Fair Games and Full Completeness for Multiplicative Linear Logic without the MIX-Rule

(Working Draft)

J. M. E. Hyland[1]
Department of Pure Mathematics and Mathematical Statistics
16 Mill Lane, Cambridge CB2 1SB
jmeh1@phx.cam.ac.uk

and

C.-H. L. Ong[2]
Computer Laboratory, New Museums Site
Pembroke Street, Cambridge CB2 3QG
Luke.Ong@cl.cam.ac.uk

July 22, 1993

Abstract

We introduce a new category of finite, fair games, and winning strategies, and use it to provide a semantics for the multiplicative fragment of Linear Logic (MLL) in which formulae are interpreted as games, and proofs as winning strategies. This interpretation provides a categorical model of MLL which satisfies the property that every (history-free, uniformly) winning strategy is the denotation of a unique cut-free proof net. Abramsky and Jagadeesan first proved a result of this kind and they refer to this property as *full completeness*. Our result differs from theirs in one important aspect: the MIX-rule, which is not part of Girard's Linear Logic, is invalidated in our model. We achieve this sharper characterization by considering *fair* games. A finite, fair game is specified by the following data:

- moves which Player can play,
- moves which Opponent can play, and
- a collection of finite sequences of maximal (or terminal) positions of the game which are deemed to be fair.

Notably, positions of a game are a derived notion. The maximal positions of a compound game are obtained by an appropriate interleaving of the maximal positions of the respective constituent games. At any position in a finite, fair game, a player can make a move if, and only if, the move can be extended to a maximal position.

[1]Partially supported by CEC ESPRIT Project CLICS-II.
[2]On leave from the National University of Singapore and supported by a fellowship from Trinity College, Cambridge.

1 Introduction

Abramsky and Jagadeesan [AJ92] (henceforth abbreviated as AJ) have recently described a Game semantics for the multiplicative fragment of Linear Logic (MLL) in which formulae are interpreted by games, and proofs by winning strategies. Their work builds on an important work of Blass [Bla92] but improves on it in a number of ways. First, unlike Blass semantics which is in essence a Game semantics for Affine Logic, Weakening is invalidated in the AJ setting. Secondly, Blass's model actually only characterises classical propositional tautology, and it does not form a category, whereas the AJ semantics yields a categorical model of MLL. The most notable contribution of the work of AJ is the proof of what they call a Full Completeness Theorem which establishes an extremely tight correspondence between syntax of proofs (= proof nets) and semantics (= games and winning strategies). The force of the result is conveyed clearly in the Introduction of their paper, and from which we quote:

> The usual completeness theorems are stated with respect to provability; a full completeness theorem is with respect to proofs. This is best formulated in terms of a categorical model of the logic, in which formulae denote objects, and proofs denote morphisms. One is looking for a model $\mathbb{C}$ such that:
>
> - **Completeness**: $\mathbb{C}(A, B)$ is non-empty only if $A \vdash B$ is provable in the logic.
> - **Full Completeness**: Any $f : A \to B$ is the denotation of $A \vdash B$. (This amounts to asking that the unique functor form the relevant free category to $\mathbb{C}$ to be full, whence our terminology). One may even ask for there to be a *unique* cut-free such proof, *i.e.* that the above functor be faithful.

The conceptual innovation therein notwithstanding, the Full Completeness Theorem of AJ is still not the best possible proof-theoretic characterization of the multiplicative fragment of Linear Logic. The Game semantics of AJ does not characterise proofs of the *pure* MLL; rather it characterises proofs of the system MLL+MIX, which is the multiplicative fragment of Linear Logic extended with the following MIX-rule [Gir87]:

$$(\mathbf{MIX}) \qquad \frac{\vdash \Gamma \qquad \vdash \Delta}{\vdash \Gamma, \Delta}.$$

This leaves open the challenge of sharpening AJ's Full Completeness Theorem to fit the pure, multiplicative fragment of Linear Logic. What is needed to achieve this is a modified Game semantics which invalidates the MIX-rule. But how should the AJ game-theoretic setting be modified? and what needs to be done? We will now illustrate the key to our solution. (The discussion that follows is carried out in the framework of Blass, Abramsky and Jagadeesan, and we assume a modicum of familiarity with it. Readers who are not acquainted with their work should probably skip it, and resume at the next section. Nonetheless, this paper is intended to be a self-contained treatment leading up to our full completeness result.)

Motivation Consider Blass-style games which are two-person (called Player and Opponent) games played according to the following ground rules:

- Only the Opponent can start a play.
- Thereafter, the play alternates strictly between Player and Opponent: each player makes one move at a time.
- The outcome of each play is never ambiguous: no draw is possible.

- A player wins a play either when the other player is unable to respond with a move, or when the resultant infinite play belongs to a designated collection of infinite plays which he is deemed to win.

Using the Linear Logic par $\propto$ or the tensor $\otimes$ connectives, two games can be composed to form a new game to be thought of as two games played out "in parallel". Each play of the compound game is necessarily put together by interleaving two plays, one from each constituent game, always observing the *switching convention*:

> In a par game, only Player is allowed to switch game; whereas in a tensor game, only Opponent is allowed to switch game.

Take Γ to be the formula $A^{\perp} \propto A$, and Δ to be $B^{\perp} \propto B$, where A and B are thought of as games in the AJ sense. Applying the MIX-rule as stated above, the sequent $\vdash A^{\perp} \propto A, B^{\perp} \propto B$ is a theorem of the system MLL + MIX, which is equivalent to each of the two following sequents

(1) $\pi \quad \vdash (A \otimes B)^{\perp} \propto (A \propto B)$,

(2) $\pi : A \otimes B \vdash A \propto B$.

By AJ's Full Completeness Theorem, we know that there is a winning strategy π which may be understood as a rule that takes any winning strategy σ for the game $A \otimes B$, and transforms it into a winning strategy $\pi(\sigma)$ for the game $A \propto B$. (Alternatively, we may read π as a winning strategy for the game $(A \otimes B)^{\perp} \propto (A \propto B)$.) Further, the rule $\pi(\sigma)$ to be followed by Player is strikingly simple: it does no more than mimicking Opponent's move at every step. That this "copy-cat" strategy is the essence of defeating all counter strategies in a generic or uniform way is a key idea in Game semantics.

Let us now pit the transformed strategy $\pi(\sigma)$ for the game $A \propto B$ against a counter strategy τ. Suppose following τ, Opponent chooses to start a play by making a move, say a_1 in the game A. To work out how he should respond to a_1 according to the strategy $\pi(\sigma)$, Player pretends that he is playing a "shadow" game against Opponent in the game $A \otimes B$, and in the shadow play, Opponent is regarded to have made the move a_1. Player now applies the winning strategy σ to the shadow play, and let us say the prescribed response is a_2 which must remain a move in game A because of the switching convention for tensor game. Now, Player flits from the shadow play in the game $A \otimes B$ to the real play in the game $A \propto B$, and respond to the move a_1 with a_2. Suppose Opponent then makes the move a_3. Since the real game $A \propto B$ is a par game, the Opponent's move a_3 has to remain in game A according to the switching convention. To decide how to respond, Player flits back to the shadow play, and pretends that Opponent has just made the move a_3 in the game $A \otimes B$. Suppose σ then prescribes the move a_4 as a response. Note that a_4 has to be in the constituent game A according to the switching convention; and so, the process goes on.

What is clear is that whether this play terminates after finitely many moves or not, the resultant play in either the real or shadow game necessarily projects onto one of the two constituent games to the exclusion of the other; and this is the same game as the one in which Opponent chose to start the play (which is A in the above example). We may describe this phenomenon as a case of an *unfair* play, in the sense that no attention is paid to the game B throughout the play. The word fairness is deliberately reminiscent of the notion of fairness in Computer Science (as in, for example, fair merge).

The moral of the example is this: a Game semantics which invalidates the MIX-rule is one in which only fair (in the above sense) plays are allowed.

Outline of the Paper The rest of the paper is organised as follows. Section 2 gives a brief introduction to the basics of the multiplicative fragment of Linear Logic, the notion of proof nets, and Danos and Regniers' graph-theoretic characterization of the soundness condition for proof nets. Section 3 introduces the category of finite, fair games, and winning strategies, as well as the subcategory of history-free strategies. We then show how an arbitrary MLL sequent defines a functor over appropriate categories of games. The development in this section leads up to the definition of a categorical model of MLL. Section 3 concludes with a proof of the soundness part of the Full Completeness Theorem. Section 4 is devoted to the proof of the completeness part of the Theorem. Finally, comparisons with related work are carried out in Section 5. We conclude with an indication of how the framework of fair games can be extended to full Classical Linear Logic, thus setting out further directions.

We end this introductory section with a few remarks on the chronology of the work reported here. This work was inspired by the work of Abramsky and Jagadeesan, which was presented at the start of the Cambridge – Imperial Joint Seminar Series on Game Semantics in October 1992. The proof of our Full Completeness Theorem was completed in February, 1993, and was first announced at a meeting of the Joint Seminar Series at Imperial College, London. Details of the work were also presented as part of the first author's series of invited lectures in the Dutch Logic Year at Utrecht, in March 1993.

2 MLL Sequents and Proof Nets

The formulae of the multiplicative fragment of Linear Logic (MLL) are built up from propositional atoms $\alpha, \beta, \gamma, \cdots$ and their linear negation $\alpha^\perp, \beta^\perp, \gamma^\perp, \cdots$ by the binary connectives tensor $\otimes$, and par $\propto$. MLL formulae are ranged over by meta-variables $A, B, C, \cdots$. A *literal* of MLL is either a propositional atom, or the linear negation of a propositional atom. We read $\alpha^\perp$ as "α-perp". The sequent calculus of MLL is presented as follows:

$$\text{(Identity)}\ \frac{}{\vdash \alpha^\perp, \alpha} \qquad \text{(Exchange)}\frac{\Gamma}{\sigma(\Gamma)}$$

$$\text{(Tensor)}\ \frac{\vdash \Gamma, A \qquad \vdash \Delta, B}{\vdash \Gamma, \Delta, A \otimes B} \qquad \text{(Par)}\ \frac{\vdash \Gamma, A, B}{\vdash \Gamma, A \propto B}$$

$$\text{(Cut)}\ \frac{\vdash \Gamma, A \qquad \vdash \Delta, A^\perp}{\vdash \Gamma, \Delta}$$

In the rule (Exchange), σ denotes an operation that permutes the formulae in the sequent Γ. Note that the identity axiom applies only to propositional atoms. This simplifies our argument but does not affect MLL-provability. An MLL sequent which is derivable from the above rules is called an MLL *theorem*, or simply a theorem. The standard reference is [Gir87]; for a systematic treatment, see also Troelstra's book [Tro92].

Proof Nets Girard's *proof nets* are a representation of the canonical proofs for sequents of the multiplicative fragment of Linear Logic [Gir87]. In the same way that Prawitz's Natural Deduction system NJ (as opposed to the Sequent Calculus LJ) faithfully captures proofs of the intuitionistic fragment $(\wedge, \rightarrow, \forall)$, so proof nets embody MLL proofs in their most primitive form. Remarkably, unlike the system NJ, proof nets denote proofs of the *classical* MLL sequents (where sequents are symmetric in a left / right fashion) in a completely satisfactory way. For instance, there is a strong normalization theorem for the process of cut-elimination in proof nets.

It is convenient to introduce the definition of a proof net via the intermediate notion of a proof structure. Formally, we define a proof net to be a proof structure which is sound in a sense to be clarified shortly. A *proof structure* is the formation "tree"[3] of an MLL sequent in which a propositional atom say α, always occurs an even number of times, half of which positively *i.e.* α, and half of which negatively, meaning $\alpha^\perp$. Moreover, each proof structure is equipped with a specific way of matching the two halves, linking the occurrence of a positive atom α uniquely to an occurrence of the negated atom $\alpha^\perp$. More precisely, we may define a proof structure to be a pair (Γ, ϕ), where Γ is an MLL sequent, and ϕ is an endofunction on the set, say O of occurrences of literals in Γ which satisifies the following properties: for o ranging over the set O,

- *fixpoint freeness*: $\phi(o) \neq o$,
- *involution*: if o is the occurrence of an atom α, then $\phi(o)$ is the occurrence of $\alpha^\perp$; further, $\phi(\phi(o)) = o$.

Thus, the map ϕ specifies the axiom links of the proof structure, all other information is already conveyed by Γ.

