong_B x_B^{\tau}$.

Then, the composite bijection

$$\varphi : x^{\sigma} \parallel x_{C}^{\tau} \stackrel{\partial_{\sigma} \parallel x_{C}^{\tau}}{\simeq} x_{A}^{\sigma} \parallel x_{B}^{\sigma} \parallel x_{C}^{\tau} \stackrel{x_{A}^{\sigma} \parallel \theta \parallel x_{C}^{\tau}}{\simeq} x_{A}^{\sigma} \parallel x_{B}^{\tau} \parallel x_{C}^{\tau} \stackrel{x_{A}^{\sigma} \parallel \partial_{\tau}^{-1}}{\simeq} x_{A}^{\sigma} \parallel x_{T}^{\tau},$$

is secured, in the sense of Proposition 3.13.

This deadlock-free lemma was first proved in the context of [CCW15] – the reader is rather directed to the more recent and detailed presentation in [Cla24, Lemma 10.4.8], where it is proved with the more general hypothesis that σ and τ should be visible. This crucial lemma bridges the main conceptual difference between game semantics and relational semantics: the former is sensitive to deadlocks, whereas the latter is not.

Applying Lemma 4.1 to the data at hand, the obtained securedness hypothesis lets us reindex x^{σ} and x^{τ} by Proposition 3.13 to obtain $y^{\tau} \odot y^{\sigma} \in \mathscr{C}^+(\tau \odot \sigma)$ such that $y^{\sigma}_A \cong x^{\sigma}_A$ and $y^{\tau}_C \cong_C x^{\tau}_C$, so that $y^{\tau} \odot y^{\sigma}$ witnesses $(\mathsf{x}_A, \mathsf{x}_C) \in \mathfrak{R}(\tau \odot \sigma)$ as required.

4.3. Further structure. From the above, $\Re(\tau \odot \sigma) = \Re(\tau) \circ \Re(\sigma)$. From Proposition 3.4, it is immediate that $\Re(\mathbf{c}_A)$ is the identity relation on $\Re(A)$. All other constructions on strategies are preserved in the appropriate sense of a relative Seely functor, with respect to the isomorphisms in **Rel** induced with the bijections of Figure 6; as follows via routine verifications from Propositions 3.29, 3.30, 3.31 and 3.32. Altogether:

Theorem 4.2. The above provide the components for a relative Seely functor:

$$\mathfrak{R}(-): \mathbf{DSInn} \to \mathbf{Rel}$$
.

See [Cla24, Corollary 10.4.15] for more details. It follows in particular that we also get

 $\mathfrak{R}_!(-):\mathbf{DSInn}_! \to \mathbf{Rel}_!$

a cartesian closed functor between the induced cartesian closed categories, so that:

Corollary 4.3. Consider $\Gamma \vdash M : A$ a simply-typed λ -term.

Then, the following diagram commutes in Rel:



4.4. Relational Collapse in Colors. The above theorem holds with respect to the relational interpretation of simple types fixed in Section 2.1.1, which specified in particular that $\llbracket o \rrbracket = \{\star\}$ some singleton set. This comes into play in the existence of the unique bijection between singleton sets $s_o : \llbracket o \rrbracket_{\mathbf{Rel}} \simeq \Re(\llbracket o \rrbracket_{\mathbf{DSInn}})$. Here, we show how Corollary 6.15 must be adapted if we instead have $\llbracket o \rrbracket_{\mathbf{Rel}} = \mathcal{C}$ some arbitrary set of *colors*. For disambiguation, from now on we shall explicitly specify the interpretation of the base type in the relational interpretation with $\llbracket - \rrbracket_{\mathbf{Rel}}^{\mathcal{C}}$.

Though the relational interpretation of types changes, the interpretation of types as games remains the same; thus the connection between games and positions must be adjusted to a connection between games and positions *in colors*. This mechanism is new in the context of concurrent games (it does not apply in earlier published works involving the relational collapse of concurrent games), but a similar idea already appears in Tsukada and Ong's account of the relationship between game semantics and the relational model [TO16].

4.4.1. *Positions in colors.* We start by adjusting configurations and positions:

Definition 4.4. Consider A a board, and $x \in \mathscr{C}(A)$ a configuration.

A *C*-coloring (or just coloring, when *C* is clear from the context) of *x* is a function $\lambda : x \to C$. We write col(x) for the set of *C*-colorings of *x*, and $\mathscr{C}_{\mathcal{C}}(A)$ for the set of configurations in colors, *i.e.* pairs (x, λ) of a configuration equipped with a coloring.

Though a configuration with colors is a pair $(x, \lambda) \in \mathscr{C}_{\mathcal{C}}(A)$, we shall sometimes just write $x \in \mathscr{C}_{\mathcal{C}}(A)$ and refer to the coloring as $\lambda_x \in \operatorname{col}(x)$.

If A and B are boards, $x_A \in \mathscr{C}(A)$ and $x_B \in \mathscr{C}(B)$, then every pair of colorings $\lambda_A \in \operatorname{col}(x_A)$ and $\lambda_B \in \operatorname{col}(x_B)$ induce a coloring $\lambda_A \otimes \lambda_B \in \operatorname{col}(x_A \otimes x_B)$ simply by copairing, informing a bijection $\operatorname{col}(x_A \otimes x_B) \simeq \operatorname{col}(x_A) \times \operatorname{col}(x_B)$. Together with the bijection $-\otimes -: \mathscr{C}(A) \times \mathscr{C}(B) \simeq \mathscr{C}(A \otimes B)$, this yields a bijection $-\otimes -: \mathscr{C}_{\mathcal{C}}(A) \times \mathscr{C}_{\mathcal{C}}(B) \simeq \mathscr{C}_{\mathcal{C}}(A \otimes B)$. We have similar bijections for $\vdash, \&, \multimap$ and !, defined in the obvious way.

Definition 4.5. Two configurations with colors $x, y \in \mathscr{C}_{\mathcal{C}}(A)$ are symmetric if there is some $\theta : x \cong_A y$ that preserves colors. This is an equivalence relation, and a **position with** colors is a symmetry class of configurations with colors of null payoff, written $x \in \mathfrak{R}^{\mathcal{C}}(A)$. All the above bijections between configurations in colors are compatible with symmetry, and yield the bijections of Figure 7. For the base type o, a coloring consists simply in the choice of a color for the unique move q, so that we indeed have $\Re^{\mathcal{C}}(\llbracket o \rrbracket_{\mathbf{DSInn}}) \simeq \llbracket o \rrbracket_{\mathbf{Rel}}^{\mathcal{C}} = \mathcal{C}$. Altogether, the bijection (4.2) extends in the presence of colors to give

$$s_A : [A]^{\mathcal{C}}_{\mathbf{Rel}} \simeq \mathfrak{R}^{\mathcal{C}}(A)$$

for every simple type A, extending our earlier situation in the presence of colors.

4.4.2. *Experiments*. Next, we must associate to any strategy a set of positions in colors. This rests on the following notion of *experiment*, the name being inspired from the notion with the same name in proof nets [Gir87]. Intuitively, an experiment is a coloring of a configuration of a strategy; except that the axiom links must be preserved: a Player move must have the same color as that given to its (necessarily unique) causal predecessor.

Definition 4.6. Consider A a board, $\sigma : A$ a strategy, and $x \in \mathscr{C}^+(\sigma)$.

A (C-)coloring of x is a function $\lambda : x \to C$ subject to:

monochrome: for all $s_1^- \to_{\sigma} s_2^+$, we have $\lambda(s_1^-) = \lambda(s_2^+)$.

As above, we write $\operatorname{col}(x)$ for the set of colorings of $x \in \mathscr{C}^+(\sigma)$. An **experiment** on σ is $x \in \mathscr{C}^+(\sigma)$ together with a coloring, and we write $\mathscr{C}^+_{\mathcal{C}}(\sigma)$ for the set of experiments.

As for configurations in colors, experiments are pairs $(x, \lambda) \in \mathscr{C}_{\mathcal{C}}(\sigma)$; nevertheless we shall often just write $x \in \mathscr{C}_{\mathcal{C}}(\sigma)$ and refer to the coloring as $\lambda_x \in \operatorname{col}(x)$.

Given $x^{\sigma} \in \mathscr{C}^+(\sigma)$, recall that the display map $\partial_{\sigma} : \sigma \to A$ induces a bijection $\partial_{\sigma} : x^{\sigma} \simeq x^{\sigma}_A$. We use this bijection to transport any coloring $\lambda^{\sigma} \in \operatorname{col}(x^{\sigma})$ to $\lambda^{\sigma}_A = \lambda^{\sigma} \circ \partial_{\sigma}^{-1} \in \operatorname{col}(x^{\sigma}_A)$. Likewise, if $x^{\sigma} \in \mathscr{C}_{\mathcal{C}}(\sigma)$, we write $x^{\sigma}_A \in \mathscr{C}_{\mathcal{C}}(A)$ for the corresponding configuration in colors.

4.4.3. The colorful collapse. With this we extend the collapse to the colored setting:

$$\mathfrak{R}^{\mathcal{C}}(\sigma) = \{ \mathsf{x}_A \vdash \mathsf{x}_B \in \mathfrak{R}^{\mathcal{C}}(A \vdash B) \mid \exists x^{\sigma} \in \mathscr{C}^+_{\mathcal{C}}(\sigma), \ x^{\sigma}_A \in \mathsf{x}_A, \ x^{\sigma}_B \in \mathsf{x}_B \},$$
(4.5)

which we now aim to show still yields a relative Seely functor. As perhaps expected, the noteworthy cases are preservation of composition, and of copycat.

We focus first on composition; fix $\sigma : A \vdash B$ and $\tau : B \vdash C$, along with $x^{\sigma} \in \mathscr{C}^+(\sigma)$ and $x^{\tau} \in \mathscr{C}^+(\tau)$ causally compatible. Given colorings $\lambda^{\sigma} \in \operatorname{col}(x^{\sigma})$ and $\lambda^{\tau} \in \operatorname{col}(x^{\tau})$, we say that they are **matching** if $\lambda_B^{\sigma} = \lambda_B^{\tau}$. Preservation of composition rests on:

Lemma 4.7. Fix $\sigma : A \vdash B, \tau : B \vdash C$ with $x^{\sigma} \in \mathscr{C}^+(\sigma)$ and $\mathscr{C}^+(\tau)$ causally compatible. Then, there is a bijection:

$$\begin{split} &-\odot - \quad : \quad \{(\lambda^{\sigma}, \lambda^{\tau}) \in \operatorname{col}(x^{\sigma}) \times \operatorname{col}(x^{\tau}) \mid matching\} \quad \simeq \quad \operatorname{col}(x^{\tau} \odot x^{\sigma}) \,,\\ & satisfying \; (\lambda^{\tau} \odot \lambda^{\sigma})_A = \lambda^{\sigma}_A \; and \; (\lambda^{\tau} \odot \lambda^{\sigma})_C = \lambda^{\tau}_C. \end{split}$$

Proof. In this proof, we use terminology and notations from [Cla24].

Recall that $\tau \odot \sigma$ is the restriction to events occurring in A, C of the *interaction* $\tau \circledast \sigma$, an event structure with display map $\partial_{\tau \circledast \sigma} : \tau \circledast \sigma \to A \parallel B \parallel C$ representing the interaction of σ and τ without hiding. Along with this, any $x^{\tau} \odot x^{\sigma} \in \mathscr{C}(\tau \odot \sigma)$ has a witness $x^{\tau} \circledast x^{\sigma}$, obtained as the down-closure $x^{\tau} \circledast x^{\sigma} = [x^{\tau} \odot x^{\sigma}]_{\tau \circledast \sigma}$, with display $\partial_{\tau \circledast \sigma} (x^{\tau} \circledast x^{\sigma}) = x_A^{\sigma} \parallel x_B \parallel x_C^{\tau}$ for $x_B = x_B^{\sigma} = x_B^{\tau}$. Reciprocally, one recovers $x^{\tau} \odot x^{\sigma}$ from $x^{\tau} \circledast x^{\sigma}$ by restricting it to visible events, *i.e.* those events in A, C – see [Cla24, Section 6.2.2].

From left to right, fix $x^{\tau} \odot x^{\sigma} \in \mathscr{C}^+(\tau \odot \sigma)$ with matching colorings λ^{σ} and λ^{τ} – as they are matching, they induce a coloring $\lambda^{\tau} \circledast \lambda^{\sigma}$ on $x^{\tau} \circledast x^{\sigma}$, which restricts to a coloring $\lambda^{\tau} \odot \lambda^{\sigma}$ for $x^{\tau} \odot x^{\sigma}$; but we must check that it is *monochrome*. The main observation is that every immediate causality $p_1^- \to_{\tau \odot \sigma} p_2^+$ in $x^{\tau} \odot x^{\sigma}$ arises from a sequence

$$p_1^- \to_{\tau \circledast \sigma} q_1 \to_{\tau \circledast \sigma} \ldots \to_{\tau \circledast \sigma} q_n \to_{\tau \circledast \sigma} p_2^+$$

where all the q_i s occur in B. In that case, all those causal links arise from immediate causal links in σ and τ [Cla24, Lemma 6.2.14], so the coloring $\lambda^{\tau} \circledast \lambda^{\sigma}$ is preserved along the chain and $(\lambda^{\tau} \circledast \lambda^{\sigma})(p_1^-) = (\lambda^{\tau} \circledast \lambda^{\sigma})(p_2^+)$, hence $(\lambda^{\tau} \odot \lambda^{\sigma})(p_1^-) = (\lambda^{\tau} \odot \lambda^{\sigma})(p_2^+)$ as well.

From right to left, fix a coloring λ on $x^{\tau} \odot x^{\sigma}$. The main observation we make is that every $q \in x^{\tau} \circledast x^{\sigma}$ is uniquely sandwiched between visible events. More precisely, there are unique p_1^-, p_2^+ visible, necessarily with the polarities indicated, and a unique causal path

$$p_1^- = q_0 \twoheadrightarrow_{\tau \circledast \sigma} \ldots \twoheadrightarrow_{\tau \circledast \sigma} q_{n+1} = p_2^+$$

such that q appears in the q_i s, and every q_1, \ldots, q_n occurs in B (is *hidden*) – this follows easily from Lemmas 6.1.16 and 6.2.15 from [Cla24], remarking that as σ and τ are *forestial*, so is $\tau \circledast \sigma$. This lets us complete λ to the interaction $x^{\tau} \circledast x^{\sigma}$ by setting $\lambda(q) = \lambda(p_1^-) = \lambda(p_2^+)$; this projects to matching λ^{σ} on x^{σ} and λ^{τ} on x^{τ} . These projected colorings are *monochrome*, because if *e.g.* $(p_1)^-_{\sigma} \to_{\sigma} (p_2)^+_{\sigma}$, then a straightforward case analysis shows that p_1 and p_2 must appear in the same sandwhich as above, and therefore receive the same color.

That these operations are inverses is an easy verification.

This immediately entails that the "colorful collapse" preserves composition; we shall see the proof just later. For preservation of copycat, the corresponding main lemma is:

Lemma 4.8. For any --board A and any $x \in \mathcal{C}(A)$, there is a bijection

$$\mathbf{c}_{-}$$
 : $\operatorname{col}(x) \simeq \operatorname{col}(\mathbf{c}_{x})$

such that for all $\lambda \in col(x)$, we have $\mathbf{c}_{\lambda} = [\lambda, \lambda]$ the co-pairing.

Proof. Immediate causal links from negative moves to positive moves in \mathbf{c}_x have shape

$$(i,a) \rightarrow \mathbf{c}_x (j,a)$$

where $i \neq j$ (see [Cla24, Lemma 6.4.3]). Thus, given $\lambda_A \in \operatorname{col}(x)$, then the co-pairing $\lambda = [\lambda_A, \lambda_A] : \mathbf{c}_x \to \mathcal{C}$ satisfies *monochrome*, therefore giving a valid coloring on \mathbf{c}_x . Reciprocally, for all $a \in x$, we always have either $(1, a) \to_{\mathbf{c}_x} (2, a)$ or $(2, a) \to_{\mathbf{c}_x} (1, a)$ depending on the polarity of a (see also [Cla24, Lemma 6.4.3]). Thus, if $\lambda \in \operatorname{col}(\mathbf{c}_x)$, we must have $\lambda(i, a) = \lambda(j, a)$ by *monochrome*, so that $\lambda = [\lambda_A, \lambda_A]$ for some $\lambda_A \in \operatorname{col}(x)$.

Putting these two lemmas together, we have:

Proposition 4.9. The colorful collapse defined above yields a functor:

 $\mathfrak{R}^{\mathcal{C}}(-) \quad : \quad \mathbf{DSInn} \quad \to \quad \mathbf{Rel}$

with a bijection $s_o: \mathfrak{R}^{\mathcal{C}}(\llbracket o \rrbracket_{\mathbf{DSInn}}) \simeq \mathcal{C}$.

Proof. This proposition builds on the functoriality of the earlier (colorless) collapse. In this proof, we leave implicit the material already covered in Sections 4.2 and 4.3, and focus on the preservation of colorings.

For preservation of copycat, consider $(\mathsf{x}_A, \mathsf{y}_A) \in \mathfrak{R}^{\mathcal{C}}(\mathbf{c}_A)$. By definition, it is witnessed by an experiment in \mathbf{c}_A , *i.e.* a pair $(\mathbf{c}_x, \lambda) \in \mathscr{C}^+_{\mathcal{C}}(\mathbf{c}_x)$ such that $(\mathbf{c}_x, \lambda)_l \in \mathsf{x}_A$ and $(\mathbf{c}_x, \lambda)_r \in \mathsf{y}_A$. By Lemma 4.8, $\lambda = [\lambda_A, \lambda_A]$ for $\lambda_A \in \operatorname{col}(x)$, so that $(\mathbf{c}_x, \lambda)_l = (x, \lambda_A) \in \mathscr{C}_{\mathcal{C}}(A)$ and $(\mathbf{c}_x, \lambda)_r = (x, \lambda_A)$ as well. Hence we have (x, λ_A) in both symmetry classes x_A and y_A , which must therefore be equal. Reciprocally, given $\mathsf{x}_A \in \mathfrak{R}^{\mathcal{C}}(A)$, taking $(x, \lambda) \in \mathscr{C}_{\mathcal{C}}(A)$ a representative, it is immediate by Lemma 4.8 again that $(\mathbf{c}_x, \mathbf{c}_\lambda) \in \mathscr{C}_{\mathcal{C}}^+(\mathbf{c}_A)$ provides an experiment which projects on the left and right to $(x, \lambda) \in \mathsf{x}_A$ as required.

