The true concurrency of Herbrand's theorem

Aurore Alcolei[1], Pierre Clairambault[1], Martin Hyland[2], and Glynn Winskel[3]

[1]Univ Lyon, CNRS, ENS de Lyon, UCB Lyon 1, LIP, Lyon, France
[2]Department of Pure Mathematics and Mathematical Statistics, University of Cambridge, Cambridge, U.K.
[3]Computer Laboratory, University of Cambridge, Cambridge, U.K.

***Abstract*—Herbrand's theorem, widely regarded as a cornerstone of proof theory, exposes some of the constructive content of classical logic. In its simplest form, it reduces the validity of a first-order purely existential formula to that of a finite disjunction. More generally, it gives a reduction of first-order validity to propositional validity, by understanding the structure of the assignment of first-order terms to existential quantifiers, and the causal dependency between quantifiers.**

In this paper, we show that Herbrand's theorem in its general form can be elegantly stated as a theorem in the framework of concurrent games. The causal structure of concurrent strategies, paired with annotations by first-order terms, is used to specify the dependency between quantifiers. Furthermore concurrent strategies can be composed, yielding a compositional proof of Herbrand's theorem, simply by interpreting classical sequent proofs in a well-chosen denotational model.

I. Introduction

"What more do we know when we have proved a theorem
by restricted means than if we merely know it is true?"

Kreisel's question is the driving force for much modern Proof Theory. This paper is concerned with Herbrand's Theorem, perhaps the earliest result in this direction. It is a simple consequence of completeness and compactness. So it is an example of information being extracted from the bare fact of provability. Usually by contrast one thinks in terms of extracting information from the proofs themselves, typically - as in Kohlenbach's proof mining - via some form of functional interpretation. This has the advantage that information is extracted compositionally in the spirit of functional programming. Specifically information for $\vdash A$ and $\vdash A \to B$ can be composed to give information for $\vdash B$; or in terms of the sequent calculus we can interpret the cut rule.

It seems to be folklore that there is a problem for Herbrand's Theorem. That is made precise in Kohlenbach [16] which shows that one cannot hope directly to use collections of Herbrand terms for $\vdash A$ and $\vdash A \to B$ to give a collection for $\vdash B$. That leaves the possibility of making some richer data compositional; and that possibility is realised indirectly in Gerhardy and Kohlenbach [10] with data provided by Shoenfield's version [25] of Gödel's Dialectica Interpretation [13]. Now functional interpretations make no pretence to be faithful to the structure of proofs as encapsulated in systems like the sequent calculus and it is compelling to seek some compositional form of Herbrand's Theorem arising directly from proofs. We present such a version in this paper: it is based on concurrent strategies in games as for example in [23], [4]. We stress that the essential point is the composition of strategies. Were we interested only in cut-free sequent calculus our strategies would essentially be Miller's expansion trees [21] enriched with explicit acyclicity witnesses.

Beyond the term information, our strategies aim to represent transparently the causal dependency between quantifiers implicitly carried by sequent proofs, avoiding the excess sequentialization caused by a negative translation. As a consequence, some phenomena known from the proof theory of classical logic reflect in our model: our interpretation does not preserve cut elimination (lest the model collapse to a boolean algebra [12]) – although we will see that cut elimination is preserved in a sense for *first-order MLL* [11]. Likewise, just as classical proofs can lead to arbitrary large cut-free proofs [8], our interpretation may yield *infinite* winning strategies, from which *finite* sub-strategies can nonetheless always be extracted. This reflects the fact that symmetric proof systems for classical logic are in general weakly, rather than strongly, normalizing.

Related work: The literature is rich in related work. Generalizations of Miller's expansion trees supporting cuts include Heijltjes' *proof forests* [14], McKinley's *Herbrand nets* [19], and Hetzl and Weller's more recent *expansion trees with cuts* [15]. In all three cases, a generalization of expansion trees allowing cuts is given along with a weakly normalizing cut reduction procedure. Intuitions from games are often mentioned, but the methods used are syntactic.

On the game-theoretic front, our model is closely related to Laurent's model for the first-order $\lambda\mu$-calculus [17], from which we differ by treating a symmetric proof system with an involutive negation, avoiding sequentiality. Also related is Mimram's categorical construction of a games model for a linear first-order logic without propositional connectives [22].

Outline: In Section II we recall Herbrand's theorem, and introduce the game-theoretic language leading to our compositional reformulation of the theorem. The rest of the paper describes the interpretation of proofs as strategies: in Section III we give the interpretation of propositional MLL, in Section IV we deal with quantifiers, and finally, in Section V, we add contraction and weakening and complete the interpretation.

II. From Herbrand to winning Σ-strategies

A **signature** is a pair $\Sigma = (\Sigma_f, \Sigma_p)$, with Σ_f a countable set of **function symbols** ($\mathsf{f}, \mathsf{g}, \mathsf{h}, etc.$ range over function symbols), and Σ_p a countable set of **predicate symbols** ($\mathsf{P}, \mathsf{Q}, etc.$ range over predicate symbols). There is an **arity**

function $\mathrm{ar} : \Sigma_f \uplus \Sigma_p \to \mathbb{N}$ where $\uplus$ is the usual set-theoretic union, where the argument sets are disjoint. For a relative gain in simplicity in some arguments and examples, we assume that Σ has at least one constant symbol, *i.e.* a function symbol of arity 0. We use $\mathsf{a}, \mathsf{b}, \mathsf{c}, \dots$ to range over constant symbols.