To identify the "sound" proof structures, Girard introduced the notions of a *switching*, and a *trip* taken by an information particle in a proof structure with a given switching. Roughly speaking, a *long trip* is a round trip which visits every node of the proof structure (with a given switching); and a *short trip* is a trip which fails to do so. A proof structure is defined to be a *proof net* if there is no switching with respect to which short trip is admitted. Girard proved an important characterization theorem for proof nets as follows (see [Gir87] for the definitive account):

Theorem 2.1 (Girard) *If π is a proof of an* MLL *sequent Γ, then we can naturally associate with π a proof net* PN(π) *which has exactly the same multiset of terminal formulae as the multiset of those occurring in Γ. Given any proof net β, there is a proof π of an* MLL *sequent such that $\beta =$* PN(π). $\Box$

Danos-Regnier Graph More recently, Danos and Regnier [DR89] (see [Bel93] for a very readable account) found an alternative soundness condition for proof nets. They avoid the need to reason with the bi-directional flows of information particles through a formula — an intuitive but somewhat complicated conceptual device which was used to great effect by Girard. Instead, Danos and Regnier introduce a simplied notion of switching, which we call DR-switching. Given such a switching, say S, they then examine structural properties of an associated undirected graph $\mathrm{Gr}(\Gamma, \phi, S)$ of the proof structure (Γ, ϕ). Recall the following MLL *links*:

axiom link	tensor link	par link
$\dfrac{}{A \qquad A^\perp}$	$\dfrac{A \qquad B}{A \otimes B}$	$\dfrac{A \qquad B}{A \propto B}$

Formally, we define a *DR-switching* S for an MLL sequent Γ to be an assignment of L or R to each occurrence of $\propto$ in Γ (the labels L and R refer respectively to the left and the right premise of the par link in question). Unlike Girard's switching, DR-switching does not concern tensor links.

Given a DR-switching S for a proof structure (Γ, ϕ), the associated undirected *DR-graph* $\mathrm{Gr}(\Gamma, \phi, S)$ is defined as follows:

- the vertices are the formulae of (Γ, ϕ): one vertex for each formula,

[3] The structure is not really a tree as is normally understood since it may have many roots. The leaves of a proof structure are labelled with the literals, and the roots are labelled with the formulae of the sequent.

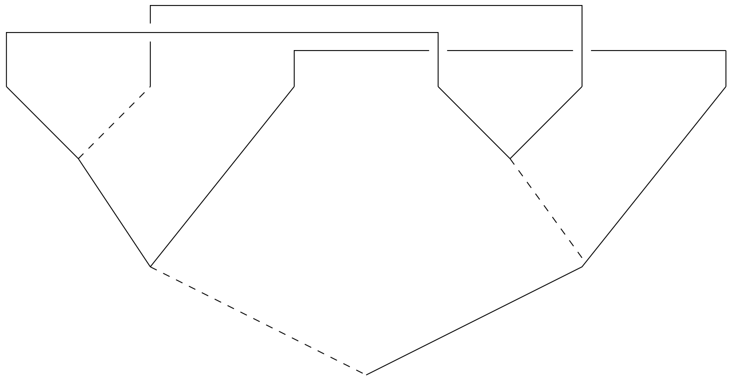

Figure 1: The DR-graph of the sequent $((A^{\perp} \propto B^{\perp}) \otimes C^{\perp}) \propto ((A \otimes B) \propto C)$ with a switching S. (Each dotted line segment corresponds to one of the two edges of a par link not indicated by the DR-switching in question).

- there is an edge between vertices P and Q whenever
  (1) P and Q are the conclusions of an axiom link, or
  (2) P is a premise, and Q the conclusion of a tensor link, or
  (3) Q is the conclusion of a par link, and P is the one premise of the par link indicated by the DR-switching.

Note that a DR-graph may be non-planar. For an example, see the graph in Figure 1.

Example Consider the sequent $\Gamma(\alpha, \beta, \gamma) \equiv ((\alpha^{\perp} \propto \beta^{\perp}) \otimes \gamma^{\perp}) \propto ((\alpha \otimes \beta) \propto \gamma)$. For a precise representation, we prefer to write the sequent in a tediously explicit way which gives a new name to each occurrence of the same literal, as well as a new name to each occurrence of the par connective. For example, the same sequent Γ is written as follows:

$$((\alpha_1^{\perp} \propto_1 \alpha_2^{\perp}) \otimes \alpha_3^{\perp}) \propto_2 ((\alpha_4 \otimes \alpha_5) \propto_3 \alpha_6)$$

with the axiom links represented in terms of the following fixpoint-free involution $\phi : \{1, \cdots, 6\} \rightarrow \{1, \cdots, 6\}$:

$$\phi \quad : \quad \begin{cases} 1 & \longleftrightarrow & 4 \\ 2 & \longleftrightarrow & 5 \\ 3 & \longleftrightarrow & 6. \end{cases}$$

Consider the following DR-switching $S : \{1, 2, 3\} \rightarrow \{L, R\}$

$$S \quad : \quad \begin{cases} 1 & \mapsto & L \\ 2 & \mapsto & R \\ 3 & \mapsto & R. \end{cases}$$

The DR-graph $\mathrm{Gr}(\Gamma, \phi, S)$ is shown in Figure 1.

In the sequel, we shall appeal to the following useful characterization [DR89]:

Theorem 2.2 (Danos-Regnier) *A proof structure* (Γ, ϕ) *is a proof net if, and only if, for every DR-switching* S*, the associated DR-graph* $\mathrm{Gr}(\Gamma, \phi, S)$ *is connected and acyclic.* $\square$

3 Finite, Fair Games

Notation Let M be a set of formal symbols. We let s and t range over M^*, the set of finite sequences of M; and write s_i for the i-th element of the sequence s, and $|s|$ for its length. For $1 \leqslant i \leqslant |s|$, the initial sequence of s truncated at length i is written $s \upharpoonright i$; so $s \upharpoonright i \equiv s_1 \cdots s_i$. For $s, t \in M^*$, we write $s \leqslant t$ just in case s is an initial subsequence of t i.e. $|s| \leqslant |t|$ and $s_i = t_i$ for each i such that $1 \leqslant i \leqslant |s|$. Clearly $\langle M^*, \leqslant \rangle$ is a set partially ordered by $\leqslant$; the least element is the empty sequence ϵ. We write the concatenation of two sequences s and t as $s \cdot t$. In fact, $\langle M^*, \leqslant \rangle$ is a *tree*. Recall that a poset $\langle T, \leqslant \rangle$ is a *tree* if for each $t \in T$ the set $\downarrow t \quad \{ s \in T : s < t \}$ of all prefixes of t is well-ordered by $<$ (see *e.g.* [Dev80]). A *path* in the tree $\langle T, \leqslant \rangle$ is a linearly ordered subset of T closed under prefixes. We largely follow the notational convention in [AJ92].

Finite, Fair Games

We introduce an abstract notion of a two-person game which are played according to the following ground rules:

- Only the Opponent can start a play.
- Thereafter, the play alternates strictly between Player and Opponent: each player makes one move at a time.
- Each play is bound to terminate after finitely many moves.
- The outcome of each play is never ambiguous: no draw is possible.
- A player loses a play whenever he is unable to respond to a move made by the other player.

Definition 3.1 A (finite) *fair game* is a structure $\langle M_A, \lambda_A, F_A \rangle$ satisfying the following conditions:

Moves The set M_A of *moves* has cardinality at least 2. Moves are usually ranged over by meta-variables m, a, b, m_i, a_j, b_k *etc.*

Labelling Function $\lambda_A : M_A \to \{ P, O \}$ is the labelling function which indicates whether a move is to be made by Player or by Opponent. We shall also write $\circ$ for Player and $\bullet$ for Opponent. We define the set M_A^+ of Player's moves as $\lambda_A^{-1}(P)$, and the set M_A^- of Opponent's moves as $\lambda_A^{-1}(O)$; also, $\overline{P} = O$ and $\overline{O} = P$.

$M_A^{\circledast}$ is defined to be the set of finite, alternately-labelled sequences of moves. We use meta-variables s, t *etc.* to range over $M_A^{\circledast}$. Formally,

$$M_A^{\circledast} \quad \stackrel{\text{def}}{=} \quad \{ s \in M_A^* : \forall i . 1 \leqslant i < |s| . \lambda_A(s_i) = \overline{\lambda_A(s_{i+1})} \}.$$

For any $s \equiv a_1 \cdots a_n \in M_A^{\circledast}$ with $n \geqslant 1$, and for $\alpha, \beta \in \{ \bullet, \circ \}$, we say that s *satisfies the shape* $\alpha \cdots \beta$ *in the game* A just in case $\lambda_A(a_1) = \alpha$ and $\lambda_A(a_n) = \beta$. So, for example, s satisfies $\circ \cdots \bullet$ means that the sequence s begins with a Player's move and ends with an Opponent's move. Similarly, if s satsfies $\cdots \bullet$, then s is a sequence that ends with an Opponent's move, in which case $n \geqslant 1$. Note that a sequence of length one satisfies either $\bullet \cdots \bullet$ or $\circ \cdots \circ$.

Fair Positions F_A, which must be a non-empty anti-chain in the poset $\langle M_A^{\circledast}, \leqslant \rangle$, is the set of *maximal, fair positions* (or plays) such that

(1) every element $s \in F_A$ is a sequence of even length,

(2) there are elements $s, t \in F_A$ such that s satisfies $\circ \cdots \bullet$ and that t satisfies $\bullet \cdots \circ$.

(We will justify the use of the adjective "maximal" shortly.) Hence, we see immediately that $\epsilon \notin F_A$, and that no element of F_A can satisfy $\circ \cdots \circ$ or $\bullet \cdots \bullet$.

Positions Given the above data, we define the collection P_A of *positions* of the game A as the least prefix-closed subset of $M_A^{\circledast}$ containing F_A. By construction, F_A is just the subset of P_A consisting of all $\leqslant$-maximal elements. It is in this sense that we call F_A the collection of *maximal*, fair positions of A.

Note that in our formulation of games, the notion of a position is a derived one; it is the notion of maximal positions (deemed to be fair) which is prior. It is helpful to bear in mind the following characterization of a position in a game A:

> A finite, alternately-labelled sequence of moves is a valid position in a finite, fair game if, and only if, it can be extended to a maximal, fair position.

Finally, we require the game A to satisfy a *finiteness* condition:

$$\text{"the tree } \langle\, P_A, \leqslant \rangle \text{ has no infinite paths"};$$

or equivalently, it has no infinite sequence of $<$-increasing elements. There is no requirement for the set F_A to have finite cardinality, nor the tree P_A to be finitely-branching.

In the following, unless otherwise stated, whenever the notion of a game is mentioned, we mean a finite, fair game. Also, whenever we mention a maximal position, we mean a maximal, fair position. Note that a finite, fair game $\langle\, M_A, \lambda_A, F_A \,\rangle$ is an instance of a game $\langle\, M_A, \lambda_A, P_A, W_A \,\rangle$ in the Abramsky-Jagadeesan sense [AJ92]. Since the length of sequences in P_A is bounded, the set W_A of winning infinite sequences is vacuously empty. At this juncture, we should acknowledge some useful ideas of Huth [Hut93]).

Strategies, Counter Strategies and Winning Strategies

In the development that follows, Player and Opponent play against one another in a game by pitting strategies against counter strategies. We always adopt the Player's perspective, and reserve the word "strategy" to mean Player's strategy, and "counter strategy" to mean Opponent's strategy. We think of a strategy or a counter strategy abstractly as a rule that prescribes a response for the player concerned at various positions. A strategy is said to be *total* if it prescribes a response for every possible position in the game; otherwise, it is said to be *partial*. A winning strategy is one which defeats all counter strategies. We first formalise the notion of a strategy.

Definition 3.2 A *strategy* σ for Player in the game A is defined to be a non-empty prefix-closed subset of P_A satisfying the following conditions:

(s1) *O-to-start*: for any $a \cdot s \in \sigma$, $\lambda_A(a) = O$;

(s2) *Determinacy*: for any $s \in \sigma$, if it is Player's turn to play at position s, and that $s \cdot a$ and $s \cdot b \in \sigma$ then $a = b$;

(s3) *Contingent Completeness*: for any $s \in \sigma$, if it is Opponent's turn to move at position s, and that $s \cdot a \in P_A$, then $s \cdot a \in \sigma$.

Note that our notion of strategy is partial, *i.e.* suppose σ is an arbitrary strategy and $s \in \sigma$ is a position with Player to move, our definition does not require σ to respond with a P-move. However, if σ has a response *i.e.* $s \cdot a \in \sigma$ for some $a \in M_A$, then the determinacy condition above ensures that the response a is not only legal (*i.e.* $s \cdot a \in P_A$), but also unique. In fact, given any strategy σ, there is a unique partial function $\hat{\sigma}$ mapping from the set of positions at which Player is to move, to the set of Player's moves such that

$$s \cdot a \in \sigma \quad \Longleftrightarrow \quad [\hat{\sigma} \text{ is defined at } s, \text{ and } \hat{\sigma}(s) = a].$$

A *counter-strategy* (*i.e.* a strategy for Opponent to play) is defined in exactly the same way as a strategy except that in conditions (s2) and (s3), the words Player and Opponent are interchanged.