For preservation of composition, consider first $(\mathsf{x}_A, \mathsf{x}_C) \in \mathfrak{R}^{\mathcal{C}}(\tau \odot \sigma)$. By definition, it is witnessed by an experiment in $\tau \odot \sigma$, which by Lemma 4.7 must have the form $(x^{\tau} \odot x^{\sigma}, \lambda^{\tau} \odot \lambda^{\sigma})$, where $(x^{\sigma}, \lambda^{\sigma}) \in \mathscr{C}^+_{\mathcal{C}}(\sigma)$ and $(x^{\tau}, \lambda^{\tau}) \in \mathscr{C}^+_{\mathcal{C}}(\tau)$ are compatible experiments, *i.e.* x^{σ}, x^{τ} are causally compatible and $\lambda^{\sigma}, \lambda^{\tau}$ are matching. Writing $x_B = x_B^{\sigma} = x_B^{\tau}$ and $\lambda_B = \lambda_B^{\sigma} = \lambda_B^{\tau}$, we form $(x_B, \lambda_B) \in \mathscr{C}_{\mathcal{C}}(B)$. Its symmetry class x_B is then in $\mathfrak{R}^{\mathcal{C}}(B)$, and $(\mathsf{x}_A, \mathsf{x}_B) \in \mathfrak{R}^{\mathcal{C}}(\sigma)$ is witnessed by the experiment $(x^{\sigma}, \lambda^{\sigma})$ while $(\mathsf{x}_B, \mathsf{x}_C) \in \mathfrak{R}^{\mathcal{C}}(\tau)$ is witnessed by the experiment $(x^{\tau}, \lambda^{\tau})$. Hence $(\mathsf{x}_A, \mathsf{x}_C) \in \mathfrak{R}^{\mathcal{C}}(\tau) \circ \mathfrak{R}^{\mathcal{C}}(\sigma)$ as required.

Reciprocally consider $(\mathbf{x}_A, \mathbf{x}_C) \in \mathfrak{R}^{\mathcal{C}}(\tau) \circ \mathfrak{R}^{\mathcal{C}}(\sigma)$, so there is $\mathbf{x}_B \in \mathfrak{R}^{\mathcal{C}}(B)$ such that $(\mathbf{x}_A, \mathbf{x}_B) \in \mathfrak{R}^{\mathcal{C}}(\sigma)$ and $(\mathbf{x}_B, \mathbf{x}_C) \in \mathfrak{R}^{\mathcal{C}}(\tau)$, respectively witnessed by $(x^{\sigma}, \lambda^{\sigma}) \in \mathscr{C}_{\mathcal{C}}(\sigma)$ and $(x^{\tau}, \lambda^{\tau}) \in \mathscr{C}_{\mathcal{C}}(\tau)$. Now, we must have $(x^{\sigma}_B, \lambda^{\sigma}_B) \in \mathbf{x}_B$ and $(x^{\tau}_B, \lambda^{\tau}_B) \in \mathbf{x}_B$, hence there must be $\theta_B : x^{\sigma}_B \cong_B x^{\tau}_B$ compatible with the coloring, *i.e.* $\lambda^{\tau}_B \circ \theta_B = \lambda^{\sigma}_B$. Now, by Proposition 3.13 (along with Lemma 4.1), there are $y^{\tau} \odot y^{\sigma} \in \mathscr{C}^+(\tau \odot \sigma)$ and $\varphi^{\sigma} : x^{\sigma} \cong_{\sigma} y^{\sigma}$ and $\varphi^{\tau} : x^{\tau} \cong_{\tau} y^{\tau}$ such that $\varphi^{\tau}_B \circ \theta_B = \varphi^{\sigma}_B$. Now the idea is to equip y^{σ} and y^{τ} with adequate colorings, to turn them into experiments: we simply set $\mu^{\sigma} = \lambda^{\sigma} \circ (\varphi^{\sigma})^{-1}$ and $\mu^{\tau} = \lambda^{\tau} \circ (\varphi^{\tau})^{-1}$. As φ^{σ} and φ^{τ} , $\mu^{\sigma}) \in \mathscr{C}^+_{\mathcal{C}}(\sigma)$ and $(y^{\tau}, \mu^{\tau}) \in \mathscr{C}^+_{\mathcal{C}}(\tau)$ are valid experiments. Additionally, by construction,

$$\mu_B^{\sigma} = \lambda_B^{\sigma} \circ (\varphi_B^{\sigma})^{-1} = \lambda_B^{\tau} \circ \theta_B \circ (\varphi_B^{\sigma})^{-1} = \lambda_B^{\tau} \circ (\varphi_B^{\tau})^{-1} = \mu_B^{\tau}$$

ensuring that μ^{σ} and μ^{τ} are matching, so that we can form the experiment $\mathcal{E} = (y^{\tau} \odot y^{\sigma}, \mu^{\tau} \odot \mu^{\sigma})$. Finally, this experiment indeed witnesses that $(\mathsf{x}_A, \mathsf{x}_C) \in \mathfrak{R}^{\mathcal{C}}(\tau \odot \sigma)$ since

$$\varphi_A^{\sigma}: (x^{\sigma}, \lambda^{\sigma})_A \cong_A (y^{\sigma}, \mu^{\sigma})_A = \mathcal{E}_A, \qquad \varphi_C^{\tau}: (x^{\tau}, \lambda^{\tau})_C \cong_C (y^{\tau}, \mu^{\tau})_C = \mathcal{E}_C$$

so that $\mathcal{E}_A \in \mathsf{x}_A$ and $\mathcal{E}_C \in \mathsf{x}_C$ as required.

4.4.4. Further structure. It remains to extend $\mathfrak{R}^{\mathcal{C}}(-)$ to a relative Seely functor. All the structural isomorphisms for that are the same as in the colorless case, colored as for copycat in the obvious way. Preservation of tensor and promotion are handled by two lemmas:

Lemma 4.10. Fix $\sigma : A \vdash B$, $\tau : C \vdash D$ strategies, and $x^{\sigma} \in \mathscr{C}^+(\sigma), x^{\tau} \in \mathscr{C}^+(\tau)$. Then, there is a bijection:

 $-\otimes -$: $\operatorname{col}(x^{\sigma}) \times \operatorname{col}(x^{\tau}) \simeq \operatorname{col}(x^{\sigma} \otimes x^{\tau})$

such that $(\lambda^{\sigma} \otimes \lambda^{\tau})_{A \otimes C} = \lambda^{\sigma}_{A} \otimes \lambda^{\tau}_{C}$ and $(\lambda^{\sigma} \otimes \lambda^{\tau})_{B \otimes D} = \lambda^{\sigma}_{B} \otimes \lambda^{\tau}_{D}$.

We omit the direct proof. Because of this lemma, it is straightforward that $\mathfrak{R}^{\mathcal{C}}(\sigma \otimes \tau)$ and $\mathfrak{R}^{\mathcal{C}}(\sigma) \otimes \mathfrak{R}^{\mathcal{C}}(\tau)$ coincide up to s^{\otimes} , as required by relative Seely functors.

Finally, we need a corresponding observation for the *promotion*:

Lemma 4.11. Fix $\sigma : !S \vdash A$, and $\lfloor x^{\sigma,i} \mid i \in I \rfloor \in \mathscr{C}^+(\sigma^{\dagger})$. Then, there is a bijection:

$$[-]$$
 : $\prod_{i \in I} \operatorname{col}(x^{\sigma,i}) \simeq \operatorname{col}(\lfloor x^{\sigma,i} \mid i \in I \rfloor)$

such that for all $(\lambda^{\sigma,i})_{i\in I} \in \prod_{i\in I} \operatorname{col}(x^{\sigma,i})$, writing $\lambda_{!S}^{\sigma,i} = \lfloor \lambda_j^{\sigma,i} \mid j \in J_i \rfloor$, we have

$$\lfloor \lambda^{\sigma,i} \mid i \in I \rfloor_{!S} = \lfloor \lambda_j^{\sigma,i} \mid \langle i,j \rangle \in \Sigma_{i \in I} J_i \rfloor, \lfloor \lambda^{\sigma,i} \mid i \in I \rfloor_A = \lfloor \lambda_A^{\sigma,i} \mid i \in I \rfloor.$$

Again there is no subtlety here besides the heavy notation, the coloring $\lfloor \lambda^{\sigma,i} \mid i \in I \rfloor$ is defined via the obvious co-pairing, and verifications are direct. Via this lemma, it follows that $\mathfrak{R}^{\mathcal{C}}(\sigma^{\dagger})$ and $\mathfrak{R}^{\mathcal{C}}(\sigma)^{\dagger}$ coincide up to $s^{!}$, as required by relative Seely functors.

Together with a few additional verifications for structural isomorphisms, this yields:

Theorem 4.12. The above provide the components for a relative Seely functor:

$$\mathfrak{R}^{\mathcal{C}}(-):\mathbf{DSInn}
ightarrow\mathbf{Rel}$$
 .

Again, it follows in particular that we also get $\mathfrak{R}^{\mathcal{C}}_{!}(-) : \mathbf{DSInn}_{!} \to \mathbf{Rel}_{!}$ a cartesian closed functor between the induced cartesian closed categories, so that:

Corollary 4.13. Consider $\Gamma \vdash M : A$ a simply-typed λ -term. Then, the following diagram commutes in **Rel**:



This is the same statement as in the colorless case, except that this fact we have fixed an arbitrary set C for the interpretation of the base type. Note however that neither **DSInn** nor its interpretation of the base type has changed: it is only the collapse that we changed.

5. FROM GAMES TO THE LINEAR SCOTT MODEL

Now, we have finally finished setting up the scene, and we can finally carry on with the journey announced in the introduction. So far, we have introduced the relational model **Rel**, the (relative) Seely category **DSInn** of thin concurrent games, and the relational collapse

$\mathfrak{R}(-):\mathbf{DSInn}\to\mathbf{Rel}$

that preserves the interpretation of the simply-typed λ -calculus. This collapse is *quantitative*: **Rel** still records the multiplicity of resource consumption. In the rest of this paper, we set to construct a corresponding *qualitative* collapse, targetting the *linear Scott model*.

In this section, we first recall the linear Scott model. Then, we construct *cartesian morphisms*, these morphisms between configurations that allow contraction and weakening of resources. We establish a few important properties of cartesian morphisms. Then, we will show that from a board equipped with cartesian morphisms we are able to construct a preorder, in a way compatible with the interpretation of types in the linear Scott model.

5.1. The Linear Scott Model. We first recall the linear Scott model, following [Ehr12].

5.1.1. The basic category. The linear Scott model can be presented in two different ways: as a category of (linear) functions between certain complete lattices, or as a category of relations between certain preordered sets. In this paper we pick the latter presentation, because it is more homogenous with the relational model and facilitates the relationship; but we shall also include a discussion about the domain-theoretic presentation.

We work with the following category:

Definition 5.1. ScottL has: (1) objects, preorders $(|A|, \leq_A)$; (2) morphisms from A to B, relations $\alpha \subseteq |A| \times |B|$ which are *down-closed*: if $(a, b) \in \alpha$ and $a \leq_A a', b \leq_B b'$, then $(a', b') \in \alpha$. Composition is relational composition, and identities $\mathrm{id}_A = \{(a, a') \mid a' \leq_A a\}$.

By a slight abuse of notation, we often write only A for the support set. If A is a preorder, we write A^{op} for the **opposite** preorder, with same support but reversed preorder. The **product** preorder has $(a, b) \leq_{A \times B} (a', b')$ iff $a \leq_A a'$ and $b \leq_B b'$.

As explained in the introduction, $a \leq_A a'$ expresses that the resources in a can be "contracted" into the resources in a'; and the resources in a' can be also "weakened" by not appearing in a. For instance, we shall see that if A is discrete, then

$$[a,a] \leq_{!A} [a,b]$$

where both copies of a are contracted into a and b is weakened. However, that intuition is incomplete as the preorder is crucially reversed in contravariant position.

5.1.2. Seely category. Firstly, **ScottL** has a symmetric monoidal structure: if A and B are preorders, then $A \otimes B = A \times B$. The monoidal unit is $1 = (\{\star\}, =)$, and the functorial action of \otimes is as in the relational model, while structural morphisms are the *down-closure* of their relational counterparts: for instance, for associativity we have

$$\begin{array}{rcl} \alpha_{A,B,C}^{\mathbf{ScottL}} & : & (A \otimes B) \otimes C & \cong & A \otimes (B \otimes C) \\ & & = & \{(((a,b),c) & , & (a',(b',c'))) & \mid & a' \leq_A a, \ b' \leq_B b, \ c' \leq_C c\} \end{array}$$

and likewise for the other components [Ehr12]. Likewise, the cartesian structure of **Rel** adapts to **ScottL** transparently, with A & B = A + B the disjoint union of the two preorders; $\top = (\emptyset, \emptyset)$ is terminal. The pairing operation is the same as in **Rel**, and projections in **ScottL** are obtained as the down-closure of those in **Rel**. We additionally set $A \multimap B = A^{\text{op}} \times B - \text{again}$, currying is as in **Rel**, and evaluation in **ScottL** is the down-closure of evaluation in **Rel**. Altogether, this makes **ScottL** a cartesian symmetric monoidal closed category.

Now, we get to the more critical definition of the exponential. Given A, we set

$$|!A| = \mathscr{M}_f(|A|), \qquad \mu \leq_{!A} \nu \ \Leftrightarrow \ \forall a \in \mathrm{supp}(\mu), \ \exists a' \in \mathrm{supp}(\nu), \ a \leq_A a'$$

where the support supp (μ) of a finite multiset $\mu \in \mathcal{M}_f(X)$ is simply the set of $x \in X$ with non-zero multiplicity. Again, this preorder is built on the same set as the exponential for the plain relational model. We specify the additional components of the exponential with:

$$\begin{array}{rcl} & & & |\alpha & = & \{(\mu,\nu) \in !A \times !B \mid \forall b \in \operatorname{supp}(\nu), \ \exists a \in \operatorname{supp}(\mu), \ (a,b) \in \alpha \} \\ & & & & \\ & & & & \\ & & & \\ & & & & \\ & & & & \\ & & &$$

and $see_1 : 1 \cong ! \top$ the obvious isomorphism. Altogether, we have:

Theorem 5.2. The above components make **ScottL** a Seely category.

Though the exponential is built from finite multisets, this model does not actually record quantitative information as morphisms are down-closed – note that !A is isomorphic (in **ScottL**) to the preorder obtained with the same definitions but built on finite sets instead of finite multisets. This Seely category **ScottL** will be the target of our qualitative collapse.

5.1.3. The Linear Scott Model and Scott Domains. Though the rest of the paper will mostly focus on **ScottL** as presented above, we take a small detour to present its relationship with the standard category of Scott domains and continuous functions.

Morphisms in ScottL as functions. If A is a preorder, we write $\mathcal{D}(A)$ for the set of down-closed subsets of |A|; it is a complete lattice. If $a \in A$, we write $[a]_A = \{a' \in A \mid a' \leq_A a\} \in \mathcal{D}(A)$; likewise if $X \subseteq |A|$, we write $[X]_A \in \mathcal{D}(A)$ for its down-closure.

We shall now construct a new category, which is an alternative presentation of **ScottL** in terms of functions. Its objects are the same as **ScottL**, *i.e.* preorders. A morphism from A to B is a *linear map* from $\mathcal{D}(A)$ to $\mathcal{D}(B)$, *i.e.* a function $f : \mathcal{D}(A) \to \mathcal{D}(B)$ such that

$$f(\bigcup_{i\in I} x_i) = \bigcup_{i\in I} (f(x_i))$$

for any family $(x_i)_{i \in I}$ of $x_i \in \mathcal{D}(A)$ – in particular, f is monotone and $f(\emptyset) = \emptyset$. We write **ScottFun** for the category with preorders as objects and linear maps.

The main observation here is that **ScottL** and **ScottFun** are isomorphic categories:

Proposition 5.3. The following constructions yield an isomorphism of categories:

$$\begin{array}{rcl} \operatorname{fun}_{A,B} & : & \operatorname{\mathbf{Scott}}\mathbf{L}[A,B] & \to & \operatorname{\mathbf{Scott}}\mathbf{Fun}[A,B] \\ & \alpha & \mapsto & x \in \mathcal{D}(A) \mapsto \alpha \, x \end{array}$$
$$\operatorname{tr}_{A,B} & : & \operatorname{\mathbf{Scott}}\mathbf{Fun}[A,B] & \to & \operatorname{\mathbf{Scott}}\mathbf{L}[A,B] \\ & f & \mapsto & \{(a,b) \mid b \in f([a]_A)\} \end{array}$$

where αx refers to relational composition, and where tr is called the **linear trace**.

A full subcategory of Scott domains. Now, we redo the analogous construction but for the Kleisli category ScottL_!. We noted above that for every preorder A, $\mathcal{D}(A)$ is a complete lattice; in fact it is a Scott domain, whose compact elements are the finitely generated elements of $\mathcal{D}(A)$, that is, those $x \in \mathcal{D}(A)$ such that $x = [X]_A$ for some finite $X \subseteq |A|$. Let us write ScottP for the category whose objects are preorders, and morphisms from A to B are Scott-continuous functions from $\mathcal{D}(A)$ to $\mathcal{D}(B)$. Writing also Scott for the usual category of Scott domains and Scott continuous functions, this yields a full and faithful (identity on morphisms) functor K: ScottP \hookrightarrow Scott.

Given $X \subseteq |A|$ for a preorder A, we set $X^! = \mathcal{M}_f(X)$. With this, we have:

Proposition 5.4. The following constructions yield an isomorphism of categories:

 $\begin{array}{rcl} \mathsf{Fun}_{A,B} & : & \mathbf{ScottL}_{!}[A,B] & \to & \mathbf{ScottP}[A,B] \\ & \alpha & \mapsto & x \in \mathcal{D}(A) \mapsto \alpha \, x^{!} \end{array}$

$$\begin{array}{rcl} \operatorname{Tr}_{A,B} & : & \operatorname{\mathbf{ScottP}}[A,B] & \to & \operatorname{\mathbf{ScottL}}_{!}[A,B] \\ & f & \mapsto & \{(\mu,b) \mid b \in f([\operatorname{supp}(\mu)]_{A})\} \end{array}$$

Altogether we get a full and faithful functor

$K \circ \mathsf{Fun} : \mathbf{ScottL}_! \to \mathbf{Scott}$

which is easily shown to preserve the cartesian closed structure; so that the Kleisli category $\mathbf{ScottL}_{!}$ is indeed a category of Scott domains and Scott-continuous functions.

Cartesian Morphisms. In Ehrhard's presentation of the linear Scott model [Ehr12], each type is interpreted as a preorder whose support set is nothing but the standard relational interpretation of the type. Thus, we expect to extract a preorder from a board A by simply attaching to $\Re(A)$ (and, later on, to $\Re^{\mathcal{C}}(A)$) an adequate preorder relation. As argued in the introduction, this preorder relation will arise from a notion of *cartesian morphisms* between configurations, maps allowing the contraction and weakening of resources.

We now move back to games, aiming for the definition of cartesian morphisms.

5.2. Mixed Boards. In the introduction, we sketched cartesian morphisms as certain forest morphisms leaving the identity of moves (*i.e.* the type component) unchanged, but that may alter *copy indices*. Unfortunately, we are not able to define such morphisms in the setting of *boards* as presented in Definition 3.15 – we need to refine our notion of game.

5.2.1. Arenas. Those boards arising from the interpretation of simple types have a specific shape – namely, they are themselves an *expansion* of a simpler structure called an *arena*:

Definition 5.5. An arena comprises $A = (|A|, \text{pol}_A, \leq_A)$ with |A| a countable set of moves, $\text{pol}_A : |A| \to \{-,+\}$ a polarity function, \leq_A a causality partial order, such that:

for all $a_1, a_2, a \in A $, if $a_1, a_2 \leq_A a$, then $a_1 \leq_A a_2$ or $a_2 \leq_A a_1$,
there is no infinite descending $a_1 >_A a_2 >_A a_3 >_A \ldots$,
if $a \in A $ is minimal for \leq_A , then $\text{pol}_A(a) = -$,
for all $a_1, a_2 \in A $, if $a_1 \twoheadrightarrow_A a_2$, then $\text{pol}_A(a_1) \neq \text{pol}_A(a_2)$.

Additionally we fix, for each move $a \in A$, a set Ind(a) which is either \mathbb{N} or $\{*\}$.

This is close to the usual notion of Hyland-Ong arenas, with a slight change in presentation so as to remain close to boards. The new component is the set Ind(a), which specifies which moves are duplicable by specifying the admissible copy indices¹⁰.

We briefly present the main constructions on arenas. First, for the atom, we write \underline{o} for the arena with exactly one (negative) move \mathbf{q}^- , with $\operatorname{Ind}(\mathbf{q}^-) = \{*\}$: this move is not duplicable. If A is an arena, the *exponential* !A has the same components as A (we do not duplicate the moves), but we set $\operatorname{Ind}_{!A}(a^-) = \mathbb{N}$ for every a^- minimal: we set the initial moves as duplicable. The *parallel composition* $A \parallel B$ adapts transparently to arenas, with the $\operatorname{Ind}(-)$ function inherited. Note that any arena may be written as $A \cong \prod_{i \in I} A_i$ with A_i well-opened. If A and B are arenas with B well-opened, then $A \multimap B$ is an arena (again with Ind inherited); this extends to B not well-opened with

$$A \multimap (\parallel_{i \in I} B_i) = \parallel_{i \in I} A \multimap B_i.$$

Altogether, this lets us interpret any simple type as an arena with $\llbracket o \rrbracket_{\mathbf{Ar}} = \underline{o}$ and $\llbracket A \to B \rrbracket = ! \llbracket A \rrbracket_{\mathbf{Ar}} \multimap \llbracket B \rrbracket_{\mathbf{Ar}}$: moves are not explicitly duplicated, but simply marked with

¹⁰This data is not required for the interpretation of simple types, where all non-initial moves are necessarily duplicable. We need this here because we wish to phrase the collapse as a relative Seely functor, also handling the linear structure; and hence some moves may not be duplicable in the arena.



Figure 9: Interpretation of $(o \rightarrow o) \rightarrow o \rightarrow o$ as an arena

the admissible copy indices. We show in Figure 9 the interpretation of the simple type $(o \rightarrow o) \rightarrow o \rightarrow o$ as an arena – the reader familiar with Hyland-Ong games will recognize here the familiar arena for that type, where each move corresponds to an atom occurrence.

5.2.2. Mixed boards. Mixed boards are boards as in Definition 3.15, except that moves are labeled by moves from an underlying arena, and various conditions are satisfied to ensure the links between the two¹¹. This leads to the following slightly bulky definition – here pred(-) denotes the unique predecessor of a non-minimal move, exploiting that boards are *forestial*:

Definition 5.6. A mixed board is (A, \underline{A}) with A a --board, \underline{A} an arena, with

$$|\mathsf{bl}_A:|A| \to |\underline{A}|, \qquad \quad \mathsf{ind}_A:|A| \to \mathbb{N} \uplus \{*\}$$

with lbl_A a **label** function preserving polarities and ind_A an **indexing** function such that $\mathsf{ind}_A(a) \in \mathsf{Ind}(a)$ for all $a \in A$, satisfying the following additional conditions:

rigid:	lbl_A preserves and reflects minimality, and preserves \rightarrow ,
transparent:	for any $x, y \in \mathscr{C}(A)$ and bijection $\theta : x \simeq y$,
	then $\theta \in \mathscr{S}(A)$ iff θ is an order-iso preserving lbl_A ,
local conflict:	if $a_1 \sim a_2$, we have $pred(a_1) = pred(a_2)$ and $ind(a_1) = ind(a_2)$,
invariant conflict:	if $a_1 \sim a_2$, $lbl(a_1) = lbl(b_1)$, $lbl(a_2) = lbl(b_2)$,
	$pred(b_1) = pred(b_2)$ and $ind(b_1) = ind(b_2)$, then $b_1 \sim b_2$.
jointly injective:	for $a_1, a_2 \in A$, if $lbl(a_1) = lbl(a_2)$, $ind(a_1) = ind(a_2)$,
	and $\operatorname{pred}(a_1) = \operatorname{pred}(a_2)$, then $a_1 = a_2$.
wide:	for any $a \in A$, $\underline{b} \in \underline{A}$, $lbl(a) \twoheadrightarrow_{\underline{A}} \underline{b}$, and $i \in Ind_A(\underline{b})$,
	there is $b \in A$ s.t. $pred(b) = a$, $lbl(b) = \underline{b}$ and $ind(b) = i$,
+-transparent:	for $\theta : x \cong_A y, \theta \in \mathscr{S}_+(A)$ iff for all $a^- \in x$, $\operatorname{ind}(\theta(a)) = \operatorname{ind}(a)$,
--transparent:	for $\theta : x \cong_A y, \theta \in \mathscr{S}_{-}(A)$ iff for all $a^+ \in x$, $\operatorname{ind}(\theta(a)) = \operatorname{ind}(a)$.

A mixed board is strict if A is strict and $\operatorname{Ind}(\underline{a}) = \{*\}$ for every $\underline{a} \in \underline{A}$ minimal.

This looks complicated, but those conditions are really simple structural properties expressing that the board is an expansion of the arena¹².

 $^{^{11}{\}rm Mixed}$ boards were developed in [CC21, Cla24] to make explicit the link between the interpretation of types in concurrent games and in standard Hyland-Ong games.

¹²It is not quite the case that the board is determined (up to iso) by the arena, because the arena lacks the information about which moves are related additively and multiplicatively – it conflates A & B with $A \otimes B$. In principle this could be added, but this does not seem worth the even heavier definition.



Figure 10: A structural map on $(o_1 \rightarrow o_2) \rightarrow o_3 \rightarrow o_4$

The mixed board for the board game is (o, \underline{o}) with $\mathsf{lbl} : \mathsf{q}^- \mapsto \mathsf{q}^-$ and $\mathsf{ind} : \mathsf{q}^- \mapsto * - \mathsf{by}$ abuse of notations, we shall denote this mixed board by o. The **tensor** of mixed boards has $\underline{A \otimes B} = \underline{A} \parallel \underline{B}$, with components inherited. The **with** of strict mixed boards S and Tis defined likewise. The **bang** of strict S has $\underline{!S} = \underline{!S}$ (*i.e.* with just $\mathsf{Ind}(\underline{s}) = \mathbb{N}$ for $\underline{s} \in \underline{S}$ minimal and other components unchanged). The **linear arrow** $A \multimap S$ is extended to mixed boards in the obvious way. However, note that only the objects of our category will be mixed boards – we do not give a mixed board construction for the hom-game.

5.3. Cartesian Morphisms. On mixed boards, we may now define cartesian morphisms.

5.3.1. Structural maps. From now on, fix a mixed board A. Cartesian maps are motivated as variations of symmetries that can additionally contract and weaken resources. The condition transparent characterises those as order-isomorphisms which leave the label unchanged; by weakening order-isomorphism to simply forest morphism, we obtain the following notion:

Definition 5.7. A structural map is a function $f: x \to y$, for $x, y \in \mathscr{C}(A)$, satisfying

min-preserving: for all $a \in \min(x)$, then $f a \in \min(y)$, \rightarrow -preserving: for all $a \rightarrow_A b$, then $f a \rightarrow_A f b$, *label-preserving:* for all $a \in x$, $\mathsf{lbl}_A(f a) = \mathsf{lbl}_A a$.

We write $f: x \rightsquigarrow y$ to indicate that $f: x \rightarrow y$ is a structural map.

In Figure 10 we give an example of a structural map, where $q_{3,2}^+$ and $q_{3,6}^+$ are sent to themselves, and the other assignments are forced. Note that all copy indices can be changed freely. The structural map contracts both positive and negative moves, while $q_{3,7}^+$ is not reached – it is regarded as *weakened* by this structural map.

Structural maps form a category, and one can consider the associated preorder with configurations as elements, and $x \leq y$ iff there is some structural map $f: x \rightsquigarrow y$. However, this preorder is not actually the one we need, because it is not compatible with the linear arrow construction of preorders. Indeed, recall that in **ScottL**, the linear arrow was

$$A \multimap B = A^{\mathrm{op}} \times B$$

contravariant on the left hand side, whereas structural maps on $A \multimap B$ are covariant on both sides. To recover the appropriate variance, we must take polarities into account:



Figure 11: A cartesian morphism

Definition 5.8. For a structural map $f: x \rightsquigarrow y$, we define the conditions:

total:	if $a^+ \in x$, for all $f a^+ \rightarrow b^-$ in y, there is $a^+ \rightarrow c^-$ in x s.t. $f c^- = b^-$;
	and for all b^- minimal in y there is c^- in x such that $f c^- = b^-$.
+-total:	if $a^- \in x$, for all $f a^- \to b^+$, there is $a^- \to c^+$ in x s.t. $f c^+ = b^+$,
<i>preserving:</i>	if $a^- \in x$, $\operatorname{ind}_A(f a) = \operatorname{ind}_A a$,
+-preserving:	if $a^+ \in x$, $\operatorname{ind}_A(f a) = \operatorname{ind}_A a$,

we call a structural map **positive** iff it is *--preserving* and *--total*; we call it **negative** iff it is *+-preserving* and *+-total*. For these notions, we use notations $f: x \stackrel{t}{\rightarrow} y$ and $f: x \stackrel{\tau}{\rightarrow} y$.

Intuitively, positive structural maps can only contract positive moves: *--preserving* ensures that they cannot contract negative moves (as the copy index is preserved), and *--total* entails that they cannot weaken negative moves: negative extensions must have a pre-image. Dually, negative structural maps can only contract and weaken negative moves.

Structural maps generalize symmetries, in a way compatible with polarities:

Lemma 5.9. Any symmetry $\theta : x \cong_A y$ is also a structural map $\theta : x \rightsquigarrow y$. Moreover, θ is positive (resp. negative) as a symmetry iff it is positive (resp. negative) as a structural map.

Proof. For the first part, note that by *transparent*, θ is an order-isomorphism preserving Ibl_A . It is thus evident from the definition that it is also a structural map. For the second part, if $\theta : x \cong_A^+ y$, then by +-*transparent* it is --*preserving*. It is also --*total* because it is an order-iso. The other implication proceeds similarly, and the negative case is symmetric. \Box

5.3.2. *Cartesian morphisms*. We are now in position to define our *cartesian morphisms*, which take positive structural maps covariantly and negative structural maps contravariantly.

Definition 5.10. A cartesian morphism $\chi : x \leftrightarrow y$ is any composite relation:

 $x = x_1 \stackrel{+}{\rightsquigarrow} x_2 \stackrel{-}{\leftarrow} x_3 \dots x_{n-2} \stackrel{+}{\rightsquigarrow} x_{n-1} \stackrel{-}{\leftarrow} x_n = y$ where $x_1, \dots, x_n \in \mathscr{C}(A)$.

A cartesian morphism $\chi : x \Leftrightarrow y$ is a relation between x and y, *i.e.* $\chi \subseteq x \times y$, but it is in general not functional in either direction. This, of course, is unavoidable: basic contractions are functional, but they are taken covariantly or contravariantly. We give in Figure 11 an example of a cartesian morphism (we do not give the exact definition of structural maps to avoid overloading the picture, but they are almost uniquely defined). Note that there is also a cartesian morphism from the rightmost diagram to the leftmost one, so that these two are considered equivalent, qualitatively. The definition of cartesian morphism obviously yields a category, but the first property we aim for is that every cartesian morphism factors uniquely as the notation $\overleftarrow{}$ suggests. But this is not a trivial fact, establishing it requires building up a number of prerequisites.

We start with this easy property of structural maps:

Lemma 5.11. Consider $f : x \rightsquigarrow y$ a structural map, and $(a, b) \in f$ with both non-minimal. Then, f(pred(a)) = f(pred(b)) as well.

Proof. Immediate from \rightarrow -preserving and the fact that the game A is forestial.

Next we show a uniqueness property of intermediate witnesses for cartesian morphisms:

Lemma 5.12. Consider $\chi : x \stackrel{\text{art}}{\longleftrightarrow} y$ a cartesian morphism, obtained through the chain:

$$x = x_1 \quad \stackrel{f_1^+}{\rightsquigarrow} \quad y_1 \quad \stackrel{f_1^-}{\longleftarrow} \quad x_2 \quad \dots \quad x_n \quad \stackrel{f_n^+}{\rightsquigarrow} \quad y_n \quad \stackrel{f_n^-}{\longleftarrow} \quad x_{n+1} = y \,.$$

Then, for any $(a_1, a_{n+1}) \in \chi$, there is a unique sequence of witnesses

x_1	$f_1^+ {\sim}$	y_1	$f_1^- \longleftrightarrow$	x_2	 x_n	$\stackrel{f_n^+}{\leadsto} i$	$J_n \stackrel{f_n^-}{\nleftrightarrow}$	x_{n+1}
Ψ		Ψ		Ψ	 Ψ	1	Ψ	Ψ
a_1		b_1		a_2	 a_n	ł	\mathcal{D}_n	a_{n+1}

such that for every $1 \leq i \leq n$, $(a_i, b_i) \in f_i^+$ and $(a_{i+1}, b_i) \in f_i^-$.

Proof. Existence is obvious by definition of relational composition, we prove uniqueness. Note that since structural maps preserve minimality and \rightarrow , all moves in such as sequence must have the same *depth*, where the **depth** of *a* minimal is 0, and the depth of *b* where $a \rightarrow b$ is depth(*b*) = depth(*a*) + 1. It follows that χ preserves depth as well. Thus, we prove the uniqueness of the chain above by induction on depth(a_1) = depth(a_{n+1}).

Assume the depth is 0. For any $1 \le i \le n-1$, if $a_{i+1} \in \min(x_{n+1})$ is fixed, necessarily negative since A is negative, then there is a unique $b_i \in y_n$ such that $(a_{i+1}, b_i) \in f_i^-$ since f_i^- is a function. Now, in turn, we show that there is a unique $a_i \in \min(x_i)$ such that $(a_i, b_i) \in f_i^+$. It must be minimal, or contradict the minimality of b_i since f_i^+ is \rightarrow -preserving. Consider then another a'_i such that $(a'_i, b_i) \in f_i^+$ also; necessarily also minimal. Then, $|\mathsf{bl}(a'_i) = |\mathsf{bl}(b_i)$ since f_i^+ is *label-preserving*. Additionally, $\mathsf{ind}(a'_i) = \mathsf{ind}(b_i)$ since f_i^+ is -preserving - as all elements in this chain are negative. But then, a'_i and a_i are both minimal, with same label and copy index, thus $a'_i = a_i$ by *jointly injective*.