We are now in a position to define the play $\sigma|\tau$ resulting from pitting a strategy σ (followed by Player) against a counter-strategy τ (played by Opponent). Note that $\sigma \cap \tau$ is a prefix-closed chain in P_A, and each such chain has a least upper bound, since by definition, P_A does not have any infinite $<$-increasing chain. We define

$$\sigma|\tau \stackrel{\text{def}}{=} \bigsqcup \sigma \cap \tau.$$

We say the strategy σ defeats the counter-strategy τ if the resultant play $\sigma|\tau$ is a position at which O is to move (but at which τ is unable to respond with any move). A strategy for the game A is said to be *winning* if it defeats all counter-strategies in A.

Game Constructions

In contrast to the Abramsky-Jagadeesan setting [AJ92], the "empty game" (where neither Player nor Opponent has any move to make) is not a finite, fair game. The simplest possible finite, fair game $G_{\min}$ has two moves $\{ p, o \}$ where p (and respectively o) is the only move that Player (and respectively Opponent) can make. The maximal positions of $G_{\min}$ are just po and op.

In the following, we introduce two basics ways by which games can be composed: linear negation, and par. We let A and B range over finite, fair games.

Linear Negation The game $A^{\perp}$, read "A perp", is the same game as A except that Player and Opponent are interchanged. More formally,

- $M_{A^{\perp}}$ is just M_A;
- $\lambda_{A^{\perp}} \stackrel{\text{def}}{=} \overline{\lambda_A}$;
- $F_{A^{\perp}}$ is F_A.

It is easy to check that $(A^{\perp})^{\perp}$ is identical to A.

Par Game The game $A \propto B$, read "A par B", is a game which is obtained from the game A and the game B by playing the two games in parallel, and in a fair manner. By parallel play, we do not mean parallelism in the so-called *true concurrency* sense *i.e.* the possibility of a player making two moves at precisely the same moment in time (though this is an idea which definitely deserves further study in Game semantic terms). Instead, a sequence of moves is a parallel play of two games if it is obtained by interleaving (in a suitable way) two plays, one move from each game respectively. The interleaving view of parallism has proved to be a successful and powerful perspective in the study of Concurrency in Theoretical Computer Science (see *e.g.* the book by N. Francez [Fra86], or [Par81]).

Fairness is a notion of fundamental importance in the organization of concurrent computations. Intuitively, a way of performing or organising two computations in parallel is said to be fair if equal (in some sense) attention, or computing power, or resourses is given to both. We have chosen a strict sense of fairness in the interpretation of a par game: fairness is expressed in modal terms as the affirmation of a possible event in the future. In our setting, to play a par game fairly is to allow a move to be made by a player only on condition that that particular move *can* lead to a fair, maximal position in the par game. The formal definition follows.

- $M_{A \propto B}$ is just the disjoint union $M_A + M_B$;
- $\lambda_{A \propto B}$ is the source tupling[4] $[\lambda_A, \lambda_B]$;
- $F_{A \propto B}$ is defined to be the set of all finite, alternately-labelled sequences $s \in M^{\circledast}_{A \propto B}$ satisfying:
  (1) the projection of s onto moves in the game A, (written $s \upharpoonright A$) is in F_A, and similarly for projection onto B, i.e. $s \upharpoonright B \in F_B$;
  (2) if two successive moves in s are in different components, then it is Player who has switched component. Formally, we refer to this condition as the *switching convention for par game.*

Whenever (1) holds, we observe that s can be reconstructed from $s \upharpoonright A$ and $s \upharpoonright B$ by an appropriate "interleaving" of the two.

We need to check that the par construction gives rise to a finite, fair game. It is clear that $F_{A \propto B}$ thus defined is a non-empty anti-chain in $\langle M^{\circledast}_{A \propto B}, \leqslant \rangle$, and that every element of $F_{A \propto B}$ is of even length because of condition (1), and that the elements of F_A and F_B are of even length. Suppose $s \equiv u \cdot a$ and s' are both elements of F_A with u satisfying the shape $\bullet \cdots \bullet$ (and so $|u| \geqslant 1$), and that s satisfies $\bullet \cdots \circ$, s' satisfies $\circ \cdots \bullet$; and that $t \in F_B$ satisfies $\circ \cdots \bullet$. Note that both $s' \cdot t$ and $u \cdot t \cdot a$ are elements of $F_{A \propto B}$, and $s' \cdot t$ satisfies the shape $\circ \cdots \bullet$ and $u \cdot t \cdot a$ satisfies the shape $\bullet \cdots \circ$ in the game $A \propto B$.

Tensor Game The tensor game construction is defined in terms of linear negation and par constructions. We define the tensor game $A \otimes B$ to be $(A^{\perp} \propto B^{\perp})^{\perp}$. It is easy to check that

- $M_{A \otimes B}$ is the disjoint union $M_A + M_B$;
- $\lambda_{A \otimes B}$ is the source tupling $[\lambda_A, \lambda_B]$;
- $F_{A \otimes B}$ is defined to be the set of all finite, alternately-labelled sequences $s \in M^{\circledast}_{A \propto B}$ satisfying:
  (1) $s \upharpoonright A \in F_A$ and $s \upharpoonright B \in F_B$;
  (2) if two successive moves in s are in different components, then it is Opponent who has switched component. Formally, we refer to this condition as the *switching convention for tensor game.*

Our notion of a fair game is specified quite differently from the Abramsky-Jagadeesan presentation [AJ92]. We regard the maximal, fair positions as primary, and introduce positions of the game as a derived notion, whereas in their setting, positions are prior and maximal positions do not play a prominent role. So in our setting, the maximal, fair plays of a compound game is quite appropriately defined in terms of the maximal, fair plays of the respective constituent games. However, as we shall

[4] We write $\mathsf{inl} : A \to A + B$ and $\mathsf{inr} : B \to A + B$ as the canonical injective maps. In general, whenever $f : A \to C$ and $g : B \to C$ are set-theoretic functions, we write $[f, g] : A + B \to C$ as the unique function satisfying the usual universal condition $f = [f, g] \circ \mathsf{inl}$, and $g = [f, g] \circ \mathsf{inr}$.

see shortly in the proof of the compositionality of winning strategies, it is technically convenient to characterise positions of a compound game (in our case, this just refers to a par game) in terms of the positions of the respective constituent games. We now present these characterizations.

Proposition 3.3 *For any $s \in M^{\circledast}_{A_0 \propto A_1}$, $s \in P_{A_0 \propto A_1}$ if, and only if,*

(a1) s satisfies the switching convention for the game $A_0 \propto A_1$,

(a2) $s^0 \stackrel{\text{def}}{=} s \restriction A_0 \in P_{A_0}$ and $s^1 \stackrel{\text{def}}{=} s \restriction A_1 \in P_{A_1}$,

(a3) if any one of s^0 or s^1 is a maximal, fair position in the respective constituent game, and that it satisfies the shape $\bullet \cdots \circ$, then the other is also a maximal, fair position, and satisfies the shape $\circ \cdots \bullet$.

Proof $(\Rightarrow)$: Suppose $s \in P_{A_0 \propto A_1}$, or equivalently, $s \leqslant l$ for some $l \in F_{A_0 \propto A_1}$. Since l satisfies the switching convention for the par game $A_0 \propto A_1$, *a fortiori*, the condition (a1) holds. Define $l^i \equiv l \restriction A_i \in F_{A_i}$ for $i = 0, 1$. It is easy to see that $s^i \leqslant l^i$, thus establishing (a2). To show (a3), suppose for i equals either 0 or 1, $s^i \in F_{A_i}$ and s^i satisfies the shape $\bullet \cdots \circ$ in the game A_i. There are two cases:

- *Either* $s = l$, that is to say, $s \in F_{A_0 \propto A_1}$ which implies $s^{1-i} \in F_{A_{1-i}}$. This means that s^{1-i} satisfies either $\bullet \cdots \circ$ or *po*. Since s satisfies the switching convention for the par game $A_0 \propto A_1$, the former has to be ruled out. Hence s^{1-i} satisfies the shape $\circ \cdots \bullet$ in the game A^{1-i}.
- *Or* $s < l$, in which case $s \cdot m \cdot \vec{a} = l$ and m must be an O-move in A_i because l satisfies the switching convention for the game $A_0 \propto A_1$, but this contradicts the $\leqslant$-maximality of s^i in the tree P_{A_i}.

$(\Leftarrow)$: Suppose $s \in M^{\circledast}_{A_0 \propto A_1}$ satisfies the three conditions (a1), (a2) and (a3). We tabulate all the possible cases and prove that for each case, $s \leqslant l$ for some $l \in F_{A_0 \propto A_1}$.

$s^0 \equiv s \restriction A_0$ maximal in P_{A_0}	$s^1 \equiv s \restriction A_1$ maximal in P_{A_1}	s ends with a P-move (w.l.o.g. in A_0)	s ends with an O-move (w.l.o.g. in A_0)
yes	yes	I	II
yes	no	III	IV
no	yes	V	VI
no	no	VII	VIII

In both cases I and II, we have $s^i \in F_{A_i}$. In view of (a1), we have $s \in F_{A_0 \propto A_1}$ immediately and so, $s \in P_{A_0 \propto A_1}$. Case III violates (a3). For case IV, by assumption, $s^1 \cdot t \in F_{A_1}$ for some $t \in M^{\circledast}_{A_1}$; then observe that $s \leqslant s \cdot t \in F_{A_0 \propto A_1}$. A similar argument settles cases V and VI. For case VII, there are two cases:

(1) s^0 satisfies the shape $\bullet \cdots \circ$ in the game A_0,

(2) s^0 satisfies the shape $\bullet \cdots \bullet$ in the game A_0.

For (1), by (a2), say $s^0 \cdot a_1 \cdots a_n \cdot b \in F_{A_0}$ for $n \geqslant 1$ with a_n an O-move, and that $s^1 \cdot t \in F_{A_1}$ with t satisfying the shape $\circ \cdots \bullet$ in game A_1. Then $s \cdot \vec{a} \cdot t \cdot b \in F_{A_0 \propto A_1}$. For (2), by (a2), say $s^0 \cdot a_1 \cdots a_n \in F_{A_0}$ with a_n an O-move and that $s^1 \cdot t \in F_{A_1}$ with t satisfying either the shape $\circ \cdots \bullet$ or $\bullet \cdots \bullet$ in game A_1. Then $s \cdot \vec{a} \cdot t \in F_{A_0 \propto A_1}$. The argument for case VIII is similar. $\square$

We can extend the characterization to a general "n-ary" par-term $A_1 \propto \cdots \propto A_n$, and similarly to a general tensor-term $A_1 \otimes \cdots \otimes A_n$: the order in which we parse the binary connective is immaterial.

Corollary 3.4 *Let $n \geqslant 2$.*

(i) For any $s \in M^{\circledast}_{A_1 \propto \cdots \propto A_n}$, $s \in P_{A_1 \propto \cdots \propto A_n}$ if, and only if,

- *(a1) s satisfies the switching convention for the par game $A_1 \propto \cdots \propto A_n$,*
- *(a2) for each $1 \leqslant i \leqslant n$, the projection of s onto the i-th game, written s^i, belongs to P_{A_i},*
- *(a3) if any one of $s^1, \cdots, s^n$ is a maximal, fair position in the respective game, and that it satisfies the shape $\bullet \cdots \circ$, then all the rest must also be maximal, fair positions, and satisfy the shape $\circ \cdots \bullet$.*

(ii) For any $s \in M^{\circledast}_{A_1 \otimes \cdots \otimes A_n}$, $s \in P_{A_1 \otimes \cdots \otimes A_n}$ if, and only if,

- *(a1) s satisfies the switching convention in the tensor game $A_1 \otimes \cdots \otimes A_n$,*
- *(a2) for each $1 \leqslant i \leqslant n$, the projection of s onto the i-th game, written s^i, belongs to P_{A_i},*
- *(a3) if any one of $s^1, \cdots, s^n$ is a maximal fair position in the respective game, and that it satisfies the shape $\circ \cdots \bullet$, then all the rest must also be maximal fair positions, and satisfy the shape $\bullet \cdots \circ$.* □

More generally, we can state the following sufficient condition for an alternately-labelled sequence of moves to be a position of a general finite, fair game.