Now, consider $(a_1, a_{n+1}) \in \chi$ of depth d + 1. Since A is forestial, there are unique $c_1 \rightarrow a_1$ and $c_{n+1} \rightarrow a_{n+1}$. By Lemma 5.11, we have $(c_1, c_{n+1}) \in \chi$ as well, and by induction hypothesis, there is a unique sequence of witnesses

for this. Now for negative extensions, we reason as above from right to left: if $c_{i+1} \rightarrow a_{i+1}$ negative, there is a unique $d_i \rightarrow b_i$ such that $(a_{i+1}, b_i) \in f_i^-$ since it is a function; and then there is a unique $c_i \rightarrow a_i$ s.t. $(a_i, b_i) \in f_i^+$ by *--total*, uniqueness being by *label-preserving* and *--preserving* for f_i^+ , along with *jointly injective* for A. For positive extensions, the reasoning is symmetric from left to right, concluding the proof.

From this unique witness property, we derive the following lifting property:

Lemma 5.13. For any structural map $\chi : x \stackrel{\sim}{\leftarrow} y$, we have the following lifting properties:

- (1) if $b_2^- \in \min(y)$, then there is a unique $b_1^- \in \min(x)$ such that $(b_1^-, b_2^-) \in \chi$,
- (2a) If $(a_1, a_2) \in \chi$, then for all $a_2 \rightarrow_A b_2^-$ in y, there is a unique $a_1 \rightarrow b_1^-$ in x such that $(b_1, b_2) \in \chi$,
- (2b) If $(a_1, a_2) \in \chi$, then for all $a_1 \to_A b_1^+$ in x, there is a unique $a_2 \to b_2^+$ in y such that $(b_1, b_2) \in \chi$.

Proof. Since A is negative, there is only one case for transporting minimal events. By definition, χ is obtained as a composition of structural maps:

$$c = x_1 \quad \stackrel{f_1^+}{\rightsquigarrow} \quad y_1 \quad \stackrel{f_1^-}{\longleftarrow} \quad x_2 \quad \dots \quad x_n \quad \stackrel{f_n^+}{\rightsquigarrow} \quad y_n \quad \stackrel{f_n^-}{\longleftarrow} \quad x_{n+1} = y$$

(1) For $1 \le i \le n$, for any $a_{i+1}^- \in x_{n+1}$ minimal, then there is a unique $b_i^- \in y_n$ such that $(a_{i+1}^-, b_i^-) \in f_i^-$, because it is a function; and b_i^- is minimal since f_i^- preserves minimality. Then, there is some $a_i^- \in x_i$ such that $(a_i^-, b_i^-) \in f_i^+$ by *--total*. Again it must be minimal: a predecessor would yield a predecessor for b_i^- by \rightarrow -preserving. It is unique since f_i^+ is label-preserving and --preserving and by jointly injective for A. Iterating this, we get a unique sequence of witnesses terminating in a_{n+1} , and in particular unique $(a_1, a_{n+1}) \in \chi$.

(2a) Now consider $(a_1, a_{n+1}) \in \chi$, and $a_{n+1} \rightarrow c_{n+1}^-$. By Lemma 5.12, there is

a unique sequence of witnesses. Note that by Lemma 5.11, if there is indeed some $a_1 \rightarrow c_1^$ such that $(c_1^-, c_{n+1}^-) \in \chi$, then its sequence of witnesses (unique by Lemma 5.12) must be above the (unique) sequence of witnesses for (a_1, a_{n+1}) . Hence it suffices to show that there is a unique sequence of witnesses above the above ending in c_{n+1}^{-} , and this is what we shall do. Thus assume $a_{i+1} \rightarrow c_{i+1}^-$. As just above, there is a unique $(c_{i+1}, d_i^-) \in f_i^-$ because it is a function, and we do have $b_i \rightarrow d_i^-$ by \rightarrow -preserving. Then, there is a unique $(c_i^-, d_i^-) \in f_i^+$ by --total, --preserving, label-preserving, \rightarrow -preserving for f_i^+ and jointly injective for A. Iterating this we get a unique sequence of witnesses terminating in c_{n+1}^{-} above the unique sequence witnessing (a_1, a_{n+1}) ; and in particular unique (c_1, c_{n+1}) as required.

(2b) The reasoning is symmetric.

Now, we are almost ready to prove our factorization result. As we will need to build structural maps gradually, we need to cover intermediate cases where total constraints are not satisfied, and shall therefore need the following notions:

Definition 5.14. A partial positive map, written $f: x \stackrel{\neq p}{\to} y$, is a structural map satisfying --preserving (but not --total). Likewise, a partial negative map, written $f:\overline{x}$, y, is a structural map satisfying +-preserving (but not +-total).

These satisfy the following lifting lemmas:

Lemma 5.15. Consider $f: x \stackrel{\pm p}{\leadsto} y$ a partial positive structural map. Then: (1) for $y \vdash_A b_2^-$ minimal, there is a unique $x \vdash_A b_1^-$ minimal such that

$$f \uplus \{(b_1^-, b_2^-)\} : x \uplus \{b_1^-\} \stackrel{\scriptscriptstyle \pm p}{\rightsquigarrow} y \uplus \{b_2^-\}$$

(2) for
$$a \in x$$
, for $y \vdash_A b_2^-$ with $f a \twoheadrightarrow_A b_2^-$, there is a unique $x \vdash_A b_1^-$ with $a \twoheadrightarrow_A b_1^-$ s.t
 $f \uplus \{(b_1^-, b_2^-)\} : x \uplus \{b_1^-\} \stackrel{\stackrel{\scriptscriptstyle p}}{\to} y \uplus \{b_2^-\}$

Proof. We detail (2), as (1) is similar but simpler.

Existence. First, we note that f is label-preserving, so that $|\mathsf{b}| a = |\mathsf{b}| (f a)$. Thus, by rigid, $|\mathsf{b}| a \to_{\underline{A}} |\mathsf{b}| b_2^-$. Thus by wide, there is indeed some b_1^- such that $a \to_A b_1^-$, $|\mathsf{b}| (b_1^-) = |\mathsf{b}| (b_2^-)$ and $\mathsf{ind}(b_1^-) = \mathsf{ind}(b_2^-)$. Now, we must justify that $x \vdash_A b_1^-$. Clearly, $x \uplus \{b_1^-\}$ is down-closed. If it was inconsistent, there would be some $b \in x$ such that $b \sim b_1^-$, in which case $a \to b$ and $\mathsf{ind}(a) = \mathsf{ind}(b)$ by local conflict. Since A is alternating, b is negative as well, and f is --preserving, so that by invariant conflict we have $fb \sim b_2^-$ as well, contradiction.

Uniqueness. Straightforward by jointly injective and that f is positive.

Lemma 5.16. Consider $f : x \stackrel{\sim p}{\sim} y$ a negative structural map. Then for all $a \in x$, for all $y \vdash_A b_2^+$ with $f a \rightarrow_A b_2^+$, there is a unique $x \vdash_A b_1^+$ with $a \rightarrow_A b_1^+$ such that

$$f \uplus \{(b_1^+, b_2^+)\} : x \uplus \{b_1^+\} \stackrel{\sim_{\mathcal{B}}}{\to} y \uplus \{b_2^+\}$$

Proof. Same as for Lemma 5.15.

Now, we are finally in position to prove our factorization result. For the proof, we introduce the following notation: if $R \subseteq A \times B$ is a relation from A to B, we write $R^{\perp} \subseteq B \times A$ for the reverse relation. We will also apply this to functions, regarded as functional relations.

Lemma 5.17. Consider A a mixed board, and $\chi : x \stackrel{\neg}{\leftrightarrow} z$ a cartesian morphism. Then, there are unique $y \in \mathscr{C}(A), \chi_{-} : y \stackrel{\neg}{\rightarrow} x$ and $\chi_{+} : y \stackrel{\neg}{\rightarrow} z$, such that



where the bottom path is composed relationally, i.e. we ask $\chi_+ \circ \chi_-^{\perp} \subseteq \chi$. Finally, the inclusion is actually an equality.

Proof. First, remark that we can find y, χ_{-} and χ_{+} such that

$$\begin{array}{c} x \xleftarrow{(\overset{\chi}{-}+) \longrightarrow} z \\ \chi^{(\overset{\sim}{-}p)} \cup \downarrow (\overset{\downarrow}{+}p) \\ y \end{array}$$

with $\chi_{-}: y \stackrel{\sim p}{\sim} x$ and $\chi_{+}: y \stackrel{+p}{\sim} z$ – indeed, it suffices to take y and χ_{-}, χ_{+} to be empty. Let us call such data a *solution*; solutions are partially ordered by componentwise inclusion. There is an upper bound to the size of y, because it is a tree whose depth is bounded by both that of x and z, and its width is bounded by the maximum of those of x and z (by --preserving and +-preserving of χ_{+} and χ_{-}). Hence, there is a solution y, χ_{-}, χ_{+} where yhas maximal size. We show that it satisfies the conditions of the lemma.

We must first check that χ_+ and χ_- are not partial. First, we show that χ_+ is --total. Thus take $a^+ \in y$, and $a_2^+ = \chi_+ a^+ \Rightarrow b_2^-$ in z. Assume, seeking a contradiction, that b_2^- has no lifting in y. There is a candidate for the lifting: by wide, there is some $a^+ \Rightarrow b^-$ such that $|\mathsf{bl}(b^-) = |\mathsf{bl}(b_2^-)$ and $\mathsf{ind}(b^-) = \mathsf{ind}(b_2^-)$, and it is unique by jointly injective. Moreover, we have $y \vdash_A b^-$, otherwise there is $c^- \sim b^-$, and by local conflict we also have $a^+ \Rightarrow c^-$

and $\operatorname{ind}(c^{-}) = \operatorname{ind}(b^{-})$. But as χ_{+} is *--preserving* and preserves labels, by *invariant conflict* we have $\chi_{+} c^{-} \cdots \chi_{+} b^{-} = b_{2}^{-}$, but $\chi_{+} c^{-} \in z$, which contradicts $z \in \mathscr{C}(A)$. Hence, we form

$$y \uplus \{b^-\} \in \mathscr{C}(A) \,, \qquad \chi_+ \uplus \{(b^-, b_2^-)\} : y \uplus \{b^-\} \stackrel{\scriptscriptstyle \pm p}{\scriptstyle \leftarrow} z \,.$$

To extend χ_{-} accordingly, write $a_{1}^{+} = \chi_{-} a^{+}$. By hypothesis, we have $(a_{1}^{+}, a_{2}^{+}) \in \chi$. Hence, by Lemma 5.13, there is a (unique) $a_{1}^{+} \rightarrow b_{1}^{-}$ such that $(b_{1}^{-}, b_{2}^{-}) \in \chi$. We form

$$\chi_{-} \uplus \{ (b^{-}, b_{1}^{-}) \} : y \uplus \{ b^{-} \} \stackrel{\neg p}{\leadsto} z$$

which is clearly a partial negative structural map. The required inclusion is still satisfied, so this contradicts the maximality of the solution y, χ_{-}, χ_{+} – the case where b^{-} is minimal is similar but simpler. The proof that χ_{-} is +-total is symmetric.

For *uniqueness*, consider we have



and show $y' \subseteq y$, and that for every $a \in y'$, we have $\chi_{-} a = \xi_{-} a$ and $\chi_{+} a = \xi_{+} a$.

Consider $a \in y'$ minimal not satisfying this property. If it is minimal in A then it is in particular negative. Then write $a' = \xi^+ a$; by definition of structural maps we have a'minimal with $|\mathbf{b}| a' = |\mathbf{b}| a$, by --preserving we have $\mathrm{ind} a' = \mathrm{ind} a$ – so that in fact a = a'by jointly injective. And by --total, there is $a'' \in y$ such that $\chi_+ a'' = a$, and for the same reason a'' = a, so that $a \in y$. Moreover, we have seen that $\chi_+ a = \xi_+ a = a$; but we also know that $(\chi_- a, \chi_+ a), (\xi_- a, \xi_+ a) \in \chi$ by hypothesis. But now, by Lemma 5.13, there is a unique b such that $(b, a) \in \chi - \mathrm{so} \ \chi_- a = \xi_- a = b$ as required. Now if $a \in y'$ is not minimal in A, it has a unique predecessor $b \rightarrow_A a$, for which we know by induction that $b \in y$ and $\chi_- b = \xi_- b$ and $\chi_+ b = \xi_+ b$. Based on that, if a is negative, we can directly replay the argument above; and the symmetric argument if a is positive. In particular, this shows that $y' \subseteq y$ with $\xi_-, \chi_-, \xi_+, \chi_+$ compatible with the inclusion. But the argument is symmetric, so we also have $y \subseteq y'$ – which concludes uniqueness.

Finally, we must show that the inclusion between the two sides of the diagram is actually an equality. Consider $(b_1, b_2) \in \chi$ minimal (note that they must have the same depth) such that $(b_1, b_2) \notin \chi_+ \circ \chi_-^{\perp}$. If (b_1, b_2) are minimal, then they are negative (b_1^-, b_2^-) . By --total for χ_+ , there is some $b^- \in y$ such that $\chi_+ b^- = b_2^-$. Writing $b_1' = \chi_- b^-$, we must have $(b_1', b_2) \in \chi$, but by Lemma 5.13 this implies $b_1' = b_1^-$, so $(b_1^-, b_2^-) \in \chi_+ \circ \chi_-^{\perp}$ after all, contradiction. So now if they are not minimal, there are $c_1 \rightarrow b_1$ and $c_2 \rightarrow b_2$ with $(c_1, c_2) \in \chi$ (by Lemma 5.11) and $(c_1, c_2) \in \chi_+ \circ \chi_-^{\perp}$ (by minimality of (b_1, b_2)), so there is $c \in y$ such that $\chi_+ c = c_2$ and $\chi_- c = c_1$. Now we distinguish by cases. If b_1, b_2 are negative, then exploiting c, the reasoning is the same as for b_1, b_2 minimal above, using the same lemmas – and if b_1, b_2 are positive, then the reasoning is symmetric.

From this lemma, it follows in particular that as accounced, every cartesian morphism $\chi: x \leftrightarrow y$ can be written uniquely as the relational composition $x \leftrightarrow dy$.



Figure 13: Colored preorder-isos

5.4. Reconstructing the Preorder. Now, in this section we reconstruct a preorder from these cartesian morphisms, and show how this is compatible with all the constructions on preorders involved in the relative Seely category structure. Fix here a mixed board A.

5.4.1. Uncolored case. From the relational collapse, we know that we must equip the set $\mathfrak{R}(A)$ of positions of A defined in (4.1) with a preorder derived from cartesian morphisms. The obvious route is to start is to define it on configurations: for $x, y \in \mathscr{C}(A)$, we set

$$x \stackrel{+-}{\leadsto} y \qquad \Leftrightarrow \qquad \exists \chi : y \stackrel{-+}{\leadsto} x,$$

noting the inversion in directions. This is compatible with symmetry:

Lemma 5.18. Consider $x, x', y, y' \in \mathscr{C}(A)$ such that $x \cong_A x'$ and $y \cong_A y'$. Then, $x \xleftarrow{} y$ iff $x' \xleftarrow{} y'$.

Proof. Fix a cartesian morphism $\chi : y \stackrel{\leftarrow}{\leftrightarrow} x$ which by Lemma 5.17 factors as $\chi_+ \circ \chi_-^{\perp}$ for $\chi_- : z \stackrel{\leftarrow}{\rightarrow} y$ and $\chi_+ : z \stackrel{\leftrightarrow}{\rightarrow} x$. Fix also symmetries $\theta : x \cong_A x'$ and $\vartheta : y \cong_A y'$, by Lemma 3.9 they factor as $\theta = \theta_+ \circ \theta_-$ and $\vartheta = \vartheta_+ \circ \vartheta_-$ for θ_+, ϑ_+ positive symmetries and θ_-, ϑ_- negative symmetries. Now, with all this data we can form

$$x' \stackrel{ heta_+}{\Leftarrow} \cdot \stackrel{ heta_-^{-1}}{\longrightarrow} x \stackrel{ ilde \chi_+}{\leftarrow} \cdot \stackrel{ ilde \chi_-}{\longrightarrow} y \stackrel{ ilde \vartheta_-^{-1}}{\longrightarrow} y$$

converting symmetries to structural maps by Lemma 5.9; so that $x' \nleftrightarrow y'$.

We write $x \leftrightarrow y$ if for any $x \in x$ and $y \in y$, we have $x \leftrightarrow y$ – by the lemma above, this does not depend on the choice of x and y, and immediately forms a preorder. Thus:

Definition 5.19. For any mixed board A, we set $\mathfrak{S}(A) = (\mathfrak{R}(A), \stackrel{+-}{\nleftrightarrow})$.

This is compatible with all the constructions on preorders involved in the relative Seely structure of **ScottL**. More precisely, the bijections of Figure 6 can be verified to be compatible with the preorder, *i.e.* to yield preorder-isomorphisms as described in Figure 12. For most of them, the corresponding verification is straightforward; we just detail two.

Lemma 5.20. Fix A, S mixed boards with S strict, $x_A \multimap x_S, y_A \multimap y_S \in \mathfrak{R}(A \multimap S)$. Then, $x_A \multimap x_S \xleftarrow{\to} y_A \multimap y_S$ iff $y_A \xleftarrow{\to} x_A$ and $x_S \xleftarrow{\to} y_S$.

Proof. By Lemma 5.18, it suffices to reason on representatives.