Corollary 3.5 *Given a sequent $\Gamma(A_1, \cdots, A_l)$ where all the A_js are finite, fair games. We may write the set $M_{\Gamma(\vec{A})}$ as the disjoint union $M_{C_1} + \cdots + M_{C_n}$ where the C_is are all the literals in the sequent $\Gamma(\vec{A})$ (so each C_i is either A_j, or $A_j^{\perp}$ for some j). Take any alternately-labelled sequence s of moves in the game $\Gamma(\vec{A})$, whenever the following conditions are satisfied:*

(1) s follows the switching convention,

(2) for each $1 \leqslant i \leqslant n$, the projection $s \upharpoonright C_i \in P_{C_i}$,

(3) for each $1 \leqslant i \leqslant n$, the projection $s \upharpoonright C_i$ is not a maximal, fair position of the game C_i,

then s is a position of the game $\Gamma(\vec{A})$, i.e. $s \in P_{\Gamma(\vec{A})}$. □

A Category of Games and Winning Strategies

We can now define a category $\mathbb{G}$ of games where

- *objects* are finite, fair games, and
- *morphisms* are winning strategies; more precisely, given two such games A and B, a morphism σ from A to B is just a winning strategy σ of the game $A^{\perp} \propto B$.

[***I intend to modify this paragraph.***] Since $A \multimap B$ is equivalent to $A^{\perp} \propto B$, a morphism from A to B is a winning strategy say σ, for the game $A \multimap B$. Such a strategy may be understood as a way of transforming a winning strategy for the game A, say θ, to a winning strategy $\sigma(\theta)$ for the game B. Let θ be a winning strategy for the game A, and σ a winning strategy for the game $A \multimap B$. We now give an idea of how the transformed strategy $\sigma(\theta)$ is defined in terms of its interaction with a counter strategy for the game B. Suppose Opponent starts in the game B by making the move b_1. Player regards the Opponent as having made the move b_1 in the par game $A^{\perp} \propto B$, and plays the strategy σ against Opponent. Suppose σ prescribes a move in B, then we regard this move as Player's move in the original game B as prescribed by $\sigma(\theta)$. Now suppose σ switches game to $A^{\perp}$, by making the move, say a_1. We regard a_1 as Opponent's move in the game A, and obtain a Player's response a_1' in the same game A by consulting θ. We then regard a_1' as Opponent's move in the game $A^{\perp} \propto B$, and so, σ has a response to it. Since our games are finite, the interplay between σ and θ in the game A must end eventually. In fact, It must end with Opponent's move say, a'' in $A^{\perp} \propto B$ because θ is a winning strategy (for Player). Since σ is a winning strategy, it has a responce to a'' in B. And at this position, either the Opponent is defeated, or it has a response, and that σ's response is necessarily in the component B. The play is then confined to B until Opponent is defeated.

We first define how winning strategies compose, and then show that the above data specify a category.

Composition of Strategies If σ is a strategy for the game $A^{\perp} \propto B$ and τ a strategy for the game $B^{\perp} \propto C$, we define

$$S \stackrel{\text{def}}{=} \{ s \in \mathcal{L}(A, B, C) : s \upharpoonright A, B \in \sigma, s \upharpoonright B, C \in \tau \};$$

and define $\sigma;\tau \stackrel{\text{def}}{=} \{ s \upharpoonright A, C : s \in S \}$.

Proposition 3.6 *If σ is a winning strategy for the game $A^{\perp} \propto B$, and that τ is a winning strategy for the game $B^{\perp} \propto C$, then $\sigma;\tau$ is a winning strategy for the game $A^{\perp} \propto C$.*

Proof For any $s \in \sigma;\tau$ with $|s| \geqslant 1$, write $\overline{s}$ for an element in S such that $\overline{s} \upharpoonright A, C = s$ (such an element exists by assumption). We have

$$\begin{array}{rcccccc} s^1 & \stackrel{\text{def}}{=} & \overline{s} \upharpoonright A, B & \in & \sigma & \subseteq & P_{A^{\perp} \propto B}; \\ s^2 & \stackrel{\text{def}}{=} & \overline{s} \upharpoonright B, C & \in & \tau & \subseteq & P_{B^{\perp} \propto C}. \end{array}$$

It is easy to see that:

$$\begin{array}{lrcl} \text{(p1)} & s \upharpoonright A & = & s^1 \upharpoonright A \\ \text{(p2)} & s \upharpoonright C & = & s^2 \upharpoonright C \\ \text{(p3)} & s^1 \upharpoonright B & = & s^2 \upharpoonright B \end{array}$$

• $\sigma;\tau$ is clearly non-empty and prefix-closed.

• $\sigma;\tau \subseteq P_{A^{\perp} \propto C}$.
Let $s \in \sigma;\tau$, we show $s \in P_{A^{\perp} \propto C}$ by showing that conditions (a1), (a2) and (a3) of Proposition 3.3 are satisfied. Suppose for some $1 \leqslant i < |s|$, s_i and s_{i+1} are moves belonging to different games. W.l.o.g., say $s_i \in M_{A^{\perp}}$ and $s_{i+1} \in M_C$. Let $s_i \cdot b_1 \cdots b_n \cdot s_{i+1}$ for $n \geqslant 1$ be the appropriate segment in $\overline{s}$ which gets projected under games A, C onto $s_i s_{i+1}$, with $\vec{b} \subseteq M_B$. Now $\vec{b} \cdot s_{i+1}$ is a segment of s^2. Since s^2 satisfies the switching convention for the game $B^{\perp} \propto C$, we have $\lambda_{B^{\perp} \propto C}(s_{i+1}) = P$ i.e. $\lambda_{A^{\perp} \propto C}(s_{i+1}) = P$. Hence (a1) holds. To see (a2), because $s^1 \in P_{A^{\perp} \propto B}$, applying Proposition 3.3 to the game $A^{\perp} \propto B$, we get $s^1 \upharpoonright A \in P_{A^{\perp}}$. Hence by (p1), we have $s \upharpoonright A \in P_{A^{\perp}}$; similarly $s \upharpoonright C \in P_C$.

It remains to show (a3). W.l.o.g. suppose $s \upharpoonright A \in F_{A^\perp}$ and $s \upharpoonright A$ satisfies the shape $\bullet \cdots \circ$ in the game $A^\perp$. By (p1), we have $s^1 \upharpoonright A \in F_{A^\perp}$ and $s^1 \upharpoonright A$ satisfies the shape $\bullet \cdots \circ$ in the game $A^\perp$. Applying Proposition 3.3 to the game $A^\perp \propto B$, we have $s^1 \upharpoonright B \in F_B$ and $s^1 \upharpoonright B$ satisfies the shape $\bullet \cdots \circ$ in the game B; which is equivalent to $s^2 \upharpoonright B^\perp \in F_{B^\perp}$ and that $s^2 \upharpoonright B^\perp$ satisfies the shape $\circ \cdots \bullet$ in the game $B^\perp$ by (p3). Applying the same proposition to the game $B^\perp \propto C$, we have $s^2 \upharpoonright C \in F_C$ and $s^2 \upharpoonright C$ satisfies the shape $\bullet \cdots \circ$ in the game C. The result then follows from (p2).

• $\sigma;\tau$ satisfies (s1).
Let $s \in \sigma;\tau$, then $\overline{s}$ cannot begin with a move in B. For suppose it does, say it is an O-move in B. Then $\overline{s} \upharpoonright B^\perp, C \in \tau$ begins with a P-move in $B^\perp$, contradicting condition (s1) on the strategy τ. It cannot be a P-move in B either for it would then contradict condition (s1) on the strategy σ since $\overline{s} \upharpoonright A^\perp, B \in \sigma$.

Suppose s begins with a move m in the game $A^\perp$, then $\overline{s} \upharpoonright A^\perp, B \in \sigma$ begins with m which, by condition (s1) applied to σ, must be an O-move in the game $A^\perp$. A similar argumentapplies to the case of s beginning with a move in the game C.

• $\sigma;\tau$ satisfies (s2). We show that a stronger property holds:

> *for any $s \in \sigma;\tau$, if s satisfies the shape $\cdots \bullet$ in the game $A^\perp \propto C$, then there is a unique P-move m such that $s \cdot m \in \sigma;\tau$.*

Suppose $s \in \sigma;\tau$ ends with an O-move a in $A^\perp$, say. Then s^1 ends with the same O-move a in $A^\perp$. Since $s^1 \in \sigma$, there is a unique P-move m_1 in either game $A^\perp$ or B such that $s^1 \cdot m_1 \in \sigma \subseteq P_{A^\perp \propto B}$. There are two cases.

i. *Either* $m_1 \in M_{A^\perp}$; in which case, $\overline{s} \cdot m_1 \upharpoonright A, B = s^1 \cdot m_1 \in \sigma$. Hence, we have $s \cdot m_1 = \overline{s} \cdot m_1 \upharpoonright A, C \in \sigma;\tau$. Take m to be m_1.

ii. *Or* we have $m_1 \in M_B$. Then, by condition (s3) on τ, we have $\overline{s} \cdot m_1 \upharpoonright B, C = s^2 \cdot m_1 \in \tau$ such that $s^2 \cdot m_1$ ends with an Opponent's move in the game $B^\perp \propto C$. Hence, there is a unique $m_2 \in M_{B^\perp}$ or M_C such that $s^2 \cdot m_1 \cdot m_2 \in \tau$. Continuing in this fashion, because of the finiteness condition, we have for some finite n or n', either

$$s^1 \cdot m_1 \cdot m_2 \cdots m_n \cdot a \quad \in \quad \sigma \qquad \text{or} \qquad s^2 \cdot m_1 \cdot m_2 \cdots m_{n'} \cdot c \quad \in \quad \tau;$$

with $\vec{m} \subseteq M_B$, a is a P-move in $M_{A^\perp}$ or c a P-move in M_C. W.l.o.g. say the former, then we have $\overline{s} \cdot \vec{m} \cdot a \in S$, and so, $\overline{s} \cdot \vec{m} \cdot a \upharpoonright A, C = s \cdot a \in \sigma;\tau$. Take m to be a in this case.

• $\sigma;\tau$ satisfies (s3). We need to show: suppose $s \in \sigma;\tau$ such that s satisfies the shape $\cdots \circ$ in the game $A^\perp \propto C$ and $s \cdot m \in P_{A^\perp \propto C}$ then $s \cdot m \in \sigma;\tau$. W.l.o.g. suppose s ends with a P-move in $A^\perp$ and that m (can only be) an O-move in $A^\perp$. It suffices to show:

Claim: $s^1 \cdot m \in P_{A^\perp \propto B}$.
For then, because $s^1 \in \sigma$, we apply condition (s3) on the strategy σ to get $s^1 \cdot m \in \sigma$. Now, because $\overline{s} \cdot m \in S$, we have $s \cdot m = \overline{s} \cdot m \upharpoonright A, C \in \sigma;\tau$.

To prove the claim, we show that $s^1 \cdot m$ satisfies conditions (a1), (a2) and (a3) of Proposition 3.3 applied to the game $A^\perp \propto B$. Since $\overline{s} \cdot m \upharpoonright A, B \in \sigma \subseteq P_{A^\perp \propto B}$, by (p1), we know that $s^1 \cdot m$ satisfies the switching convention for the game $A^\perp \propto B$ *i.e.* (a1) holds. To see (a2), because $s \cdot m \in P_{A^\perp \propto C}$, by Proposition 3.3 (a2), $s \cdot m \upharpoonright A \in P_{A^\perp}$, and so because of (p1), $s^1 \cdot m \in P_{A^\perp}$. Because $s^1 \in \sigma \subseteq P_{A^\perp \propto B}$, so $s^1 \upharpoonright B = s^1 \cdot m \upharpoonright B \in P_B$. Finally, note that (a3) is vacuously true.

Now that we have shown that $\sigma;\tau$ is a strategy, the stronger variant of condition (s2) implies that $\sigma;\tau$ is a winning strategy. $\square$

Identity morphism Given any game A, the identity morphism is just the "copy-cat" strategy id_A for the game $A^\perp \propto A$ where id_A consists of all positions t of the game $A^\perp \propto A$ satisfying the properties:

- t begins with an Opponent's move,
- for any position s of the game $A^\perp \propto A$ such that $s \leqslant t$, if s is of even length, then $s \restriction A = s \restriction A^\perp$.

The identity strategy may equivalently be defined as the prefix-closure of the set consisting of sequences of the form:

$$\langle\, i, a_1 \,\rangle \cdot \langle\, i', a_1 \,\rangle \cdot \langle\, i', a_2 \,\rangle \cdot \langle\, i, a_2 \,\rangle \cdot \langle\, i, a_3 \,\rangle \cdots \langle\, i', a_{n-1} \,\rangle \cdot \langle\, i', a_n \,\rangle \cdot \langle\, i, a_n \,\rangle$$

where

- $a_1 a_2 \cdots a_n$ is a maximal position in the game A,
- $\langle\, i, a_1 \,\rangle$ is an Opponent's move in the game $A^\perp \propto A$, and that
- i and i' label different components of the game.

To show that $\mathbb{G}$ defines a category, it remains to show that composition is associative. [***I have yet to fill up the gap here.***]

Proposition 3.7 *$\mathbb{G}$ is a category.* □

Note that in contrast to the Abramsky-Jagadeesan setting, given any (finite, fair) game A, there is always a winning strategy for Player. This is essentially because maximal positions in a game are of even length, so if a play is started by one player (which, by convention, is always the Opponent), then it is necessarily the other player who ends the play. So given any position s at which Player is to make a move. Since it is always the Opponent who starts a play, s satisfies the shape $\bullet \cdots \bullet$, and so, it has odd length, and cannot be a maximal position. In fact, s is, by definition, a prefix of some maximal position, say t, so there is some move which Player can make. In particular, this means that given any two games A and B, there is always a winning strategy for the game $A^\perp \propto B$. In other words, there is always a morphism between any two objects of the category $\mathbb{G}$.

A Category of Games and History-free, Winning Strategies

Given any game A, we say that a strategy σ is *history-free* if there is a partial function $\sigma_f : M_A^- \rightharpoonup M_A^+$ from Opponent's moves to Player's moves such that for any position $s \cdot a$ at which Player is to move,

$$\hat{\sigma}(s \cdot a) \quad = \quad \begin{cases} f(a) & \text{if } f(a) \text{ is defined, and } s \cdot a \cdots f(a) \in P_A, \\ \text{undefined} & \text{else.} \end{cases}$$

It is easy to see that in this case, there is always a least partial function inducing σ; we write $\sigma = \sigma_f$, always meaning this unique f. Restricting to history-free strategy, it is no longer the case that every game has a winning strategy.

Proposition 3.8 *$\mathbb{G}_{\mathsf{hf}}$ is a subcategory of $\mathbb{G}$.*

Proof [***I have yet to fill up the gap here.***] □

$\mathbb{G}_e$ **as a $*$-Autonomous Category** [Bar79, Bar91] [See89]

Proposition 3.9 *(i)* $\mathbb{G}$ *is a $*$-Autonomous Category.*

(ii) $\mathbb{G}_e$ *is a sub-$*$-autonomous category of* $\mathbb{G}$.

Proof [***I have yet to fill up the gap here.***] $\square$

Generic Games and History-Free, Uniformly Winning Strategies

Up to this point, we have seen how fair games may be composed by the use of the linear logic connectives of tensor and par. Given any MLL sequent $\Gamma(\alpha_1, \cdots, \alpha_l)$ where $\alpha_1, \cdots, \alpha_l$ include all the propositional atoms occuring in Γ, and given an l-tuple of games $\vec{A} \equiv A_1, \cdots, A_l$, one can ask whether the composite game $\Gamma(\vec{A})$ (obtained by replacing the $\alpha_1, \cdots, \alpha_l$ by the games $A_1, \cdots, A_l$ respectively) has a history free, winning strategy. Since we are ultimately interested in proof-theoretic information about the MLL sequent Γ, and not about the prospect of winning a particular game $\Gamma(\vec{A})$, we should ask whether a history-free,winning strategy for $\Gamma(\vec{A})$ exists *uniformly* for each l-tuple $\vec{A}$.

Following Abramsky and Jagadeesan [AJ92], we formalise the notion of uniformity as a natural transformation, though we work with a different category. Given any two games A and B, we say that there is an *embedding* e from A to B, denoted $e : A \rightarrowtail B$, just in case

- $e : M_A \to M_B$ is an injective function,
- $e \circ \lambda_A = \lambda_B$,
- $e^*(F_A) \subseteq F_B$.

Quite naturally, we define a new category $\mathbb{G}_e$ of finite, fair games. Morphisms of the category are just the embeddings.

Lemma 3.10 *Tensor, par, and linear negation extend to covariant functors over* $\mathbb{G}_e$.

Any MLL formula $A(\alpha_1, \cdots, \alpha_l)$ with l propositional atoms induces a functor which we shall simply write (by abuse of notation) $A(\vec{\alpha}) : (\mathbb{G}_e)^l \to \mathbb{G}_e$. In the same way, any MLL sequent $\Gamma(\alpha_1, \cdots, \alpha_l)$ induces a functor $\Gamma(\vec{\alpha}) : (\mathbb{G}_e)^l \to \mathbb{G}_e$, with Γ interpreted simply as $\propto\Gamma$.

One may regard the sequent $\Gamma(\alpha_1, \cdots, \alpha_l)$, *qua* the induced functor over the category $\mathbb{G}_e$, as a "game scheme", or a *generic game*. What then does it mean to say that the generic game $\Gamma(\vec{\alpha})$ has a (history-free) winning strategy? Formally, we define a (history-free) winning strategy for the generic game $\Gamma(\vec{\alpha})$ to be a collection $\{\sigma_{\vec{A}}\}$ indexed over the objects of $(\mathbb{G}_e)^l$ such that for each l-tuple $\vec{A}$ of games, $\sigma_{\vec{A}}$ is a (history-free) winning strategy for the game $\Gamma(A_1, \cdots, A_l)$.

Given a sequent $\Gamma(\vec{\alpha})$, we can define two functors $M^-_{\Gamma(\alpha_1,\cdots,\alpha_l)}, M^+_{\Gamma(\alpha_1,\cdots,\alpha_l)} : (\mathbb{G}_e)^l \to \mathbb{SET}^{\rightharpoonup}$ by the following actions:

- *objects*: $M^-_{\Gamma(\alpha_1,\cdots,\alpha_l)}$ and $M^+_{\Gamma(\alpha_1,\cdots,\alpha_l)}$ map $\vec{A}$ to $M^-_{\Gamma(\vec{A})}$ and $M^+_{\Gamma(\vec{A})}$ respectively,
- *morphisms*: $M^-_{\Gamma(\alpha_1,\cdots,\alpha_l)}$ sends $\vec{e}$ to $M^-_{\Gamma(\vec{e})}$, similarly for $M^+_{\Gamma(\vec{e})}$.

If $\sigma = \{ \sigma_{\vec{A}} \}$ is a family of history-free strategies, then each $\sigma_{\vec{A}}$ is characterised by a partial map

$$f_{\vec{A}} \quad : \quad M^{-}_{\Gamma(\vec{A})} \quad \rightharpoonup \quad M^{+}_{\Gamma(\vec{A})}.$$

We say that the history-free strategy σ for the generic game $\Gamma(\vec{\alpha})$ is *uniform* if f is a natural transformation $f : M^{-}_{\Gamma(\vec{\alpha})} \to M^{+}_{\Gamma(\vec{\alpha})}$; that is to say, for each morphism $\vec{e} : \vec{A} \to \vec{B}$ in $(\mathbb{G}_{\mathrm{e}})^{l}$, the following diagram commutes:

$$\begin{tikzcd}
M^{-}_{\Gamma(\vec{A})} \arrow[r, "\sigma_{\vec{A}}"] \arrow[d, "\vec{e}"'] & M^{+}_{\Gamma(\vec{A})} \arrow[d, "\vec{e}"'] \\
M^{+}_{\Gamma(\vec{B})} \arrow[r, "\sigma_{\vec{B}}"] & M^{+}_{\Gamma(\vec{B})}
\end{tikzcd}$$

Theorem 3.11 (Soundness) *For any* MLL *sequent $\Gamma(\vec{\alpha})$, if $\Gamma(\vec{\alpha})$ is an* MLL*-theorem, then the generic game $\Gamma(\vec{\alpha})$ has a history-free, uniformly winning strategy.* □

4 Full Completeness Theorem

This section is devoted to a proof of the following theorem:

Theorem 4.1 (Full Completeness) *Given any* MLL *sequent Γ, if σ is a history-free, uniformly winning strategy for the generic game Γ, then σ is the denotation of a unique proof net.* □

A reduction argument First, two definitions. A sequent Γ is said to be *binary* if every propositional atom occuring in Γ does so precisely twice: once positively, and once negatively. Note that any binary sequent is a proof structure; and a proof structure is equivalent to a binary sequent by a renaming the propositional atoms. A sequent $\Gamma \equiv A_1, \cdots, A_n$ is said to be *semi-simple* if each formula A_i is built up from the literals by the tensor connective only. For example, the sequent $(\alpha_1 \otimes (\alpha_2^{\perp} \otimes \alpha_3)) \otimes \alpha_3^{\perp}, \alpha_1^{\perp} \otimes \alpha_2$ is semi-simple, but $(\alpha_1 \propto \alpha_2) \otimes \alpha_3, \alpha_1^{\perp} \otimes (\alpha_2^{\perp} \otimes \alpha_3)$ is not.

Our proof follows a reduction argument similar to AJ's approach, and it uses the following intermediate arguments:

- For any binary sequent Γ, there is a finite set of binary, semi-simple sequents $\{ \Gamma_1, \cdots, \Gamma_n \}$ such that $\vdash \Gamma$ if, and only if, for every i, $\vdash \Gamma_i$.
- If a binary, semi-simple sequent Γ (*qua* a generic game), has a history-free, uniformly winning strategy, then the associated DR-graph is disconnected and acyclic.

Copy-Cat Strategies

We are now in a position to state and prove a result crucial to our proof of the Full Completeness Theorem. The proposition may be stated informally in the following way:

> *for any sequent Γ, any history-free, uniformly winning strategy σ of Γ (regarded as a generic game) turns the sequent Γ into a proof structure; that is to say, the strategy σ defines a unique fixpoint-free, involution ϕ over the set of occurrences of literals in Γ, thereby specifying a set of axiom links matching the literals in Γ completely.*

To achieve this, we first show that such a strategy can only be the "copy-cat strategy".

Proposition 4.2 (Copy-Cat) *For any* MLL *sequent* $\Gamma(\alpha_1, \cdots, \alpha_l)$*, if the generic game (i.e. the induced functor* $\Gamma : (\mathbb{G}_e)^l \to \mathbb{G}_e$*) has a history-free, uniformly winning strategy* σ*, then the strategy must be a copy-cat strategy.*

Of course, we need to spell out precisely what it means for a history-free, uniformly winning strategy of a generic game Γ to be a copy-cat strategy. It is essentially this: σ determines a fixpoint-free, involution ϕ of $\{1, \cdots, n\}$ where n is the number of literals occuring in Γ. We can immediately infer that n is even.

We assume that for any l-tuple of games $A_1, \cdots, A_l$, there is an associated history-free winning strategy $\sigma_{\vec{A}}$ for the game $\Gamma(\vec{A})$ satisfying the uniformity condition. Let n be the number of literals in $\Gamma(\alpha_1, \cdots, \alpha_l)$. To say that σ is a copy-cat strategy is to assert the following: there is a map $\phi : \{1, \cdots, n\} \to \{1, \cdots, n\}$ satisfying the following conditions:

(c1) **copy-cat**
for any l-tuple of games $A_1, \cdots, A_l$ (suppose the set $M_{\Gamma(\vec{A})}$ of moves is the disjoint union $M_{C_1} + \cdots + M_{C_n}$, where each C_i is either A_j or $A_j^{\perp}$ for some $1 \leqslant j \leqslant l$), and for any $1 \leqslant i \leqslant n$, the following conditions are satisfied:

- for any $m \in C_i$ such that Player is to move at the position $\langle i, m \rangle$ in the game $\Gamma(\vec{A})$,

$$\sigma_{\vec{A}}(\langle i, m \rangle) \quad = \quad \langle \phi(i), m \rangle;$$

  note that $\phi(i)$ is independent of m,

- $$C_{\phi(i)} \quad = \quad \begin{cases} A_j^{\perp} & \text{if } C_i = A_j \text{ for some } j, \\ A_j & \text{else if } C_i = A_j^{\perp}. \end{cases}$$

(c2) **fixpoint-free involution**
for each $1 \leqslant i \leqslant n$, $\phi(i) \neq i$ (which already follows from (c1)), and $\phi(\phi(i)) = i$.

Note that the map ϕ is determined by the strategy, and is independent of the particular l-tuple $\vec{A}$ of games.