If. Fix y_A, x_A, x_S, y_S such that $y_A \stackrel{\leftarrow}{\leftarrow} x_A$ and $x_S \stackrel{\leftarrow}{\leftarrow} y_S$. By definition, this means that there are $\chi_A : x_A \stackrel{\leftarrow}{\leftarrow} y_A$ and $\chi_S : y_S \stackrel{\leftarrow}{\leftarrow} x_S$. By Lemma 5.17, there are $\chi_A^- : z_A \stackrel{\leftarrow}{\to} x_A$ and $\chi_A^+ : z_A \stackrel{\leftarrow}{\to} y_A$, with also $\chi_S^- : z_S \stackrel{\leftarrow}{\to} y_S$ and $\chi_S^+ : z_S \stackrel{\leftarrow}{\to} x_S$. We may then form

$$\chi_A^+ \multimap \chi_S^- : z_A \multimap z_S \stackrel{\sim}{\twoheadrightarrow} y_A \multimap y_S, \qquad \chi_A^- \multimap \chi_S^+ : z_A \multimap z_S \stackrel{\prec}{\twoheadrightarrow} x_A \multimap x_S$$

defined in the obvious way – observe the inversion of polarities, to match that A is dualized on the left hand side. Altogether, we have $x_A \multimap x_S \xleftarrow{+-} y_A \multimap y_S$ as required.

Only if. Assuming $x_A \multimap x_S \stackrel{+-}{\longleftrightarrow} y_A \multimap y_S$, we have structural maps

 $\chi^-_{A\multimap S} : z_A \multimap z_S \quad \bar{\leadsto} \quad y_A \multimap y_S, \qquad \xi^+_{A\multimap S} : z_A \multimap z_S \quad \dot{\breve{\leadsto}} \quad x_A \multimap x_S.$

First, as $\chi_{A \to S}^-$ preserves labels, it decomposes uniquely into $\chi_A \to \chi_S$ for $\chi_A : z_A \rightsquigarrow y_A$ and $\chi_S : z_S \rightsquigarrow y_S$ structural maps. Furthermore, the negativity of $\chi_{A \to S}^-$ ensures that χ_A is positive and χ_S is negative. From the same reasoning on $\xi_{A \to S}^+$, we get $\xi_A^- : z_A \xrightarrow{\sim} x_A$ and $\xi_S^+ : z_S \xrightarrow{\pm} x_S$, which may be directly assembled to witness $y_A \xleftarrow{\leftarrow} x_A$ and $x_S \xleftarrow{\leftarrow} y_S$. \Box

This precisely ensures that the bijection $s_{A,S}^{-\infty}$ extends to an isomorphism of preorders as in Figure 12 – note how the contravariant on the left hand side matches the dualization of the moves in A in the construction of $A \multimap S$. The other noteworthy case is the exponential:

Lemma 5.21. Consider S a strict mixed board, and $x = [x_i \mid i \in I], y = [y_j \mid j \in J] \in \mathfrak{R}(!S)$. Then, $x \xleftarrow{} y$ iff for all $i \in I$, there is $j \in J$ such that $x_i \xleftarrow{} y_j$.

Proof. By Lemma 5.18, it suffices to reason on representatives.

If. Working with representatives, we have $I, J \subseteq_f \mathbb{N}$. Assume that for all $i \in I$, there is $j \in J$ such that $x_i \stackrel{+}{\longleftrightarrow} y_j$. Fix a function $f: I \to J$ such that for all $i \in I$, we have $x_i \stackrel{+}{\longleftrightarrow} y_{f(i)}$. So, we have $\chi_i^+: z_i \stackrel{+}{\Leftrightarrow} x_i$ and $\chi_i^-: z_i \stackrel{-}{\Leftrightarrow} y_{f(i)}$. Form $z = \lfloor z_i \mid i \in I \rfloor$. Setting

$$\begin{array}{rcl} \chi^+ & : & \lfloor z_i \mid i \in I \rfloor & \to & \lfloor x_i \mid i \in I \rfloor \\ & & (i,s) & \mapsto & (i,\chi_i^+(s)) \,, \end{array}$$

it is a positive structural map. Critically, it leaves i unchanged, as required by --preserving and it reaches all minimal moves in x, as required by --total. Likewise, setting

$$\begin{array}{rcl} \chi^- & : & \lfloor z_i \mid i \in I \rfloor & \to & \lfloor y_j \mid j \in J \rfloor \\ & & (i,s) & \mapsto & (f(i), \chi_i^-(s)) \end{array}$$

it follows that χ^- is +-preserving and +-total, so a positive structural map. Altogether, we have constructed a witness z and structural maps ensuring that $x \leftrightarrow y$ as required.

Only if. Assuming $x \stackrel{+}{\leftrightarrow} y$, we have some $z = \lfloor z_k \mid k \in K \rfloor$ with maps

$$\chi^+ : \lfloor z_k \mid k \in K \rfloor \stackrel{*}{\leadsto} \lfloor x_i \mid i \in I \rfloor, \qquad \chi^- : \lfloor z_k \mid k \in K \rfloor \stackrel{*}{\leadsto} \lfloor y_j \mid j \in J \rfloor.$$

Now, any $(k, s) \in \lfloor z_k \mid k \in K \rfloor$ minimal must be negative and ind(k, s) = k, so that $\chi^+(k, s) = (k, s')$ by *--preserving*. In fact, for every $(k, s) \in \lfloor z_k \mid k \in K \rfloor$, there is

$$(k, s_0) \rightarrow _! S (k, s_1) \rightarrow _! S \dots \rightarrow _! S (k, s)$$

a (unique) sequence of justifiers with s_0 minimal. Since χ^+ is \rightarrow -preserving, we must have $(k, s'_0) = \chi^+(k, s_0) <_{!S} \chi^+(k, s) = (k', s')$, but this entails that k = k': χ^+ preserves the first component k; and also $K \subseteq I$. In fact, this is an equality since χ^+ is also +-total, so that $z = \lfloor z_i \mid i \in I \rfloor$. Additionally, it follows that χ^+ decomposes into the data, for each $i \in I$, of $\chi_i^+: z_i \stackrel{t}{\rightarrow} x_i$. Decomposing χ^- similarly, we get a function $f: I \to J$ and $\chi_i^-: z_i \stackrel{t}{\rightarrow} y_{f(i)}$. But altogether, assembling all the data we get $x_i \stackrel{t}{\longleftrightarrow} y_{f(i)}$ for all $i \in I$ as required. \Box

Thus, we have established that the construction of a preorder from cartesian morphisms is compatible with the constructions on objects involved in the relative Seely structure. 5.4.2. Colored case. In the colored case, we must first extend cartesian morphisms:

Definition 5.22. Consider A a mixed board, and $x, y \in \mathscr{C}_{\mathcal{C}}(A)$ configurations in colors.

A structural map $f : x \rightsquigarrow y$ is simply a structural map in the previous sense, preserving colors, *i.e.* $\lambda_y(f(a)) = \lambda_x(a)$ for all $a \in x$. It is **positive** (resp. **negative**) if it is positive (resp. **negative**) as a plain structural map. A **cartesian morphism** is defined as in Definition 5.10, with positive and negative structural maps preserving colors.

All the development of Section 5.3 adapts transparently in the presence of colors, which are preserved everywhere. The induced preorder is again compatible with (color-preserving) symmetries, so that it yields a preorder on positions in colors; altogether extending $\mathfrak{R}^{\mathcal{C}}(A)$ into a preordered set $\mathfrak{S}^{\mathcal{C}}(A)$. Finally, as before the bijections matching the relative Seely constructions in games and in the linear Scott model are obviously compatible with colors, yielding the preorder-isomorphisms of Figure 13.

6. Qualitative Collapse

We have built, from any mixed board A, a preorder $\mathfrak{S}(A)$ (or $\mathfrak{S}^{\mathcal{C}}(A)$ in the colored case) in a way compatible with the constructions involved in the relative Seely structure of **ScottL**. We shall now extend that to strategies, as usual focusing first on the uncolored case.

6.1. Introduction. The basic idea for our collapse to the linear Scott model is simple: we shall simply take the down-closure of the relational collapse of (4.3), *i.e.*, without colors:

$$\mathfrak{S}(\sigma) = [\mathfrak{R}(\sigma)]_{\mathfrak{S}(A)^{\mathrm{op}} \times \mathfrak{S}(B)} \in \mathbf{ScottL}[\mathfrak{S}(A), \mathfrak{S}(B)]$$

$$(6.1)$$

which means that $(\mathsf{x}_A, \mathsf{x}_B) \in \mathfrak{S}(\sigma)$ provided we can find $\mathsf{y}_A \in \mathfrak{S}(A), \mathsf{y}_B \in \mathfrak{S}(B)$ with $\mathsf{y}_A \xleftarrow{+-} \mathsf{x}_A$ and $\mathsf{x}_B \xleftarrow{+-} \mathsf{y}_B$ along with a witness $y^{\sigma} \in \mathscr{C}^+(\sigma)$ such that $y^{\sigma}_A \in \mathsf{y}_A$ and $y^{\sigma}_B \in \mathsf{y}_B$.

It is obvious that this is indeed a morphism in **ScottL**. We shall now make our first steps towards proving functoriality, introducing the main difficulties. First, we note:

Lemma 6.1. Consider A a mixed board. Then, $\mathfrak{S}(\mathbf{c}_A) = \mathrm{id}_{\mathfrak{S}(A)}$.

Proof. This is straightforward: the latter comprises all pairs (x_1, x_2) with $x_2 \stackrel{\leftarrow}{\longleftrightarrow} x_1$, while the former amounts to all (y_1, y_2) such that there exists $y \in \mathfrak{S}(A)$ with a specific representative $y \in y$ with $y_2 \stackrel{\leftarrow}{\longleftrightarrow} y$ and $y \stackrel{\leftarrow}{\longleftrightarrow} y_1$ – clearly, this concerns the same pairs.

We also easily have oplax functoriality:

Lemma 6.2. Consider A, B and C mixed boards, and $\sigma : A \vdash B$, $\tau : B \vdash C$. Then, $\mathfrak{S}(\tau \odot \sigma) \subseteq \mathfrak{S}(\tau) \circ \mathfrak{S}(\sigma)$.

Proof. Consider $(\mathsf{x}_A, \mathsf{x}_C) \in \mathfrak{S}(\tau \odot \sigma)$. This means there are $\mathsf{y}_A \stackrel{t}{\longleftrightarrow} \mathsf{x}_A$ and $\mathsf{x}_C \stackrel{t}{\longleftrightarrow} \mathsf{y}_C$ with $(\mathsf{y}_A, \mathsf{y}_C) \in \mathfrak{R}(\tau \odot \sigma)$, which we know is in $\mathfrak{R}(\tau) \circ \mathfrak{R}(\sigma)$. But $\mathfrak{R}(\sigma) \subseteq \mathfrak{S}(\sigma)$ and likewise for τ , hence $(\mathsf{y}_A, \mathsf{y}_C) \in \mathfrak{S}(\tau) \circ \mathfrak{S}(\tau)$. But then, $(\mathsf{x}_A, \mathsf{x}_C) \in \mathfrak{S}(\tau) \circ \mathfrak{S}(\tau)$ as well by down-closure. \Box

Again, the final inequality $\mathfrak{S}(\tau) \circ \mathfrak{S}(\sigma) \subseteq \mathfrak{S}(\tau \odot \sigma)$ is more problematic, let us see why. Consider $(\mathsf{y}_A, \mathsf{y}_C) \in \mathfrak{S}(\tau) \circ \mathfrak{S}(\sigma)$. Unfolding definitions, there must be witnesses $x^{\sigma} \in \mathscr{C}^+(\sigma)$ and $x^{\tau} \in \mathscr{C}^+(\tau)$ such that $x_A^{\sigma} \in \mathsf{x}_A \stackrel{+-}{\longleftrightarrow} \mathsf{y}_A$, $x_B^{\tau} \stackrel{+-}{\longleftrightarrow} x_B^{\sigma}$, and $x_C^{\tau} \in \mathsf{x}_C$ with $\mathsf{y}_C \stackrel{+-}{\longleftrightarrow} \mathsf{x}_C$. Recall that the situation in Section 4.2.2, concerning lax preservation of composition for the relational collapse, was somewhat analogous, except that we had $x_B^{\sigma} \cong_B x_B^{\tau}$ rather than $x_B^{\sigma} \stackrel{-+}{\longleftrightarrow} x_B^{\tau}$ – we needed to synchronize x^{σ} and x^{τ} through a symmetry, whereas we now



Figure 14: Example resolution of a cartesian matching problem

want to synchronize them through a *cartesian morphism*. In the symmetry case, this was handled by Proposition 3.13, which we do not (yet!) have for cartesian morphisms.

Let us sum up below what we need, for fixed $\sigma : A \vdash B$ and $\tau : B \vdash C$. What we have: witnesses $x^{\sigma} \in \mathscr{C}^+(\sigma)$ and $x^{\tau} \in \mathscr{C}^+(\tau)$ along with a cartesian morphism $\chi : x_B^{\sigma} \xleftarrow{\to} x_B^{\tau}$. What we want: some $y^{\tau} \odot y^{\sigma} \in \mathscr{C}^+(\tau \odot \sigma)$ such that $y_A^{\sigma} \xleftarrow{\to} x_A^{\sigma}$ and $x_C^{\tau} \xleftarrow{\to} y_C^{\tau}$. In the sequel, we shall refer to this data as, respectively, a *cartesian (matching) problem*, and a *solution* to the problem. We call this a cartesian problem because both strategies are actively trying to duplicate and erase each other, and a solution is a situation where both strategies have reached a state where they have all the resources they need, not more and not less.

Example 6.3. To illustrate this operation, we detail an example, using two λ -terms

$$\vdash M_{\sigma} = \lambda f x. \underbrace{f(\dots(f)_{n} x) \dots}_{n} : (o \to o) \to o \to o,$$
$$g: (o \to o) \to o \to o \vdash M_{\tau} = \lambda x. \underbrace{g(\dots(g)_{m} x) \dots}_{m} : (o \to o) \to o \to o.$$

The reader may recognize M_{σ} the Church integer for n and M_{τ} that for m, though on different types. Interpreting those (with the adequate promotion for M_{σ}), we obtain

 $\sigma \in \mathbf{DSInn}[1, B], \quad \tau \in \mathbf{DSInn}[B, C]$

strategies that we wish to compose – here, $B = ! \llbracket (o \to o) \to o \to o \rrbracket$ and $C = \llbracket (o \to o) \to o \to o \rrbracket$. As events in B in C correspond to atoms in those types, we find it convenient to write B as $(a \to b) \to c \to d$ and C as $(e \to f) \to g \to h$ to ease the correspondence.

With these notations, the upper part of Figure 14 presents a cartesian problem involving σ and τ . On the upper left part, we have the typical (unduplicated) configuration of σ , iterating *n* times its argument, displayed onto *B* as the configuration indicated (omitting copy indices). Likewise, on the upper right corner, we have the typical (unduplicated) configuration of τ , made larger than for σ because of η -expansion, displayed to *B* as shown. There is a cartesian morphism as shown, linking all pairs of moves with the same label.

Resolving this problem involves performing all the necessary duplications: τ makes m copies of σ , but the first copy of σ makes n copies of the m-1 remaining calls of σ , and so on... The solution appears in the bottom part of the diagram, consisting in the indicated expansions of the configurations of σ and τ , whose display on B now match.

The example above illustrates that solving a cartesian problem can involve an exponential blowup in the size of the configurations – indeed, we know that the Church integer for mapplied to that for n normalizes to the Church integer for n^m , witnessing the n^m calls to the event f^+ in the duplicated version of the configuration for τ in the diagram. In general, the situation is far worse: the size of the solution is not elementary in the size of the problem, witnessing the usual bounds in the normalization of the simply-typed λ -calculus¹³.

From this explosion, it is clear that the resolution of a cartesian problem will be non-trivial. In particular, we rely on a non-trivial termination argument, introduced next.

6.2. Bounding Interactions. Here we provide an upper bound on the size of solutions to cartesian problems – this relies on our earlier work on the size of interactions in Hyland-Ong games [Cla11, Cla13, Cla15], which we shall import into the realm of thin concurrent games.

6.2.1. Structural maps in strategies. In this endeavour, our first step will be to formalize what it means for one configuration to be an *expansion* of another, as in the left and right hand sides of Figure 14. This involves adapting *structural morphisms* to strategies:

Definition 6.4. Consider A, B mixed boards, $\sigma : A \vdash B$ a strategy, $x, y \in \mathscr{C}(\sigma)$. A **partial structural map** is a function $f : x^{\sigma} \to y^{\sigma}$ satisfying:

$$\begin{array}{ll} \textit{min-preserving:} & \text{for all } s \in \min(x^{\sigma}), \text{ then } f \, s \in \min(y^{\sigma}), \\ \rightarrow \textit{-preserving:} & \text{for all } s \twoheadrightarrow_{\sigma} t, \text{ then } f \, s \twoheadrightarrow_{\sigma} f \, t, \\ & \textit{valid:} & \partial_{\sigma} f \text{ has form } f_A \vdash f_B \text{ for } f_A : x_A^{\sigma} \rightsquigarrow y_A^{\sigma} \text{ and } f_B : x_B^{\sigma} \rightsquigarrow y_B^{\sigma}. \end{array}$$

where $\partial_{\sigma} f$, the **display** of f to $A \vdash B$, is obtained as the composition

 $x^{\sigma}_{A} \vdash x^{\sigma}_{B} \quad \stackrel{\partial^{-1}_{\sigma}}{\simeq} \quad x^{\sigma} \quad \stackrel{f}{\to} \quad y^{\sigma} \quad \stackrel{\partial_{\sigma}}{\simeq} \quad y^{\sigma}_{A} \vdash y^{\sigma}_{B} \, .$

We write $f: x^{\sigma} \rightsquigarrow y^{\sigma}$. It is a **structural map**, written $f: x^{\sigma} \rightsquigarrow y^{\sigma}$, if additionally *total:* if $f \mathrel{s} \twoheadrightarrow_{\sigma} t^{+}$, there is a (necessarily) unique $\mathrel{s} \twoheadrightarrow_{\sigma} u^{+}$ s.t. $f(u^{+}) = t^{+}$.

¹³It is direct to extend Figure 14 into an example of that, exploiting the explosion of the application $n \dots n$ of Church integers – here τ is n, and σ is the tuple of ns typed appropriately.

So in essence, a structural map between configurations of a strategy is a forest-morphism which displays to a structural map in the game, in the sense of Definition 5.7.

Morally, we wish that a structural map can only send an event of σ to one that is "the same", in the sense of the symmetry. For deterministic sequential innocent strategies this follows from the definition above (which is indeed fine-tuned for deterministic sequential innocent strategies, but would be poorly behaved beyond that) – we now introduce a lemma that expresses that. For that, we recall from [Cla24] the concept of a **grounded causal chain (gcc)** of a strategy σ : it is a finite set $\rho = \{\rho_1, \ldots, \rho_n\} \subseteq_f |\sigma|$ that forms

$$\rho_1 \twoheadrightarrow_{\sigma} \ldots \twoheadrightarrow_{\sigma} \rho_n$$

a sequence with $\rho_1 \in \min(\sigma)$. In the general framework of [Cla24], a gcc may not be a configuration as it may not be down-closed. But here, because strategies are sequential innocent, we have $\rho \in \mathscr{C}(\sigma)$ indeed – it is then simply a *branch* of σ . We write $gcc(\sigma)$ for the set of gccs of σ . Additionally, if $x^{\sigma} \in \mathscr{C}(\sigma)$, we write $gcc(x^{\sigma})$ for the set of gccs within x^{σ} .

Lemma 6.5. Consider A, B mixed boards, $\sigma : A \vdash B$ a strategy, $x^{\sigma}, y^{\sigma} \in \mathscr{C}(\sigma)$, and $f : x^{\sigma} \xrightarrow{\mathscr{R}} y^{\sigma}$ a partial structural map.

Then, for any $\rho \in \operatorname{gcc}(x^{\sigma})$, f induces a symmetry $f_{\uparrow \rho} : \rho \cong_{\sigma} f \rho$.

Proof. By induction on ρ . If ρ is empty, then this is obvious.

Consider first $\rho \to s_1^- \to s_2^+ \in \operatorname{gcc}(x^{\sigma})$. By induction hypothesis, we have

$$f:\rho \twoheadrightarrow s_1^- \cong_{\sigma} f \, \rho \twoheadrightarrow f \, s_1^-$$

a symmetry, which extends on the left to $\rho \rightarrow s_1^- \rightarrow s_2^+$, a configuration since σ is deterministic sequential innocent. Now by symmetry, there must be an extension

$$f \cup \{(s_2^+, t_2^+)\} : \rho \twoheadrightarrow s_1^- \twoheadrightarrow s_2^+ \cong_{\sigma} f \rho \twoheadrightarrow f s_1^- \twoheadrightarrow t_2^+$$

with in particular $f s_1^- \to t_2^+$. But now we also have $f s_1^- \to f s_2^+$ since f is rigid; and this implies $f s_2^+ = t_2^+$ since σ is dsinn. So overall, we have, as required:

$$f: \rho \twoheadrightarrow s_1^- \twoheadrightarrow s_2^+ \cong_{\sigma} f \rho \twoheadrightarrow f s_1^- \twoheadrightarrow f s_2^+.$$

Consider now $\rho \to s_1^+ \to s_2^- \in \mathsf{gcc}(x^{\sigma})$. By induction hypothesis,

$$f:\rho \twoheadrightarrow s_1^+ \cong_{\sigma} f \rho \twoheadrightarrow f s_1^+$$

is a symmetry. But also, $f s_1^+ \to f s_2^-$ since f preserves \to . But now, we argue that

$$\partial_{\sigma} f : \partial_{\sigma} (\rho \twoheadrightarrow s_1^+) \twoheadrightarrow \partial_{\sigma} s_2^- \cong_{A \vdash B} \partial_{\sigma} (\rho \twoheadrightarrow s_1^+) \twoheadrightarrow \partial_{\sigma} (f s_2^-)$$

by transparent, because as a structural map, $\partial_{\sigma} f$ preserves labels. It follows that

$$f:\rho \twoheadrightarrow s_1^+ \twoheadrightarrow s_2^+ \cong_{\sigma} f \rho \twoheadrightarrow f s_1^+ \twoheadrightarrow f s_2^-$$

by \sim -receptive, as required.

The last case is the same where s_2^- is minimal; which is similar but simpler.

6.2.2. Pointer structures. In [CH10], the present author together with Harmer studied the termination of the simply-typed λ -calculus, through the lense of game semantics. In particular, they proved that any interaction between finite innocent strategies (in the traditional sense of Hyland-Ong games [HO00]) must be finite. This was later refined by Clairambault into a quantitative bound [Cla11, Cla13, Cla15], that we shall import here into concurrent games. As we rely on this result, we must first provide a reminder.

The result is more adequately phrased in terms of *pointer structures*:

Definition 6.6. A **pointer structure** is the data of a natural number $n \in \mathbb{N}$ together with

$$\phi: \{1, \dots, n-1\} \to \{0, \dots, n-2\}$$

a **pointer** function, which is:

contractive: for all
$$i \in \{1, ..., n-1\}$$
, $\phi(i) < i$,
alternating: for all $i \in \{1, ..., n-1\}$, *i* is even iff $\phi(i)$ is odd.

Pointer structures are what remain from Hyland-Ong games by forgetting the identity of moves in arenas, and only remembering the *pointers*. We give an example below, pictured from left to right. Instead of writing integers, we only write \circ for even numbers (reminiscent of Opponent moves) and \bullet for odd numbers (reminiscent of Player moves),



and the function **ptr** is indicated by following the edges from right to left.

Not all these pointer structures can arise in interactions between strategies; only those that satisfy an additional *visibility* condition. Given a pointer structure ϕ , we define

$$\begin{bmatrix} - & - \\$$

respectively called the *P*-view and the *O*-view. The *P*-view captures the part of a play available to an innocent strategy, while the *O*-view captures the part of a play available to an innocent environment. A pointer structure is visible when for all $1 \le i \le n-1$, we have $\phi(i) \in [i]$ when *i* is odd and $\phi(i) \in [i]$ when *i* is even – visible pointer structures are exactly those that may arise as an interaction between innocent strategies.

Finally, in a pointer structure $\phi : \{1, \ldots, n-1\} \to \{0, \ldots, n-2\}$, the **length** $|\phi|$ of ϕ is simply n. The **depth** of $1 \le i \le n-1$ is the $k \in \mathbb{N}$ such that $\phi^k(i) = 0$; we extend this to $0 \le i \le n-1$ by stating that the depth of 0 is 0. The **depth** of a pointer structure ϕ is the maximal depth of $0 \le i \le n-1$. The *P*-size of ϕ is the minimal N such that for all $0 \le i \le n-1$, we have $\#[i] \le 2N$, the *O*-size of ϕ is the minimal N such that for all $0 \le i \le n-1$, we have $\#[i] \le 2N+1$ – the depth roughly corresponds to the order of the type, the *P*-size and *O*-size to the size of the two interacting strategies.

Then, we have [Cla15, Theorem 4.17]:

Theorem 6.7. Consider ϕ visible pointer structure of depth bounded by $d \ge 3$, *P*-size bounded by $n \ge 1$ a *O*-size bounded by $p \ge 1$. Then, writing $2_0(N) = N$ and $2_{d+1}(N) = 2^{2_d(N)}$,

$$|\phi| \leq 2_{d-3} \left(\frac{p^{n+1}-1}{p-1} - 1 \right) ,$$

additionally $|\phi| = 1$ if d = 0, $|\phi| = 2$ if d = 1, and $|\phi| \le 2n$ if d = 2.

In [Cla15], this bound is also shown to be asymptotically tight. Though for the work here, this precise bound does not matter: what matters is that it exists.

6.2.3. The upper bound. We now move back to the technical setting of this paper, for interacting $\sigma : A \vdash B$, $\tau : B \vdash C$ with $x^{\sigma} \in \mathscr{C}(\sigma)$, $x^{\tau} \in \mathscr{C}(\tau)$ and a cartesian morphism $\chi : x_B^{\sigma} \xleftarrow{\tau^+} x_B^{\tau}$, we want an upper bound to the size of all solutions to the cartesian problem.

We call the τ -size of $(x^{\sigma}, \chi, x^{\tau})$ the minimal n s.t. every gcc of $x_A^{\sigma} \parallel x^{\tau}$ is smaller than 2n; its σ -size the minimal p such that every gcc of $x^{\sigma} \parallel x_C^{\tau}$ is smaller than 2p; its **depth** the minimal d + 2 such that every gcc of x_C^{τ} is smaller than d + 2, every gcc of x_B^{σ}, x_B^{τ} is smaller than d + 1, and every gcc of x_A^{σ} is smaller than d. Finally, its **branching degree** is the minimal b such that (regarded as trees), x^{σ} and x^{τ} have branching degree smaller than b.

Lemma 6.8. Consider $(x^{\sigma} \in \mathscr{C}(\sigma), \chi, x^{\tau} \in \mathscr{C}(\tau))$ as above with τ -size less than $n \geq 1$, σ -size less than $p \geq 1$, depth less than $d \geq 3$ and branching degree less than $b \geq 2$.

Then, for any $y^{\sigma} \in \mathscr{C}(\sigma)$ and $y^{\tau} \in \mathscr{C}(\tau)$ matching such that there are structural maps $\chi^{\sigma}: y^{\sigma} \xrightarrow{\mathfrak{p}} x^{\sigma}, \ \chi^{\tau}: y^{\tau} \xrightarrow{\mathfrak{p}} x^{\tau}$ with $\chi^{\sigma}_{A}: y^{\sigma}_{A} \xrightarrow{\mathfrak{p}} x^{\sigma}_{A}$ and $\chi^{\tau}_{C}: y^{\tau}_{C} \xrightarrow{\mathfrak{p}} x^{\tau}_{C}$, we have

$$\#(y^{\tau} \circledast y^{\sigma}) \leq b^{2_{d-3}\left(rac{p^{n+1}-1}{p-1}-1
ight)}.$$

Proof. Firstly, because σ and τ are deterministic sequential innocent, it follows that $y^{\tau} \circledast y^{\sigma}$ is a forest. Consider one of its branches $\rho \in gcc(y^{\tau} \circledast y^{\sigma})$, written as

$$\rho_0 \rightarrow \rho_1 \rightarrow \ldots \rightarrow \rho_{l-1}$$

for l its length. For each $1 \leq i \leq l-1$, if $\partial_{\tau \circledast \sigma} \rho_i$ is non-minimal in $A \parallel B \parallel C$, then its unique immediate predecessor is some $\partial_{\tau \circledast \sigma} \rho_j$ for j < i, we set $\phi(i) = j$. If $\partial_{\tau \circledast \sigma} \rho_i$ is minimal in A, we must have $(\rho_i)_{\sigma}$ defined. Because y^{σ} is a forest, there is a unique minimal $s \in y^{\sigma}$ such that $s <_{\sigma} (\rho_i)_{\sigma}$, and it must correspond to some unique ρ_j (such that $(\rho_j)_{\sigma} = s$), with j < i - we set $\phi(i) = j$. If $\partial_{\tau \circledast \sigma} \rho_i$ is minimal in B we set $\phi(i) = 0$, and it cannot be minimal in C because ρ_0 is, and y^{σ}_C (thus y^{σ}) has a unique minimal event since C is strict¹⁴.

Then, it follows by [Cla24, Proposition 10.2.5] that ϕ is a visible pointer structure; that its *O*-views correspond to gccs in $y^{\sigma} \parallel y_{C}^{\tau}$, and that its *P*-views correspond to gccs in $y_{A}^{\sigma} \parallel y^{\tau}$. But because χ^{σ} and χ^{τ} and their displays χ_{A}^{σ} , χ_{C}^{τ} are structural maps and hence forest morphisms, those are bounded respectively by 2p and 2n. Likewise, by construction its depth is smaller than d. Hence, we may apply Theorem 6.7 and deduce that we have

$$\#\rho \leq 2_{d-3} \left(\frac{p^{n+1} - 1}{p - 1} - 1 \right) . \tag{6.2}$$

So $y^{\tau} \circledast y^{\sigma}$ is a forest whose depth is bounded by this quantity. To deduce a bound on the size of $y^{\tau} \circledast y^{\sigma}$, we give a corresponding upper bound to the branching degree of that forest. For that, consider $p \in y^{\tau} \circledast y^{\sigma}$. We prove that the branching at p is bounded by b, reasoning by cases on the polarity of p and its component of occurrence.

If p has polarity I and occurs in A, then reasoning by cases via [Cla24, Lemma 6.2.15], any $p \rightarrow_{\tau \circledast \sigma} p'$ must satisfy that p' has polarity -, p'_{σ} is defined and

$$\partial_{\tau \circledast \sigma}(p_{\sigma}) \twoheadrightarrow_{A \parallel B \parallel C} \partial_{\tau \circledast \sigma}(p'_{\sigma})$$

so that in particular, $p_{\sigma} \rightarrow_{\sigma} p'_{\sigma}$ by [Cla24, Lemma 6.1.16]. Since χ^{σ} is rigid,

$$\chi(p_{\sigma}) \twoheadrightarrow_{\sigma} \chi(p'_{\sigma})$$

¹⁴This is the definition of the *justifier* in an interaction between visible strategies [Cla24, Section 10.2.2].

as well. Hence, this defines a map from successors of p in $y^{\tau} \circledast y^{\sigma}$ to successors of $\chi(p_{\sigma})$ in x^{σ} . We show this map is injective; for that, consider $p \to_{\tau \circledast \sigma} p', p''$ with $\chi(p'_{\sigma}) = \chi(p''_{\sigma})$. In particular, $\partial_{\sigma}(\chi(p'_{\sigma})) = \partial_{\sigma}(\chi(p''_{\sigma}))$ so they have the same label, index and predecessor. By *rigidity*, *label-preserving* and *negative*, it follows that $\partial_{\sigma}(p'_{\sigma})$ and $\partial_{\sigma}(p''_{\sigma})$ have the same predecessor, label and copy index, so that $\partial_{\sigma}(p'_{\sigma}) = \partial_{\sigma}(p''_{\sigma})$ by *jointly injective* [Cla24, Definition 12.1.1]. By *receptivity* of σ , it follows that $p'_{\sigma} = p''_{\sigma}$, so that p' = p'' also by *local injectivity* of ∂_{σ} . It follows that the set of causal successors of p in $y^{\tau} \circledast y^{\sigma}$ has cardinal less or equal than B. The case where p has polarity r and occurs in C is symmetric.

If p has polarity -, say e.g. that it occurs in C. Then, reasoning by cases based on [Cla24, Lemma 6.2.15], any $p \to_{\tau \circledast \sigma} p'$ in $y^{\tau} \circledast y^{\sigma}$ must have polarity r and satisfy $p_{\tau} \to_{\tau} p'_{\tau}$. By sequential innocence, there can be at most one such p'_{τ} , so there is at most one such p' in $\tau \circledast \sigma$. If p has polarity I and occurs in B, then again reasoning by cases on [Cla24, Lemma 6.2.15], any $p \to_{\tau \circledast \sigma} p'$ in $y^{\tau} \circledast y^{\sigma}$ satisfies $p_{\tau} \to_{\tau} p'_{\tau}$, and the same reasoning applies. The final case where p has polarity r and occurs in B is symmetric.

Altogether, $y^{\tau} \circledast y^{\sigma}$ is a forest of depth bounded by (6.2) and branching bounded by b, giving the announced upper bound on its size.

This upper bound could be improved for d = 1, 2, but again the precise quantity does not really matter here; only that once x^{σ} and x^{τ} are fixed, the solutions to the cartesian problems have a bounded size. Note also the importance of the assumptions that χ^{σ}_A is negative and χ^{τ}_C positive: this is what ensures that y^{σ} and y^{τ} do not have more duplications by the external Opponent than in x^{σ} and x^{τ} , meaning that the upper bound to the branching degree of x^{σ}, x^{τ} transports to y^{σ}, y^{τ} and to $y^{\tau} \circledast y^{\sigma}$. Finally, we note that as $\#y^{\sigma} \leq \#(y^{\tau} \circledast y^{\sigma})$ and likewise for τ , the same upper bound applies to y^{σ} and y^{τ} .

6.3. Functorial Collapse. With this, we have finished introducing the essential ingredients to show functoriality of the collapse to the linear Scott model.

6.3.1. Cartesian problems and how to solve them. First, the tools introduced just above now allow us to be more precise about what we mean by *cartesian problem*.

Definition 6.9. Consider $\sigma : A \vdash B$ and $\tau : B \vdash C$. A **cartesian (matching) problem** is the data of $x^{\sigma} \in \mathscr{C}(\sigma), x^{\tau} \in \mathscr{C}(\tau)$ and a cartesian morphism $\chi : x_B^{\sigma} \xleftarrow{\tau} x_B^{\tau}$.

A solution to this problem is given by $y^{\sigma} \in \mathscr{C}(\sigma), y^{\tau} \in \mathscr{C}(\tau), \chi^{\sigma}, \chi^{\tau}$ such that:



with $\chi_A^{\sigma}: y_A^{\sigma} \xrightarrow{\sim} x_A^{\sigma}$ and $\chi_C^{\tau}: y_C^{\tau} \xrightarrow{\leftarrow} x_C^{\tau}$.

This captures the notion of cartesian matching problem introduced in Section 6.1. First, indeed we have $y_A^{\sigma} \stackrel{+}{\longrightarrow} x_A^{\sigma}$ and $x_C^{\tau} \stackrel{+}{\longleftrightarrow} y_C^{\tau}$, so that this solution will provide the required witness for functoriality. But we have a little bit more here: we know that y^{σ} is an *expansion* of x^{σ} and likewise y^{τ} is an expansion of x^{τ} ; this is witnessed by specific structural morphisms χ^{σ} and χ^{τ} whose display is compatible with the cartesian morphism χ .

Now, the next key proposition shows how to solve cartesian problems:

Proposition 6.10. Consider $\sigma : A \vdash B$ and $\tau : B \vdash C$ deterministic sequential innocent. Then any cartesian problem for σ, τ has a unique solution.