Before we prove the two conditions, it is convenient to introduce a way of constructing new games now. For any game A, and given a move $m_0 \in M_A$ and a fresh symbol m_1 (so $m_1 \notin M_A$), we define a new game $A[m_0 + m_1]$ such that at any position, a player (whether Player or Opponent) can make the move m_1 if, and only if, it can make the move m_0. More formally, the structure $A[m_0 + m_1]$ is specified by the following data:

- $M_{A[m_0+m_1]} \stackrel{\text{def}}{=} M_A \cup \{ m_1 \}$,
- the map $\lambda_{A[m_0+m_1]}$ is defined as follows: for any $m \in M_{A[m_0+m_1]}$,

$$\lambda_{A[m_0+m_1]}(m) \quad \stackrel{\text{def}}{=} \begin{cases} \lambda_A(m) & \text{if } m \neq m_1 \\ \lambda_A(m_0) & \text{else if } m = m_1. \end{cases}$$

- $F_{A[m_0+m_1]} \stackrel{\text{def}}{=} F_A \cup \{ s[m_1/m_0] : s \in F_A \}$ where $s[m_1/m_0]$ means "in the sequence s substitute m_1 for every occurrence of m_0".

It is clear that $A[m_0 + m_1]$ thus defined, is a finite, fair game. We state an obvious fact of the game $A[m_0 + m_1]$ which we shall use in the proof of the Copy-Cat Proposition.

Lemma 4.3 *We can identify two embeddings e_0, e_1 of A into $A[m_0 + m_1]$: the first embedding e_0 is just the inclusion map $M_A \subseteq M_{A[m_0+m_1]}$; the second e_1 sends m_0 to m_1, and sends every other element to itself.* $\square$

We are now ready to prove the Copy-Cat Proposition, and we do so by proving the two conditions (c1) and (c2).

Proof Suppose σ is a history-free, uniformly winning strategy of the generic game $\Gamma(\alpha_1, \cdots, \alpha_l)$ which contains n literals. For any l-tuple $A_1, \cdots, A_l$ of games, by our assumption, $\sigma_{\vec{A}}$ is a history-free, winning strategy for $\Gamma(\vec{A})$. We may express the set $M_{\Gamma(\vec{A})}$ of moves in the game $\Gamma(\vec{A})$ as the disjoint union $M_{C_1} + \cdots + M_{C_n}$ where each C_i is either A_j, or $A_j^{\perp}$ for some j. Fix any i such that $1 \leqslant i \leqslant n$. Let us say $C_i = A_j$, for some $1 \leqslant j \leqslant l$ (the case of $C_i = A_j^{\perp}$ is completely symmetrical, as we shall see). Since $\sigma_{\vec{A}}$ is history-free, for any $m \in C_i$ such that it is Player's turn to move at $\langle\, i, m \,\rangle$ (in the game $\Gamma(\vec{A})$), we may write

$$\sigma_{\vec{A}}(\langle\, i, m \,\rangle) \quad = \quad \langle\, \phi(i), m' \,\rangle,$$

where $\phi(i) \equiv i'$ takes a value in $\{\, 1, \cdots, n \,\}$. With reference to the preceding equation, we first prove the following:

Claim 1 *If any one of the following conditions holds,*

(a) $C_{i'} = A_k^{\perp}$ or A_k where $j \neq k$,

(b) $C_{i'} = A_j$ (then $m \neq m'$),

(c) $C_{i'} = A_j^{\perp}$ and $m' \neq m$,

then σ is not a uniform strategy. Whence, we may conclude that if σ is a history-free, uniformly winning strategy, then $\sigma_{\vec{A}}(\langle\, i, m \,\rangle) = \langle\, \phi(i), m \,\rangle$, where $C_{\phi(i)} = A_j^{\perp}$.

To prove the claim, suppose (a) holds. Without loss of generality, suppose $m' \in M_{A_k}$. The case of $m' \in M_{A_k^{\perp}}$ is similar, and so we omit it. We define a new l-tuple $\vec{A}'$ of games which is obtained from $\vec{A}$ by replacing A_k with the game $A_k[m' + m'']$, for some fresh symbol $m'' \notin M_{A_k}$. Consider the two obvious embeddings: $\vec{e}_0$ and $\vec{e}_1$ of $\vec{A}$ into $\vec{A}'$ which are built up in the obvious way from the respective embeddings identified in the preceding lemma. If σ were uniform, then we would have:

$$\begin{aligned} \sigma_{\vec{A}'} \circ \vec{e_0}(\langle\, i, m \,\rangle) \quad &= \quad \vec{e_0} \circ \sigma_{\vec{A}}(\langle\, i, m \,\rangle) \quad = \quad \langle\, i', m' \,\rangle, \\ \sigma_{\vec{A}'} \circ \vec{e_1}(\langle\, i, m \,\rangle) \quad &= \quad \vec{e_1} \circ \sigma_{\vec{A}}(\langle\, i, m \,\rangle) \quad = \quad \langle\, i', m'' \,\rangle \end{aligned}$$

The l.h.s. of both equations are equal to $\sigma_{\vec{A}'}(\langle\, i, m \,\rangle)$, but the r.h.s. of both are by design distinct. Hence, σ cannot be a uniform strategy.

Suppose (b) holds. Then $m' \in M_{A_j}$. We define a new l-tuple $\vec{A}'$ of games which is obtained from $\vec{A}$ by replacing A_j with $A_j[m' + m'']$, for some $m'' \notin M_{A_j}$. By the same reasoning as case (a), we see that σ is not a uniform strategy. The case of (c) is proved in essentially the same way. So this completes the proof of the Claim.

To prove (c1), it remains to prove:

Claim 2 *$\phi(i)$ thus defined is independent of the move m.*

Suppose for the sake of a contradiction, there are two distinct moves $m_1, m_2 \in M_{C_i}$ (note that we assume $C_i = A_j$, for some j) such that both $\langle i, m_1 \rangle$ and $\langle i, m_2 \rangle$ are positions at which Player is to move in the game $\Gamma(\vec{A})$ satisfying,

$$\sigma_{\vec{A}} \quad : \quad \begin{cases} \langle i, m_1 \rangle & \mapsto & \langle i_1, m_1 \rangle \\ \langle i, m_2 \rangle & \mapsto & \langle i_2, m_2 \rangle, \end{cases}$$

and such that $i_1 \neq i_2$ but $C_{i_1} = C_{i_2} = A_j^{\perp}$. There are two cases:

(b1) *either* there is no maximal position in the game A_j which begins with the move m_1, nor is there any that begins with m_2,

(b2) *or* the negation of (b1), that is to say, $m_1 t \in F_{A_j}$ for some t, or $m_2 t \in F_{A_j}$ for some t.

We consider the case (b1) first. Let m_0 be any fresh symbol (so $m_0 \notin M_{A_j}$). Consider a new sequence $s \equiv m_1 m_0 m_2 m_0$ where m_0 is regarded as Player's move, then s is a sequence of moves alternating between Player and Opponent of even length.

We define a new game $A_j[s]$ which is obtained by simply "grafting" a new maximal position s onto A_j. More formally, $A_j[s]$ is the new game defined as:

- $M_{A_j[s]} \stackrel{\text{def}}{=} M_{A_j} \cup \{ m_0 \}$,
- $\lambda_{A_j[s]}(m) \stackrel{\text{def}}{=} \begin{cases} \lambda_{A_j}(m) & \text{if } m \neq m_0 \\ P & \text{else if } m = m_0, \end{cases}$
- $F_{A_j[s]} \stackrel{\text{def}}{=} F_{A_j} \cup \{ s \}$.

Let $\vec{A}'$ be a new l-tuple of games which is obtained from $\vec{A}$ by replacing A_j with $A_j[s]$. As before, we write $M_{\Gamma(\vec{A}')}$ as the disjoint union $M_{C_1} + \cdots + M_{C_n}$. By the assumption of uniformity, $\sigma_{\vec{A}'}$ has exactly the same response as $\sigma_{\vec{A}}$ at positions $\langle i, m_1 \rangle$ and $\langle i, m_2 \rangle$ in the game $\Gamma(\vec{A}')$. We consider the trace of a play obtained by piting the strategy $\sigma_{\vec{A}'}$ against a particular counter strategy τ for the game $\Gamma(\vec{A}')$ as follows:

$$u \equiv \langle i, m_1 \rangle \cdot \langle i_1, m_1 \rangle \cdot \langle i_1, m_0 \rangle \cdot \langle i, m_0 \rangle \cdot \langle i, m_2 \rangle;$$

or more vividly, see the sequence u of moves presented graphically in Figure 2.

Note that the sequence of moves u is a position in the game $\Gamma(\vec{A}')$, by Corollary 3.5. We now explain why u is the (partial) trace of a valid play involving the strategy $\sigma_{\vec{A}'}$ and a counter strategy τ. The counter strategy τ begins with the move $\langle i, m_1 \rangle$. The strategy $\sigma_{\vec{A}'}$ replies with $\langle i_1, m_1 \rangle$ in the game C_{i_1}, according to our assumption. Note that the only maximal position in $C_{i_1} = A_j[s]^{\perp}$ beginning with m_1 is s, by assumption (b1). The counter strategy τ then makes the move $\langle i_1, m_0 \rangle$. The strategy $\sigma_{\vec{A}'}$ is forced to respond with m_0, by Claim 1; and this has to be in the sub-game C_i (that is to say, Player responds with $\langle i, m_0 \rangle$), since for Player to make a legal m_0 move in any sub-game C_k other than C_i at this stage of play, there must be a maximal position in C_k beginning with the move m_0, and by assumption no sub-game C_k has this property. The counter strategy τ now makes the move $\langle i, m_2 \rangle$; the strategy $\sigma_{\vec{A}'}$ then responds with $\langle i_2, m_2 \rangle$ by assumption, but such a move is not legal, since in C_{i_2}, by assumption, there is no maximal position that begins with m_2.

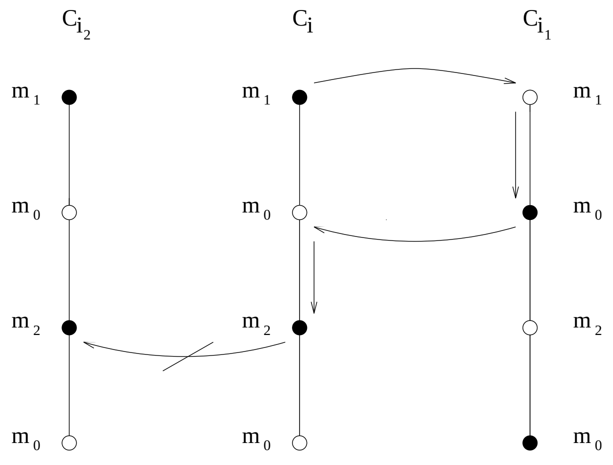

Figure 2:

Now, we consider the case (b2). In order to apply the same argument as case (b1), we consider the new l-tuple of games $\vec{A}'$ which is obtained from $\vec{A}$ by replacing A_j with A'_j, where A'_j is obtained from A_j by performing the two operations as follows:

(1) first, prune away any maximal position in A_j which begins with either the move m_1 or m_2,

(2) then graft the new position $s \equiv m_1 m_0 m_2 m_0$ onto the pruned structure (where as before, m_0 is deemed to be a Player's move).

Call the new structure $A_j^-[s]$. Note that $A_j^-[s]$ is a well-defined finite, fair game with $M_{A_j} \cup \{ m_0 \}$ as the set of moves, and

$$\{ s \in F_{A_j} : \ s \text{ does not begin with either } m_1 \text{ or } m_2 \} \cup \{ m_1 m_0 m_2 m_0 \}$$

as the anti-chain of maximal positions (which by construction has at least one position each satisfying the shapes $\circ \cdots \bullet$ and $\bullet \cdots \circ$ respectively).