Proof. Consider $x^{\sigma} \in \mathscr{C}(\sigma), x^{\tau} \in \mathscr{C}(\tau)$ and $\chi : x_B^{\sigma} \stackrel{\text{at}}{\longleftrightarrow} x_B^{\tau}$ a cartesian problem.

First, note we can find a *partial solution*, *i.e.* $y^{\sigma} \in \mathcal{C}(\sigma), y^{\tau} \in \mathcal{C}(\tau), \chi^{\sigma}$ and χ^{τ} s.t.:



with $\chi_A^{\sigma} : y_A^{\sigma} \stackrel{\sim p}{\sim} x_A^{\sigma}$ and $\chi_C^{\tau} : y_C^{\tau} \stackrel{\pm p}{\sim} x_C^{\tau}$ – indeed one can take $y^{\sigma} = y^{\tau} = \emptyset$. Such partial solutions are partially ordered by componentwise inclusion. By Lemma 6.8, there is a bound $N \in \mathbb{N}$ on the cardinal of $y^{\tau} \circledast y^{\sigma}$ for partial solutions; thus there is a partial solution of maximal size. From now on, we fix a partial solution $y^{\sigma}, y^{\tau}, \chi^{\sigma}, \chi^{\tau}$ of maximal size and prove that it actually is a total solution as seeked.

First, we prove that $\chi^{\tau} : y^{\tau} \stackrel{*}{\to} x^{\tau}$ is *total*. Consider $\chi^{\tau} s \rightarrow_{\tau} t^{+}$. Assuming there is no $s \rightarrow_{\tau} u^{+}$ in y^{τ} such that $\chi^{\tau} u^{+} = t^{+}$, we shall construct an extension of the solution, contradicting its maximality. We start with the easy case where t^{+} occurs in *C*. Since τ is a forest, $[s]_{\tau}$ is a gcc. By Lemma 6.5, this means that the restriction of χ^{τ} to $[s]_{\tau}$ is $\theta : [s]_{\tau} \cong_{\tau} [\chi^{\tau} s]_{\tau}$ a symmetry on τ . Hence, there is $s \rightarrow_{\tau} u^{+}$ in τ such that $\theta \cup \{(u^{+}, t^{+})\} \in \mathscr{S}(\tau)$ is still a symmetry. Extend y^{τ} with u^{+} and $\chi^{\tau}(u^{+}) = t^{+}$, it is a direct verification that this yields a solution, contradicting maximality.

Now, let us assume that t^+ occurs in B, *i.e.* $\partial_{\tau}(t^+) = (1, b')$. First, we update y^{τ} and χ^{τ}_{B} as above. This means that we have $\partial_{\tau}(u^+) = (1, b)$ with b negative in B, with $y_B \vdash_B b$, and $\chi^{\tau}_B(b) = b'$. By receptive, there is a unique $y^{\sigma} \vdash v^-$ such that $\partial_{\sigma}(v^-) = (2, b)$. We must now extend χ^{σ} accordingly. For that, assume w.l.o.g. that b is not minimal – the minimal case is similar but simpler. So, there is a (unique) $a \rightarrow_B b$. By *locally injective*, there is a unique $n^+ \in y^{\sigma}$ such that $\partial_{\sigma}(n^+) = (2, a)$. Write $m^+ = \chi^{\sigma}(n^+)$ with $\partial_{\sigma}(m^+) = (2, c)$ with $c \in x_B^{\sigma}$ – by definition, $c = \chi^{\sigma}_B(a)$. If we also write $d = \chi^{\tau}_B(a)$, then we have $(c, d) \in \chi_B$ by hypothesis, and by \rightarrow -preserving, $d \rightarrow_B b'$ negative. Thus, by Lemma 5.13, there is a unique $c \rightarrow_B e^-$ in x_B^{σ} such that $(e^-, b') \in \chi_B$. By *locally injective*, there is a unique $k^- \in x^{\sigma}$ such that $\partial_{\sigma}(k^-) = (2, e^-)$. Then, we may finally extend χ^{σ} with $v^- \mapsto e^-$, and verify the required conditions. Symmetrically, $\chi^{\sigma} : y^{\sigma} \stackrel{\star}{\to} x^{\sigma}$ is *total* as well.

Finally, it remains to prove that χ_A^{σ} and χ_C^{τ} are not partial; first we check that χ_C^{τ} is --total. So consider $a \in y_C^{\tau}$ and $\chi_C^{\tau} a \rightarrow_C c^-$ – the case where c^- is minimal is similar but simpler. If c^- has no predecessor for χ_C^{τ} , then Lemma 5.15 provides a unique extension of y_C^{τ} which yields an extension of y^{τ} by receptivity, contradicting maximality. Likewise, the +-totality of χ_A^{σ} follows from Lemma 5.16.

For *uniqueness*, assume we have two solutions, *i.e.*



and consider *e.g.* an event $p \in v^{\tau} \odot v^{\sigma}$ which is minimal such that it is not in $y^{\tau} \odot y^{\sigma}$. Now we reason by cases: if p is positive for σ , this is absurd by sequential innocence for σ and receptivity for τ . If it is positive for τ , the situation is symmetric. If p is negative in A, then it must be also in $y^{\tau} \odot y^{\sigma}$ because χ_A^{σ} is *--preserving* and *--total* and by receptivity of σ . If p is negative in C, then the situation is symmetric.

Thus as illustrated in Figure 14, one can always find solutions to cartesian problems: two deterministic sequential innocent strategies trying to duplicate and erase each other alongside a cartesian morphism can always find a successful resolution, and this resolution is unique if one adequately takes into account the copy indices.

6.3.2. *Functorial collapse*. The property above is the key conceptual contribution of this work. The functoriality of the collapse to the linear Scott model immediately follows.

Theorem 6.11. We have a \sim -functor

$\mathfrak{S}(-):\mathbf{DSInn}\to\mathbf{Scott}$

Proof. Preservation of identities was handled in Lemma 6.1 and oplax preservation of composition in Lemma 6.2. For lax preservation of composition, consider $(\mathsf{x}_A, \mathsf{x}_B) \in \mathfrak{S}(\sigma)$ and $(\mathsf{x}_B, \mathsf{x}_C) \in \mathfrak{S}(\tau)$. By definition, this means that taking some representatives $x_A \in \mathsf{x}_A, x_B \in \mathsf{x}_B$ and $x_C \in \mathsf{x}_C$, there are witnesses $x^{\sigma} \in \mathscr{C}^+(\sigma)$ and $x^{\tau} \in \mathscr{C}^+(\tau)$ such that

$$x_A^{\sigma} \stackrel{+}{\longleftrightarrow} x_A, \qquad x_B \stackrel{+}{\longleftrightarrow} x_B^{\sigma}, \qquad x_B^{\tau} \stackrel{+}{\longleftrightarrow} x_B, \qquad x_C \stackrel{+}{\longleftrightarrow} x_C^{\tau},$$

meaning in particular that $x_B^{\tau} \stackrel{+}{\longleftrightarrow} x_B^{\sigma}$, which means that there is a cartesian morphism $\chi : x_B^{\sigma} \stackrel{-}{\longleftrightarrow} x_B^{\tau}$. This forms a cartesian problem, which we can solve using Proposition 6.10: the solution consists in $y^{\sigma} \in \mathscr{C}(\sigma), y^{\tau} \in \mathscr{C}(\tau)$ matching, with $\chi^{\sigma}, \chi^{\tau}$ such that:



with $\chi_A^{\sigma}: y_A^{\sigma} \xrightarrow{\sim} x_A^{\sigma}$ and $\chi_C^{\tau}: y_C^{\tau} \xrightarrow{+} x_C^{\tau}$. But then, we may form $y^{\tau} \odot y^{\sigma} \in \mathscr{C}^+(\tau \odot \sigma)$, with

$$(y^{\tau} \odot y^{\sigma})_A \xrightarrow{\sim} x_A^{\sigma} \xleftarrow{+-} x_A, \qquad x_C \xleftarrow{+-} x_C^{\tau} \xleftarrow{+} y_C^{\tau} = (y^{\tau} \odot y^{\sigma})_C$$

yielding a witness for $(x_A, x_C) \in \mathfrak{S}(\tau \odot \sigma)$ as required.

6.4. Further Structure. The above takes care of the main obstactle in linking thin concurrent games with the linear Scott model, but there is still some work to do to conclude.

6.4.1. A relative Seely functor. First, the functor of Theorem 6.11 extends to a relative Seely functor. To show this we must adjoin to it a number of structural isomorphisms, which are simply (the down-closure of) those of Figure 12. All required conditions are straightforward, save for the preservation of promotion which we detail here.

Lemma 6.12. Consider S, A mixed boards with S strict, and $\sigma : !S \vdash A$.

Then, the following diagram commutes (with $t_S^!$ the down-closure of $s_S^!$):



Proof. First, as $(\mathfrak{S}(\sigma) \circ t_S^!)^{\dagger}$ and $\mathfrak{S}(\sigma^{\dagger})$ are down-closed, it is equivalent to compose them with the preorder-isomorphisms $s_S^!$ and $s_T^!$, which are bijections between multisets (quotiented lists) and symmetry classes of configurations of !T (quotiented families), for which we use similar notations, thus we will deal with these bijections implicitly in this proof.

Thus, take $([x_{A,l} | l \in L], [y_{B,k} | k \in K]) \in (\mathfrak{S}(\sigma) \circ s_S^!)^{\dagger}$. This means that for all $k \in K$, we have $([x_{A,l} | l \in L], y_{B,k}) \in \mathfrak{S}(\sigma)$. So there is $|x_{A,l} | l \in L|$ and $y_{B,k} \in y_{B,k}$ such that

$$y_{B,k} \stackrel{+-}{\longleftrightarrow} x_B^{\sigma,k}, \qquad \qquad \lfloor x_{A,j}^{\sigma,k} \mid j \in J_k \rfloor \stackrel{+-}{\longleftrightarrow} \lfloor x_{A,l} \mid l \in L \rfloor$$

for some $x^{\sigma,k} \in \mathscr{C}^+(\sigma)$ displaying to $\lfloor x_{A,j}^{\sigma,k} \mid j \in J_k \rfloor \vdash x_B^{\sigma,k}$. We may form $\lfloor x^{\sigma,k} \mid k \in K \rfloor \in \mathscr{C}^+(\sigma^{\dagger})$ displaying by definition to $\lfloor x_{A,j}^{\sigma,k} \mid \langle k, j \rangle \in \Sigma_{k \in K} J_k \rfloor \vdash \lfloor x_B^{\sigma,k} \mid k \in K \rfloor$ and

$$\lfloor y_{B,k} \mid k \in K \rfloor \stackrel{+-}{\longleftrightarrow} \lfloor x_B^{\sigma,k} \mid k \in K \rfloor, \quad \lfloor x_{A,j}^{\sigma,k} \mid \langle k,j \rangle \in \Sigma_{k \in K} J_k \rfloor \stackrel{+-}{\longleftrightarrow} \lfloor x_{A,l} \mid l \in L \rfloor$$

as needed. Reciprocally, take $([x_{A,l} \mid l \in L], [y_{B,k} \mid k \in K]) \in \mathfrak{S}(\sigma^{\dagger})$. By definition, there are

 $\lfloor x_{A,l} \mid l \in L \rfloor \in [\mathsf{x}_{A,l} \mid l \in L], \qquad \lfloor y_{B,k} \mid k \in K \rfloor \in [\mathsf{y}_{B,k} \mid k \in K]$

representatives witnessed by σ^{\dagger} , *i.e.* there is $\lfloor x^{\sigma,i} \mid i \in I \rfloor \in \mathscr{C}^+(\sigma^{\dagger})$, such that

$$\lfloor x_{A,j}^{\sigma,i} \mid j \in J_i \rfloor \stackrel{+-}{\longleftrightarrow} \lfloor x_{A,l} \mid l \in L \rfloor, \qquad \lfloor y_{B,k} \mid k \in K \rfloor \stackrel{+-}{\longleftrightarrow} \lfloor x_B^{\sigma,i} \mid i \in I \rfloor$$

where $\partial_{\sigma^{\dagger}} \lfloor x^{\sigma,i} \mid i \in I \rfloor = \lfloor x_{A,j}^{\sigma,i} \mid j \in J_i \rfloor \vdash \lfloor x_B^{\sigma,i} \mid i \in I \rfloor$. But by Lemma 5.21 (on representatives), this means that for all $k \in K$ there is $i \in I$ such that $y_{B,k} \stackrel{+-}{\longleftrightarrow} x_B^{\sigma,i}$. Thus for all $k \in K$, there is $i \in I$ such that $x^{\sigma,i}$ witnesses that $([\mathbf{x}_{A,l} \mid l \in L], \mathbf{y}_{B,k}) \in \mathfrak{S}(\sigma)^{\dagger}$. \Box

Altogether, we have the following theorem:

Theorem 6.13. The above provide the components for a relative Seely functor:

$$\mathfrak{S}(-): \mathbf{DSInn} \to \mathbf{ScottL}$$

6.4.2. Collapse in colors. Finally, we must now extend this with colors – this essentially follows the pattern of Section 4.4. Most ingredients are in place already: the relational collapse of strategies in colors was introduced in Section 4.4.3, and in particular in (4.5). As without colors, this collapse is adapted to target the linear Scott model by taking the down-closure for the preorder with colors introduced in Section 5.4.2.

Thus, we simply restate the definition in (6.1) with colors:

$$\mathfrak{S}^{\mathcal{C}}(\sigma) = [\mathfrak{R}^{\mathcal{C}}(\sigma)]_{\mathfrak{S}^{\mathcal{C}}(A)^{\mathrm{op}} \times \mathfrak{S}^{\mathcal{C}}(B)} \in \mathbf{ScottL}[\mathfrak{S}^{\mathcal{C}}(A), \mathfrak{S}^{\mathcal{C}}(B)].$$
(6.3)

With this definition, we can finally state our main theorem:

Theorem 6.14. For every set C we have a relative Seely functor:

$$\mathfrak{S}^{\mathcal{C}}(-) \quad : \quad \mathbf{DSInn} \ o \ \mathbf{ScottL}$$
 .

Proof. Preservation of identity as in Lemma 6.1, via Lemma 4.8. For composition, as in Lemma 6.2, oplax functoriality follows from that for the colorful relational collapse (Proposition 4.9). Now for lax functoriality, consider $(\mathsf{x}_A,\mathsf{x}_C) \in \mathfrak{S}^{\mathcal{C}}(\tau) \circ \mathfrak{S}^{\mathcal{C}}(\sigma)$, so there is a position in colors x_B such that $(\mathsf{x}_A,\mathsf{x}_B) \in \mathfrak{S}^{\mathcal{C}}(\sigma)$ and $(\mathsf{x}_B,\mathsf{x}_C) \in \mathfrak{S}^{\mathcal{C}}(\tau)$. Unfolding the definitions, this means that there are experiments $(x^{\sigma},\lambda^{\sigma}) \in \mathscr{C}^+_{\mathcal{C}}(\sigma), (x^{\tau},\lambda^{\tau}) \in \mathscr{C}^+_{\mathcal{C}}(\tau)$ s.t.

$$(x_A^{\sigma}, \lambda_A^{\sigma}) \stackrel{+-}{\longleftrightarrow} (x_A, \lambda_A), \qquad (x_B^{\sigma}, \lambda_B^{\sigma}) \stackrel{-+}{\longleftrightarrow} (x_B^{\tau}, \lambda_B^{\tau}), \qquad (x_C, \lambda_C) \stackrel{+-}{\longleftrightarrow} (x_C^{\tau}, \lambda_C^{\tau})$$

providing us in particular with a cartesian morphism $\chi : x_B^{\sigma} \stackrel{\text{art}}{\longleftrightarrow} x_B^{\tau}$ which preserves colors, *i.e.* it relates moves with the same color. By Proposition 6.10, this cartesian problem has a solution, giving us $y^{\sigma} \in \mathscr{C}(\sigma), y^{\tau} \in \mathscr{C}(\tau), \chi^{\sigma}, \chi^{\tau}$ such that:



which we must enrich with colors. To turn y^{σ} and y^{τ} into experiments we simply set

$$\mu^{\sigma}(s) = \lambda^{\sigma}(\chi^{\sigma} s), \qquad \qquad \mu^{\tau}(t) = \lambda^{\tau}(\chi^{\tau} t),$$

as χ^{σ} and χ^{τ} are forest morphisms it is straightforward that those are valid experiments. Furthermore, μ^{σ} and μ^{τ} are matching because in the inclusion in the diagram above and since χ preserves colors. Hence, by Lemma 4.7, we may form an experiment

$$(y^{\tau} \odot y^{\sigma}, \mu^{\tau} \odot \mu^{\sigma}) \in \mathscr{C}^{+}_{\mathcal{C}}(\tau \odot \sigma),$$

additionally from the definition, χ^σ_A and χ^τ_C preserve colors, so we still have

$$(y_A^{\sigma}, \mu_A^{\sigma}) \xleftarrow{+-} (x_A, \lambda_A), \qquad (x_C, \lambda_C) \xleftarrow{+-} (y_C^{\tau}, \mu_C^{\tau}),$$

concluding lax functoriality and the fact that we have a functor. Finally, all structural isomorphisms required for a relative Seely functor are (the down-closure of) those in Figure 13. The proof for the coherence laws is essentially undisturbed by colors. \Box

Finally, from this we deduce:

Corollary 6.15. Consider $\Gamma \vdash M : A$ a simply-typed λ -term.