Now consider a second new l-tuple of games $\vec{A}''$ which is obtained from $\vec{A}$ by replacing A_j with $A_j[s]$ (note that $A_j[s]$ is obtained from A_j by performing only operation (2), but not (1)). So, we have

$$\begin{array}{rcllll} \vec{A} & \equiv & A_1, & \cdots & A_j, & \cdots & A_l \\ \vec{A}' & \equiv & A_1, & \cdots & A_j^-[s], & \cdots & A_l \\ \vec{A}'' & \equiv & A_1, & \cdots & A_j[s], & \cdots & A_l \end{array}$$

Note that both A_j as well as $A_j^-[s]$ have an obvious embedding (which is just inclusion) into $A_j[s]$: call them e_1 and e_2 respectively. Consider the obvious extensions of e_1 and e_2,

$$\begin{array}{rcccl} \vec{e_1} & : & \vec{A} & \rightarrowtail & \vec{A}'' \\ \vec{e_2} & : & \vec{A}' & \rightarrowtail & \vec{A}''. \end{array}$$

Applying the condition of uniformity to them in turn, we get the following equations:

$$\begin{array}{rcl} \sigma_{\vec{A}''} \circ \vec{e_1} & = & \vec{e_1} \circ \sigma_{\vec{A}} \\ \sigma_{\vec{A}''} \circ \vec{e_2} & = & \vec{e_2} \circ \sigma_{\vec{A}'}. \end{array}$$

Since $\vec{e_1}$ and $\vec{e_2}$ are just inclusions, we deduce that $\sigma_{\vec{A'}}$ behaves exactly as $\sigma_{\vec{A}}$ on positions $\langle\, i, m_1 \,\rangle$ and $\langle\, i, m_2 \,\rangle$ (note that the second new tuple $\vec{A''}$ is only introduced to facilitate this reasoning). More formally, writing $M_{\Gamma(\vec{A'})}$ as the disjoint union $M_{C_1} + \cdots + M_{C_n}$, we have

$$\sigma_{\vec{A'}} \quad : \quad \begin{cases} \langle\, i, m_1 \,\rangle & \mapsto \quad \langle\, i_1, m_1 \,\rangle \\ \langle\, i, m_2 \,\rangle & \mapsto \quad \langle\, i_2, m_2 \,\rangle, \end{cases}$$

Since $\vec{A'}$ now satisfies the condition (b1), We can now apply exactly the same argument as case (b1) to the l-tuple of games $\vec{A'}$. This concludes the proof of (c1).

Finally, it remains to show (c2).

Claim 3 ϕ is an involution.

Suppose for some i, with $C_i = C_{\phi(\phi(i))} = A_j$ and $C_{\phi(i)} = A_j^{\perp}$ but $\phi(\phi(i)) \neq i$. Write $i' \equiv \phi(i)$ and $i'' \equiv \phi(\phi(i))$. Suppose

$$\sigma_{\vec{A}} \quad : \quad \begin{cases} \langle\, i, m_1 \,\rangle & \mapsto \quad \langle\, i', m_1 \,\rangle, \\ \langle\, i', m_2 \,\rangle & \mapsto \quad \langle\, i'', m_2 \,\rangle. \end{cases}$$

Note that m_1 is an Opponent's move, and m_2 a Player's move in the game A_j. Consider the new l-tuple $\vec{A'}$ of games, which is obtained from the l-tuple $\vec{A}$ of games by replacing A_j with $A_j^-[s]$, where, as before $A_j^-[s]$ is obtained from A_j by performing the operations as before:

(1) first, prune away any maximal position in A_j which begins with either the move m_1 or m_2,

(2) then graft the new position $s \equiv m_1 m_2 m_1 m_2$ onto the pruned structure.

(3) In case the game A_j has no maximal position satisfying the shape $\bullet \cdots \circ$ other than that which begins with m_2 (and which has now been pruned away), graft a new maximal position $t \equiv m_0 m_1$ onto it where m_0, a fresh symbol is deemed to be an Opponent's move.

Call the new structure $A_j^-[s]$. It is a well-defined finite, fair game. For maximal generality, suppose it is necessary to perform operation (3).

Now, as before, consider a second new l-tuple of games $\vec{A''}$ which is obtained from $\vec{A}$ by replacing the game A_j with a new game $A_j^+[s]$ which has

- the set $M_{A_j} \cup \{\, m_0 \,\}$ as moves,
- and $F_{A_j} \cup \{\, s, t \,\}$ as maximal positions.

So, we have

$$\begin{array}{rcllll} \vec{A} & \equiv & A_1, & \cdots & A_j, & \cdots \;\; A_l \\ \vec{A'} & \equiv & A_1, & \cdots & A_j^-[s], & \cdots \;\; A_l \\ \vec{A''} & \equiv & A_1, & \cdots & A_j^+[s,t], & \cdots \;\; A_l \end{array}$$

Consider the following canonical embeddings,

$$\begin{array}{rcccl} e_1 & : & \vec{A} & \rightarrowtail & \vec{A''} \\ e_2 & : & \vec{A'} & \rightarrowtail & \vec{A''} \end{array}$$

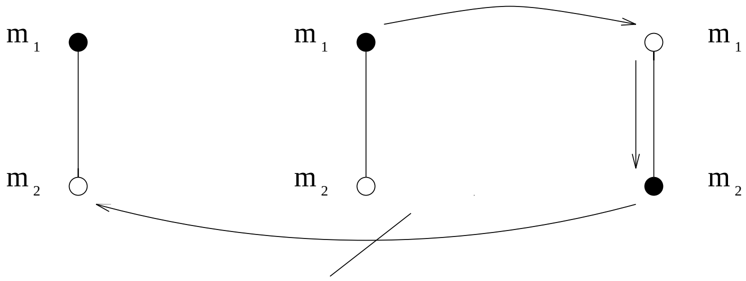

Figure 3:

and applying the condition of uniformity to them in turn as before, we deduce that $\sigma_{\vec{A}'}$ behaves exactly as $\sigma_{\vec{A}}$ on positions $\langle i, m_1 \rangle$ and $\langle i', m_2 \rangle$. More formally, writing $M_{\Gamma(\vec{A}')}$ as the disjoint union $M_{C_1} + \cdots + M_{C_n}$, we have

$$\sigma_{\vec{A}'} \quad : \quad \begin{cases} \langle i, m_1 \rangle & \mapsto \quad \langle i', m_1 \rangle \\ \langle i', m_2 \rangle & \mapsto \quad \langle i'', m_2 \rangle, \end{cases}$$

Now, consider the (partial) trace of a play which results from pitting $\sigma_{\vec{A}'}$ against a counter strategy τ:

$$u \equiv \langle i, m_1 \rangle \cdot \langle i', m_1 \rangle \cdot \langle i', m_2 \rangle;$$

or for a more vivid presentation, see Figure 3.

By Corollary 3.5 u is a valid position in the game $\Gamma(\vec{A}')$. However, at the position u, the "strategy" $\sigma_{\vec{A}'}$ is forced to respond with $\langle i, m_2 \rangle$. Since by construction $C_{i''}, = A_j^-[s]$ does not have any maximal position starting with the move m_2, the "strategy" $\sigma_{\vec{A}'}$ fails to be a strategy, and has to be rejected. □

Corollary 4.4 *For any sequent Γ, if the generic game Γ has a history-free, uniformly winning strategy σ, then*

(1) σ must be a copy-cat strategy,

(2) the copy-cat strategy σ induces a proof structure (Γ, ϕ_σ), and so, we may regard Γ as a binary sequent. □

Binary Sequents and Copy-Cat Pre-strategies A *binary sequent* is a sequent in which every propositional atom occurs precisely twice, once positively, and once negatively.

Any binary sequent $\Gamma(\alpha_1, \cdots, \alpha_l)$ defines a unique fixpoint-free involution over $\{1, \cdots, 2l\}$, call it ϕ. For any l-tuple of games $\vec{A}$, the map ϕ specifies a history-free "pre-strategy"; that is to say, it defines a map from Opponent's moves to Player's moves in exactly the same way as a copy-cat strategy: for any $1 \leqslant i \leqslant 2l$, writing $M_{\Gamma(\vec{A})} = M_{C_1} + \cdots + M_{C_{2l}}$, and for any $m \in M_{C_i}$ such that m is an Opponent's move in the game $\Gamma(\vec{A})$,

$$\langle i, m \rangle \quad \mapsto \quad \langle \phi(i), m \rangle.$$

However, this copy-cat pre-strategy does not necessarily determine a strategy for the game $\Gamma(\vec{A})$ because the copy-cat response by the Player may well be an illegal move at certain positions in the game $\Gamma(\vec{A})$. For example, consider the game $G_{\min} \otimes G_{\min}{}^{\perp}$. Suppose Opponent starts by making the move $\langle 1, o \rangle$. To copy Opponent's move, Player responds with the move $\langle 2, o \rangle$. This involves the Player changing from one component-game to another in a tensor game, which violates the switching convention. More generally, to show that a history-free pre-strategy σ is not a strategy, it suffices to pick a position sm in the game A at which Player is to move, such that smm is not a valid position.

Reduction to Semi-simple Sequents A *semi-simple* sequent $\Gamma \equiv T_1, \cdots, T_t$ is a sequent in which every formula T_i is built up from the literals by the tensor connective alone. Each such formula is called a $\otimes$-*cluster*. We shall exploit a key feature of $\otimes$-clusters: the associated DR-graph of a $\otimes$-cluster is always totally connected, regardless of the presence of any axiom link.

Let Γ be a binary sequent. Consider monotone MLL-contexts $C[\,]$, *i.e.* those with the "hole" appearing within the scope of tensors and pars but not linear negation. For such contexts, it is easy to check that whenever $A \multimap B$, then $C[A] \multimap C[B]$.

Proposition 4.5 *Given any binary sequent Γ, there are a finite number of binary, semi-simple sequents Γ_i with $1 \leqslant i \leqslant n$, satisfying the following conditions:*

(i) for each $1 \leqslant i \leqslant n$, $\vdash \Gamma \multimap \Gamma_i$,

(ii) $\vdash \Gamma$ if, and only if, for every $1 \leqslant i \leqslant n$, $\vdash \Gamma_i$.

Proof We first establish the following:

> **Claim** *Let $\Gamma = C[A \otimes (B \propto C)]$ be a binary sequent, and $\Gamma_1 = C[(A \otimes B) \propto C]$ and $\Gamma_2 = C[(A \otimes C) \propto B]$. Then*
>
> *(1) for each i, $\vdash \Gamma \multimap \Gamma_i$,*
>
> *(2) $\vdash \Gamma$ if, and only if, for every i, $\vdash \Gamma_i$.*

(1) follows immediately from following theorems of MLL, and the fact that $C[\,]$ is a monotonic context:

$$\begin{aligned} A \otimes (B \propto C) &\multimap (A \otimes B) \propto C, \\ A \otimes (B \propto C) &\multimap (A \otimes C) \propto B. \end{aligned}$$

For (2), we use the Soundness Condition for DR-graphs. Suppose Γ is not provable. Then, there are two cases:

- *either* there is a DR-switching S such that the DR-graph $\mathrm{Gr}(C[A \otimes (B \propto C)], S)$ has a cycle,
- *or* there is a DR-switching S such that the DR-graph $\mathrm{Gr}(C[A \otimes (B \propto C)], S)$ is disconnected.

Suppose the former. Referring to the figure below, if S sets the indicated par link to L (corresponding to the diagram on the left), there will be a cycle in Γ_1 with the par set to L (corresponding to the diagram on the right) as represented in Figure 4.

If S sets the par link to R (corresponding to the diagram on the left), there will be a cycle in Γ_2 with the par link set to L (see diagram on the right) as illustrated in Figure 5.

Suppose the latter, *i.e.* for some DR-switching S, $\mathrm{Gr}(C[A \otimes (B \propto C)], S)$ is disconnected. If S sets the indicated par link to L, then the associated DR-graph of Γ_1 will be disconnected; if S sets the par link to R, the associated DR-graph of Γ_2 will be disconnected.

The proposition then follows by repeated applications of the Claim, each time pushing a par connective to the top. $\square$

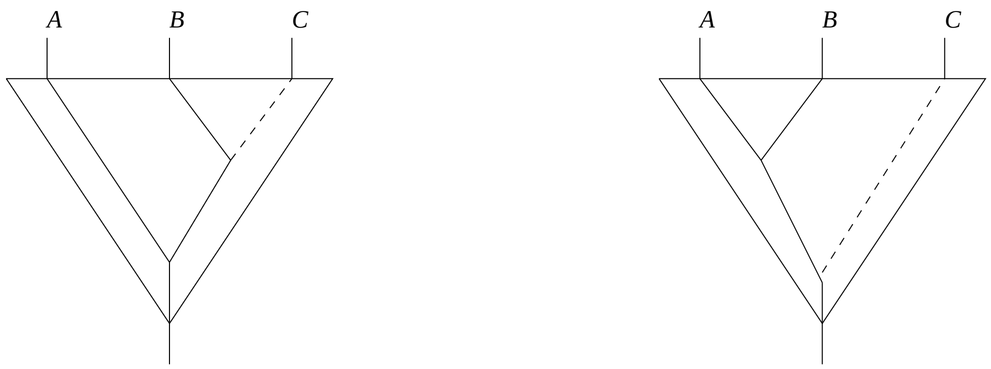

Figure 4:

Figure 5:

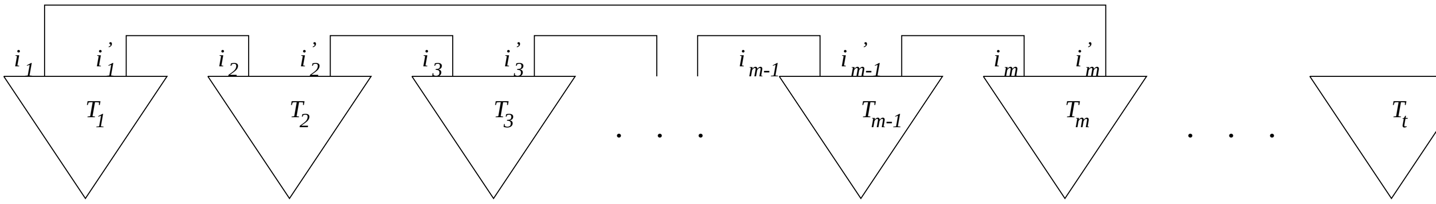

Figure 6: A cycle.

We have already observed that any binary sequent is equipped with a history-free copy-cat pre-strategy. When is this pre-strategy not a strategy? The answer is provided in the following proposition.

Proposition 4.6 *For any binary, semi-simple* MLL *sequent* Γ*, if the associated DR-graph is disconnected, or is cyclic for some switching* S*, then the generic game* Γ *does not have a history-free, uniformly winning strategy.*

Proof Let $\Gamma(\alpha_1, \cdots, \alpha_l)$ be a binary, semi-simple sequent, and let the binary linkages be specified by the fixpoint-free, involution ϕ over $\{1, \cdots, n\}$. Note that since Γ is a binary sequent, n has to be $2l$. We write $\Gamma(\vec{G_{\min}})$ for the game $\Gamma(G_{\min}, \cdots, G_{\min})$ where $G_{\min}$ is the least game with the set $F_{G_{\min}} = \{ op, po \}$ of maximal positions such that

$$\lambda_{G_{\min}} \quad : \quad \begin{cases} o \mapsto O \\ p \mapsto P. \end{cases}$$

The literals that occur in the game $\Gamma(\vec{G_{\min}})$ are represented as $C_1, \cdots, C_n$; so each C_i is either $G_{\min}$ or ${G_{\min}}^{\perp}$. We consider the cases of $\mathrm{Gr}(\Gamma, \phi)$ being cyclic, and disconnected in turn. In each case, we shall show that the copy-cat pre-strategy associated with the binary sequent ϕ is not a strategy. Since a history-free, uniformly winning strategy of any generic game is necessarily a copy-cat strategy (Copy-cat Proposition), we conclude that Γ does not have a history-free, uniformly winning strategy.

I Suppose $\mathrm{Gr}(\Gamma, \phi)$ is cyclic. Take any cycle, say specified by the following data:

$$\phi \quad : \quad \begin{cases} i'_1 & \mapsto & i_2 \\ i'_2 & \mapsto & i_3 \\ & \vdots & \\ i'_{m-1} & \mapsto & i_m \\ i'_m & \mapsto & i_1. \end{cases}$$

such that for each $1 \leqslant j \leqslant m$, the pair i_j and i'_j belong to the same $\otimes$-cluster T_j. We further assume that the m $\otimes$-clusters identified in this way are all distinct. There is no harm in assuming that $\Gamma(\vec{G_{\min}})$ is the sequent $T_1, \cdots, T_m, \cdots, T_t$ with the formulae arranged in the indicated order as in Figure 6.

Without loss of generality, we may suppose that the literals involved in the cycle have the following signs: for each $1 \leqslant j \leqslant m$, the literal occuring at i'_j is always positive, whereas the literal occuring at i_j is always negative.

Consider the counter strategy τ which switches the game $\langle\, i_j, G_{\min}{}^{\perp}\,\rangle$ to the game $\langle\, i'_j, G_{\min}\,\rangle$ in the following way: for each $1 \leqslant j \leqslant m$,

$$\tau : \langle\, i_j, o\,\rangle \mapsto \langle\, i'_j, o\,\rangle.$$

Note that since the two games $\langle\, i_j, G_{\min}{}^{\perp}\,\rangle$ and $\langle\, i'_j, G_{\min}\,\rangle$ belong to the same $\otimes$-cluster, such a move by Opponent is quite legal. Now, we pit the copy-cat pre-strategy induced by the proof structure (Γ, ϕ) against the counter strategy τ. Our intention is to show that the pre-strategy is not a strategy.

Suppose Opponent chooses to start the game with the move $\langle\, i'_1, o\,\rangle$, the following sequence may be reached:

$$s \equiv \langle\, i'_1, o\,\rangle \cdot \langle\, i_2, o\,\rangle \cdot \langle\, i'_2, o\,\rangle \cdot \langle\, i_3, o\,\rangle \cdot \cdots \cdot \langle\, i'_{m-1}, o\,\rangle \cdot \langle\, i_m, o\,\rangle \cdot \langle\, i'_m, o\,\rangle.$$

At this point, it is Player's turn to move. According to the copy-cat strategy, Player is to switch the game from $\langle\, i'_m, G_{\min}\,\rangle$ to $\langle\, i_1, G_{\min}{}^{\perp}\,\rangle$, and make the move $\langle\, i_1, o\,\rangle$. Now, the games $\langle\, i_1, G_{\min}{}^{\perp}\,\rangle$ and $\langle\, i'_1, G_{\min}\,\rangle$ are different games in the same $\otimes$-cluster. Such a move is illegal since by projecting the resultant sequence $s \cdot \langle\, i_1, o\,\rangle$ onto the $\otimes$-cluster T_1 which contains i_1 and i'_1, we see that it is a case of Player trying to switch game in a $\otimes$-cluster.

Finally, to complete the argument, we have to demonstrate that the sequence s of moves really is a position in the game $\Gamma(\vec{G_{\min}})$; that is to say, s may be extended to a maximal position in the game $\Gamma(\vec{G_{\min}})$. To do this, we classify the $\otimes$-clusters of Γ as follows:

(1) the $\otimes$-cluster T_1: only one move has been made in T_1, namely, $\langle\, i'_1, o\,\rangle$; so the points of "entry" and "exit" coincide.

(2) the $\otimes$-clusters $T_2, \cdots, T_m$: for each such cluster T_j, there is a point of "entry" $\langle\, i_j, o\,\rangle$, and a point of "exit" $\langle\, i'_j, o\,\rangle$;

(3) the $\otimes$-clusters $T_{m+1}, \cdots, T_t$: every such $\otimes$-cluster T_j is a "virgin" game; that is to say, up to the sequence s, no move has been made in T_j.

We check that each of the three conditions of Corollary 3.4(ii) are satisfied by s: the switching convention is followed; for each $1 \leqslant i \leqslant t$, the projection of s onto T_i is easily seen to be a valid position of T_i; and the third condition (a3) is vacuously satisfied. Hence, we apply the Corollary to conclude that s is indeed a valid position of the game $\Gamma(\vec{G_{\min}}) \equiv T_1 \propto \cdots \propto T_t$.

II Suppose the DR-graph $\mathrm{Gr}(\Gamma, \phi)$ is disconnected. Take a connected component (*i.e.* a maximally connected subgraph) of the graph, corresponding to say, the following group of $\otimes$-clusters $T_1, \cdots, T_m$; and we may assume that this component is acyclic, for otherwise, the argument for **I** suffices. We aim to show that the copy-cat pre-strategy defined by ϕ is not a strategy. To see this, consider a play which begins with an Opponent's move, say m_1 in T_1. At an arbitrary position of the unfolding play which ends with, say an Opponent's move m in the $\otimes$-cluster T_i, Player's response as prescribed by ϕ thought of as the copy-cat pre-strategy necessarily stays within the connected component $T_1, \cdots, T_m$, by assumption of maximal connectedness. If the arbitrary position ends with a Player's move m in the $\otimes$-cluster T_i, then Opponent's response has to be within the same cluster, according to the switching convention for the tensor game. Hence, no play resulting from piting the copy-cat pre-strategy against any counter-strategy can ever take the play out of the component $T_1, \cdots, T_m$. This means that there can be no resultant play which extends to a maximal fair position. Hence, the copy-cat pre-strategy is not a strategy. This concludes our argument. $\square$

Proof of the Full Completeness Theorem Suppose a given MLL sequent Γ regarded as a generic game has a history-free, uniformly winning strategy σ. By the Copy-Cat Proposition, and the Corollary which follows, we may regard Γ as a binary sequent to which Proposition 4.5 may be applied. Or equivalently, and more explicitly, there is a fixpoint-free involution ϕ_σ (the subscript σ emphasises that the involution is determined by the strategy σ) such that (Γ, ϕ_σ) is a proof structure. So, let $\{ \Gamma_1, \cdots, \Gamma_n \}$ be the set of binary, semi-simple sequents which characterise the proof structure (Γ, ϕ_σ) in the sense of Proposition 4.5, that is to say,

(1) for each i, $\vdash \Gamma \multimap \Gamma_i$,

(2) $\vdash \Gamma$ if, and only if, for each i, $\vdash \Gamma_i$.

Since the game semantics is sound, (1) and the assumption that Γ has a history-free, uniformly winning strategy together imply that for each i, the binary, semi-simple sequent Γ_i regarded as a generic game has a history-free, uniformly winning strategy. Now by Proposition 4.6, for each i, the associated DR-graph of each Γ_i is connected and acyclic. By Danos-Regnier Theorem, each proof structure Γ_i is a proof net. By (2), this implies that the proof structure (Γ, ϕ_σ) is in fact a proof net. $\square$

5 Further Directions

Extension to Infinite Games [Preliminary calculations suggest that the obvious extension of the definition works!]

[***Big gaps here...***]

Related Work

Old work: [Con76], [Joy77], [Hyl90], [LS91]

Lamarche: [Lam93]

Concrete Data Structure and Sequential Algorithms: [Cur93]

Acknowledgements

We would like to thank members of the Cambridge – Imperial Joint Seminars on Game Semantics for helpful discussions.

References

[AJ92] S. Abramsky and R. Jagadeesan. *Games and Full Completeness for Multiplicative Linear Logic.* Technical Report DoC 92/94, Imperial College, 1992. 32 pp. To appear in *J. Symb. Logic.*

[Bar79] M. Barr. *$*$-autonomous categories.* Springer-Verlag, 1979. Lecture Notes in Mathematics Vol. 752.

[Bar91] M. Barr. $*$-autonomous categories and linear logic. *Mathematical Structures in Computer Science*, 1:159–178, 1991.

[Bel93] G. Bellin. Proof nets for multiplicative and additive linear logic. 1993. unpublished manuscript.

[Bla92] A. Blass. A game semantics for linear logic. *Annals of Pure and Applied Logic*, 56:183–220, 1992.

[Con76] J. H. Conway. *On Numbers and Games.* Academic Press, 1976. Vol. 6 of London Mathematical Society Monograph.

[Cur93] P.-L. Curien. *Categorical Combinators, Sequential Algorithms, and Functional Programming.* Birkhäuser, second edition, 1993. Progress in Theoretical Computer Science Series.

[Dev80] K. Devlin. *Fundamentals of Contemporary Set Theory.* Springer-Verlag, 1980. Universitext.

[DR89] V. Danos and L. Regnier. The structures of multiplicatives. *Archive for Mathematical Logic*, 28:181–203, 1989.

[Fra86] N. Francez. *Fairness.* Springer, 1986.

[Gir87] J.-Y. Girard. Linear logic. *Theoretical Computer Science*, 50:1 – 102, 1987.

[Hut93] M. Huth. On fair games and parallelism. 1993. unpublished draft.

[Hyl90] J. M. E. Hyland. Conway games and linear logic. 1990. unpublished lecture.

[Joy77] A. Joyal. Remarques sur la theorie des jeux a deux personnes. *Gazette des Sciences Mathematiques du Quebec*, 1(4), 1977.

[Lam93] F. Lamarche. On the proof net problem for additives in linear logic. 1993. unpublished lecture.

[LS91] Y. Lafont and T. Streicher. Games semantics for linear logic. In *Proc. Sixth Annual IEEE Symp. Logic in Computer Science*, pages 42–51, Computer Society Press, 1991.

[Par81] D. M. Park. Concurrency on automata and infinite sequences. In P. Deussen, editor, *Conference on Theoretical Computer Science*, Springer-Verlag, Berlin, 1981. Lecture Notes in Computer Science Vol. 104.

[See89] R. Seely. $*$-autonomous categories, cofree algebras and linear logic. In J. W. Gray and A. Scedrov, editors, *Categories in Computer Science and Logic*, pages 371–382, American Mathematical Society, 1989. Vol. 92 of *Contemporary Mathematics.*

[Tro92] A. S. Troelstra. *Lectures on Linear Logic.* Centre for the Study of Language and Information, Stanford, 1992. CSLI Lecture Notes No. 29.