Then, the diagram commutes in **ScottL**, for each set C:



with t_{Γ}, t_A the down-closure of s_{Γ}, s_A :

We phrase the result with t_{Γ} , t_A because they are indeed the structural isomorphisms in **ScottL** given through the relative Seely functor. Note however that s_{Γ} and s_A are orderisomorphisms and $\mathfrak{S}^{\mathcal{C}}(\llbracket M \rrbracket_{\mathbf{DSInn}})$ and $\llbracket M \rrbracket_{\mathbf{ScottL}}^{\mathcal{C}}$ are down-closed, so that the diagram also holds if t_{Γ} , t_A are simply replaced with the order-isomorphisms s_{Γ} , s_A (this glosses over the fact that the functorial action of ! is different in **Rel** and **ScottL**. Fortunately, it is easily verified that if $\varphi : A \cong B$ is an order-isomorphism between preorders, then $!^{\mathbf{ScottL}}[\varphi]_{A^{\mathrm{op}} \times B} = [!^{\mathbf{Rel}}\varphi]_{(!A)^{\mathrm{op}} \times !B}$; here this ensures that $!^{\mathbf{ScottL}}t_{\Gamma} = [!^{\mathbf{Rel}}s_{\Gamma}]$). This concludes the link between thin concurrent games and the linear Scott model –

This concludes the link between thin concurrent games and the linear Scott model – this is really the main result of this paper. However, because we also know their relationship with the relational model, we are able to leverage it to study the direct link between the relational model and the linear Scott model, in the spirit of Ehrhard [Ehr12].

6.5. Quantitative Collapse. As announced, the interpretation of a simply-typed λ -term in ScottL₁ is simply the down-closure of its interpretation in Rel₁:

Theorem 6.16. Consider $\Gamma \vdash M : A$ a simply-typed λ -term. Then, $\llbracket M \rrbracket_{\mathbf{ScottL}_!} = [\llbracket M \rrbracket]_{\mathbf{Rel}_!}]$.

Proof. This is direct via the following calculations in **Rel**:

$$\begin{split} \llbracket M \rrbracket_{\mathbf{ScottL}_{!}} &= s_{A} \circ \mathfrak{S}^{\mathcal{C}}(\llbracket M \rrbracket_{\mathbf{DSInn}_{!}}) \circ ! s_{\Gamma}^{-1} \\ &= s_{A} \circ [\mathfrak{R}^{\mathcal{C}}(\llbracket M \rrbracket_{\mathbf{DSInn}_{!}})] \circ ! s_{\Gamma}^{-1} \\ &= [s_{A} \circ \mathfrak{R}^{\mathcal{C}}(\llbracket M \rrbracket_{\mathbf{DSInn}_{!}}) \circ ! s_{\Gamma}^{-1}] \\ &= [\llbracket M \rrbracket_{\mathbf{Rel}_{!}}] \end{split}$$

using Corollary 6.15 (inlining the remark about phrasing it with s_{Γ} , s_A); then by definition of the linear Scott collapse with colors in (6.3); then using that s_{Γ} , s_A are order-isomorphisms and hence commute with down-closure; and finally using Corollary 6.15.

In [Ehr12], Ehrhard shows that \mathbf{ScottL}_1 is the extensional collapse of \mathbf{Rel}_1 . Ehrhard presents this as a relatively complex categorical statement relating two cartesian closed categories \mathcal{C} (in this case, \mathbf{Rel}_1) and \mathcal{E} (in this case, \mathbf{ScottL}_1), the former intensional and the latter extensional; roughly speaking expressing that the interpretation of the λ -calculus in \mathcal{E} corresponds the interpretation in \mathcal{C} modulo extensional equivalence, a notion constructed by induction on types. We do not reprove this extensional collapse theorem here. In fact, Ehrhard's statement does not seem to follow easily from our results: in order to prove the extensional collapse, Ehrhard provides a cartesian closed category \mathbf{Ppl}_1 together with

$$\mathbf{Ppl}_{!} \rightarrow e(\mathbf{Rel}_{!}, \mathbf{ScottL}_{!})$$

a cartesian closed functor, where $(\mathcal{C}, \mathcal{E})$ is a category of "heterogeneous logical relations" glueing together, intuitively, an intensional model and an extensional model. There, $\mathbf{Ppl}_{!}$ is

a sort of hybrid between $\mathbf{Rel}_{!}$ and $\mathbf{ScottL}_{!}$: morphisms are relations that are well-behaved with respect to the preorder, a property phrased via biorthogonality. In particular, we have

$$\mathbf{Rel}_! \quad \leftarrow \quad \mathbf{Ppl}_! \quad o \quad \mathbf{ScottL}_!$$

cartesian closed functors, so that $\mathbf{Ppl}_{!}$ can play the same role as $\mathbf{DSInn}_{!}$ in the proof of Theorem 6.16. Thus we could imagine reproving Ehrhard's result by similarly providing

$$\mathbf{DSInn}_! \rightarrow e(\mathbf{Rel}_!, \mathbf{ScottL}_!)$$

a cartesian closed functor; however the natural definition does not seem to work. Strategies are somehow "too intensional"; the problem is that the fact that they behave well with respect to the preorder is derived combinatorially rather than maintained as an invariant – so they behave well only against other strategies, and not arbitrary elements in **Rel**.

This is not a bug! Rather than the extensional collapse statement *per se*, it is the relationship between qualitative and quantitative models implicit in Ehrhard's paper, *i.e.* Theorem 6.16, that turned out to be more influential, and that we reproduced here.

7. Conclusions

In addition to reproving Ehrhard's result by different means, this result is a stepping stone for other work in progress.

The first one is an infinitary extension of Theorem 6.14, *i.e.* an alternative collapse to the linear Scott model giving also account of infinitary executions, *i.e.* infinite configurations; we hope that such a result is key in a complete semantic understading of the decidability of higher-order model-checking [Ong06]. Ehrhard's does not extend to an infinitary version of Theorem 6.16 with respect to the infinitary relational model [GM15], and actually, it seems that surprisingly, the infinitary version of Theorem 6.16 does not hold at all in the presence of greatest fixed point. In contrast, work in progress suggests that Theorem 6.14 does extend.

The second one is a bicategorical extension: from the result of this paper, it is not too hard to send strategies to cartesian distributors [Oli21]. Proving that this is (pseudo-)functorial is more subtle; this could potentially give a bicategorical version of Ehrhard's result, thanks to the connection we already have between thin concurrent games and generalized species of structure [COP23].

Acknowledgment

The author would like to thank Charles Grellois for intense black board discussions which motivated this development in the first place.

This work was supported by the ANR project DyVerSe (ANR-19-CE48-0010-01); by the Labex MiLyon (ANR-10-LABX-0070) of Université de Lyon, within the program "Investissements d'Avenir" (ANR-11-IDEX-0007), operated by the French National Research Agency (ANR); and by the PEPR integrated project EPiQ ANR-22-PETQ-0007 part of Plan France 2030.

References

- [ACU10] Thorsten Altenkirch, James Chapman, and Tarmo Uustalu. Monads need not be endofunctors. In International Conference on Foundations of Software Science and Computational Structures, pages 297–311. Springer, 2010.
- [Aeh06] Klaus Aehlig. A finite semantics of simply-typed lambda terms for infinite runs of automata. In Zoltán Ésik, editor, Computer Science Logic, 20th International Workshop, CSL 2006, 15th Annual Conference of the EACSL, Szeged, Hungary, September 25-29, 2006, Proceedings, volume 4207 of Lecture Notes in Computer Science, pages 104–118. Springer, 2006. doi:10.1007/ 11874683_7.
- [AJM00] Samson Abramsky, Radha Jagadeesan, and Pasquale Malacaria. Full abstraction for PCF. Inf. Comput., 163(2):409–470, 2000.
- [BDER97] Patrick Baillot, Vincent Danos, Thomas Ehrhard, and Laurent Regnier. Timeless games. In Mogens Nielsen and Wolfgang Thomas, editors, Computer Science Logic, 11th International Workshop, CSL '97, Annual Conference of the EACSL, Aarhus, Denmark, August 23-29, 1997, Selected Papers, volume 1414 of Lecture Notes in Computer Science, pages 56–77. Springer, 1997. doi:10.1007/BFb0028007.
- [BM20] Davide Barbarossa and Giulio Manzonetto. Taylor subsumes scott, berry, kahn and plotkin. Proc. ACM Program. Lang., 4(POPL):1:1–1:23, 2020.
- [Bou09] Pierre Boudes. Thick subtrees, games and experiments. In Pierre-Louis Curien, editor, Typed Lambda Calculi and Applications, 9th International Conference, TLCA 2009, Brasilia, Brazil, July 1-3, 2009. Proceedings, volume 5608 of Lecture Notes in Computer Science, pages 65–79. Springer, 2009. doi:10.1007/978-3-642-02273-9_7.
- [CC21] Simon Castellan and Pierre Clairambault. Disentangling parallelism and interference in game semantics. *CoRR*, abs/2103.15453, 2021.
- [CCPW18] Simon Castellan, Pierre Clairambault, Hugo Paquet, and Glynn Winskel. The concurrent game semantics of probabilistic PCF. In *LICS*, pages 215–224. ACM, 2018.
- [CCRW17] Simon Castellan, Pierre Clairambault, Silvain Rideau, and Glynn Winskel. Games and strategies as event structures. Log. Methods Comput. Sci., 13(3), 2017.
- [CCW15] Simon Castellan, Pierre Clairambault, and Glynn Winskel. The parallel intensionally fully abstract games model of pcf. In 2015 30th Annual ACM/IEEE Symposium on Logic in Computer Science, pages 232–243. IEEE, 2015.
- [CCW19] Simon Castellan, Pierre Clairambault, and Glynn Winskel. Thin games with symmetry and concurrent hyland-ong games. Log. Methods Comput. Sci., 15(1), 2019.
- [CdV20] Pierre Clairambault and Marc de Visme. Full abstraction for the quantum lambda-calculus. Proc. ACM Program. Lang., 4(POPL):63:1–63:28, 2020. doi:10.1145/3371131.
- [CF23] Pierre Clairambault and Simon Forest. The cartesian closed bicategory of thin spans of groupoids. In 38th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2023, Boston, MA, USA, June 26-29, 2023, pages 1–13. IEEE, 2023. doi:10.1109/LICS56636.2023.10175754.
- [CH10] Pierre Clairambault and Russ Harmer. Totality in arena games. Ann. Pure Appl. Log., 161(5):673– 689, 2010.
- [Cla11] Pierre Clairambault. Estimation of the length of interactions in arena game semantics. In *FoSSaCS*, volume 6604 of *Lecture Notes in Computer Science*, pages 335–349. Springer, 2011.
- [Cla13] Pierre Clairambault. Bounding skeletons, locally scoped terms and exact bounds for linear head reduction. In TLCA, volume 7941 of Lecture Notes in Computer Science, pages 109–124. Springer, 2013.
- [Cla15] Pierre Clairambault. Bounding linear head reduction and visible interaction through skeletons. Log. Methods Comput. Sci., 11(2), 2015.
- [Cla24] Pierre Clairambault. Causal investigations in interactive semantics, 2024. Habilitation à Diriger les Recherches.
- [COP23] Pierre Clairambault, Federico Olimpieri, and Hugo Paquet. From thin concurrent games to generalized species of structures. In 38th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2023, Boston, MA, USA, June 26-29, 2023, pages 1–14. IEEE, 2023. doi: 10.1109/LICS56636.2023.10175681.
- [CP21] Pierre Clairambault and Hugo Paquet. The quantitative collapse of concurrent games with symmetry. *CoRR*, abs/2107.03155, 2021.

THE QUALITATIVE COLLAPSE OF CONCURRENT GAMES

- [dC18] Daniel de Carvalho. Execution time of λ -terms via denotational semantics and intersection types. Math. Struct. Comput. Sci., 28(7):1169–1203, 2018. doi:10.1017/S0960129516000396.
- [DE11] Vincent Danos and Thomas Ehrhard. Probabilistic coherence spaces as a model of higher-order probabilistic computation. *Inf. Comput.*, 209(6):966–991, 2011. doi:10.1016/j.ic.2011.02.001.
- [Ehr96] Thomas Ehrhard. Projecting sequential algorithms on strongly stable functions. Ann. Pure Appl. Log., 77(3):201–244, 1996. doi:10.1016/0168-0072(95)00026-7.
- [Ehr12] Thomas Ehrhard. The scott model of linear logic is the extensional collapse of its relational model. *Theor. Comput. Sci.*, 424:20–45, 2012.
- [EPT18] Thomas Ehrhard, Michele Pagani, and Christine Tasson. Full abstraction for probabilistic PCF. J. ACM, 65(4):23:1–23:44, 2018. doi:10.1145/3164540.
- [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.
- [FP09] Claudia Faggian and Mauro Piccolo. Partial orders, event structures and linear strategies. In TLCA, volume 5608 of Lecture Notes in Computer Science, pages 95–111. Springer, 2009.
- [Gir87] Jean-Yves Girard. Linear logic. Theor. Comput. Sci., 50:1–102, 1987. doi:10.1016/ 0304-3975(87)90045-4.
- [Gir88] Jean-Yves Girard. Normal functors, power series and λ-calculus. Ann. Pure Appl. Log., 37(2):129– 177, 1988. doi:10.1016/0168-0072(88)90025-5.
- [GM15] Charles Grellois and Paul-André Melliès. An infinitary model of linear logic. In Andrew M. Pitts, editor, Foundations of Software Science and Computation Structures - 18th International Conference, FoSSaCS 2015, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2015, London, UK, April 11-18, 2015. Proceedings, volume 9034 of Lecture Notes in Computer Science, pages 41–55. Springer, 2015. doi:10.1007/978-3-662-46678-0_3.
- [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. doi:10.1006/inco.2000.2917.
- [Hut93] Michael Huth. Linear domains and linear maps. In Stephen D. Brookes, Michael G. Main, Austin Melton, Michael W. Mislove, and David A. Schmidt, editors, Mathematical Foundations of Programming Semantics, 9th International Conference, New Orleans, LA, USA, April 7-10, 1993, Proceedings, volume 802 of Lecture Notes in Computer Science, pages 438–453. Springer, 1993. doi:10.1007/3-540-58027-1_21.
- [KO09] Naoki Kobayashi and C.-H. Luke Ong. A type system equivalent to the modal mu-calculus model checking of higher-order recursion schemes. In Proceedings of the 24th Annual IEEE Symposium on Logic in Computer Science, LICS 2009, 11-14 August 2009, Los Angeles, CA, USA, pages 179–188. IEEE Computer Society, 2009. doi:10.1109/LICS.2009.29.
- [LMMP13] Jim Laird, Giulio Manzonetto, Guy McCusker, and Michele Pagani. Weighted relational models of typed lambda-calculi. In 28th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2013, New Orleans, LA, USA, June 25-28, 2013, pages 301–310. IEEE Computer Society, 2013. doi:10.1109/LICS.2013.36.
- [Mel03] Paul-André Mellies. Asynchronous games 1: Uniformity by group invariance, 2003.
- [Mel05] Paul-André Melliès. Asynchronous games 4: A fully complete model of propositional linear logic. In LICS, pages 386–395. IEEE Computer Society, 2005.
- [Mel06] Paul-André Melliès. Asynchronous games 2: The true concurrency of innocence. *Theor. Comput.* Sci., 358(2-3):200–228, 2006.
- [Mel09] Paul-André Mellies. Categorical semantics of linear logic. *Panoramas et syntheses*, 27:15–215, 2009.
- [MM07] Paul-André Melliès and Samuel Mimram. Asynchronous games: Innocence without alternation. In CONCUR, volume 4703 of Lecture Notes in Computer Science, pages 395–411. Springer, 2007.
- [NPW79] Mogens Nielsen, Gordon D. Plotkin, and Glynn Winskel. Petri nets, event structures and domains. In Semantics of Concurrent Computation, volume 70 of Lecture Notes in Computer Science, pages 266–284. Springer, 1979.
- [Oli21] Federico Olimpieri. Intersection type distributors. In 36th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2021, Rome, Italy, June 29 - July 2, 2021, pages 1–15. IEEE, 2021. doi:10.1109/LICS52264.2021.9470617.

- [Ong06] C.-H. Luke Ong. On model-checking trees generated by higher-order recursion schemes. In 21th IEEE Symposium on Logic in Computer Science LICS 2006), 12-15 August 2006, Seattle, WA, USA, Proceedings, pages 81–90. IEEE Computer Society, 2006. doi:10.1109/LICS.2006.38.
- [Paq22] Hugo Paquet. Bi-invariance for uniform strategies on event structures. In *MFPS*, 2022.
- [Plo77] Gordon D. Plotkin. LCF considered as a programming language. Theor. Comput. Sci., 5(3):223–255, 1977. doi:10.1016/0304-3975(77)90044-5.
- [PSV14] Michele Pagani, Peter Selinger, and Benoît Valiron. Applying quantitative semantics to higherorder quantum computing. In Suresh Jagannathan and Peter Sewell, editors, The 41st Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '14, San Diego, CA, USA, January 20-21, 2014, pages 647–658. ACM, 2014. doi:10.1145/2535838. 2535879.
- [RW11] Silvain Rideau and Glynn Winskel. Concurrent strategies. In *LICS*, pages 409–418. IEEE Computer Society, 2011.
- [SS71] Dana S Scott and Christopher Strachey. Toward a mathematical semantics for computer languages, volume 1. Oxford University Computing Laboratory, Programming Research Group Oxford, 1971.
- [TO16] Takeshi Tsukada and C.-H. Luke Ong. Plays as resource terms via non-idempotent intersection types. In *LICS*, pages 237–246. ACM, 2016.
- [Win98] Glynn Winskel. A linear metalanguage for concurrency. In Armando Martin Haeberer, editor, Algebraic Methodology and Software Technology, 7th International Conference, AMAST '98, Amazonia, Brasil, January 4-8, 1999, Proceedings, volume 1548 of Lecture Notes in Computer Science, pages 42–58. Springer, 1998. doi:10.1007/3-540-49253-4_6.
- [Win07] Glynn Winskel. Event structures with symmetry. Electron. Notes Theor. Comput. Sci., 172:611– 652, 2007.