If $\mathcal{V}$ is a set of **variable names**, we write $\mathsf{Tm}_\Sigma(\mathcal{V})$ for the set of first-order terms on Σ with free variables in $\mathcal{V}$. We use variables $t, s, u, v, \dots$ to range over terms. **Atomic formulas** have the form $\mathsf{P}(t_1, \dots, t_n)$ or $\neg\mathsf{P}(t_1, \dots, t_n)$, where P is a n-ary predicate symbol and the t_is are terms. **Formulas** are also closed under quantifiers, and the connectives $\vee$ and $\wedge$. **Negation** is not considered a logical connective: the negation $\varphi^\perp$ of φ is obtained by De Morgan rules. We write $\mathsf{Form}_\Sigma(\mathcal{V})$ for the set of **first-order formulas** on Σ with free variables in $\mathcal{V}$, and use $\varphi, \psi, \dots$ to range over them. We also write $\mathsf{QF}_\Sigma(\mathcal{V})$ for the set of **quantifier-free** formulas. Finally, we write $\mathrm{fv}(\varphi)$ or $\mathrm{fv}(t)$ for the set of free variables in a formula φ or a term t. Formulas are considered up to α-conversion and assumed to satisfy Barendregt's convention.

A. Herbrand's theorem

Intuitionistic logic has the *witness* property: if $\exists \mathsf{x}\, \varphi$ holds intuitionistically, then there is *a single* term t such that $\varphi(t)$ holds. While this fails in classical logic, Herbrand's theorem, in its most popular form, gives a weakened classical version:

Theorem II.1. *Let $\mathscr{T}$ be a theory finitely axiomatized by universal formulas. Consider a formula of the form $\psi = \exists \mathsf{x}_1 \dots \exists \mathsf{x}_n \varphi(\mathsf{x}_1, \dots, \mathsf{x}_n)$ where $\varphi \in \mathsf{QF}_\Sigma$. Then, $\mathscr{T} \vDash \psi$ iff there are closed terms $(t_{i,j})_{1 \le i \le p, 1 \le j \le n}$ such that*

$$\mathscr{T} \vDash \bigvee_{i=1}^{p} \varphi(t_{i,1}, \dots, t_{i,n})$$

The single witness is replaced with a finite disjunction.

Example II.2. *Consider the formula $\psi = \exists \mathsf{x} \neg\mathsf{P}(\mathsf{x}) \vee \mathsf{P}(\mathsf{f}(\mathsf{x}))$ (where $\mathsf{f} \in \Sigma_f$). A valid Herbrand disjunction for ψ is*

$$(\neg\mathsf{P}(c) \vee \mathsf{P}(\mathsf{f}(\mathsf{c}))) \vee (\neg\mathsf{P}(\mathsf{f}(\mathsf{c})) \vee \mathsf{P}(\mathsf{f}(\mathsf{f}(\mathsf{c}))))$$

where c is some constant symbol.

Such a disjunction result can also be given for general formulas. A common way to do so is by reduction to the above: a formula φ is converted to prenex normal form and universally quantified variables are replaced with new function symbols added to Σ, in a process called *Herbrandization* (dual to Skolemization). For instance, the *drinker's formula*:

$$\exists \mathsf{x} \forall \mathsf{y} \neg\mathsf{P}(\mathsf{x}) \vee \mathsf{P}(\mathsf{y}) \qquad (DF)$$

yields by Herbrandization the formula ψ of Example II.2.

But Herbrand's original theorem did not use Herbrandization, although it held for general formulas. We recall here Buss' formulation [3], which is different from Herbrand's but similar in spirit. Its definition requires some machinery.

Definition II.3. *A $\vee$-**expansion** of φ is obtained by replacing a subformula ψ of φ with $\psi \vee \psi$, where the outermost connective of ψ is an existential quantifier. A **hereditary** $\vee$-**expansion** of φ is a formula obtained by a finite number of $\vee$-expansions.*

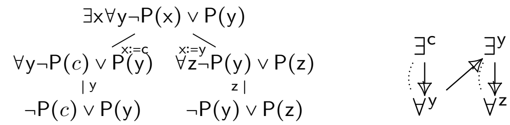

Fig. 1. An expansion tree and winning Σ-strategy for *DF*

The general version of Herbrand's theorem is stated in terms of prenex normal forms of hereditary $\vee$-expansions.

Definition II.4. *Let φ be in prenex normal form.* W.l.o.g.*, φ has the following form, with ψ quantifier-free:*

$$\forall \mathsf{x}_1 \dots \mathsf{x}_{n_1} \exists \mathsf{y}_1 \forall \mathsf{x}_{n_1+1} \dots \mathsf{x}_{n_2} \exists \mathsf{y}_2 \dots \exists \mathsf{y}_p \forall \mathsf{x}_{n_p+1} \dots \mathsf{x}_{n_{p+1}} \psi(\mathsf{x}_i, \mathsf{y}_j)$$

*with $0 \le n_1 \le n_2 \le \dots \ne n_{p+1}$. A **witness** for φ is given by terms $t_i \in \mathsf{Tm}_\Sigma(\{\mathsf{x}_1, \dots, \mathsf{x}_{\mathsf{n}_\mathsf{i}}\})$ such that $\vDash \psi(\mathsf{x}_i, t_j)$.*

A **Herbrand proof** of φ is a *witness* for a *prenexification* of a *hereditary $\vee$-expansion* of φ; which by definition of witness yields a propositional tautology. Herbrand's theorem states:

Theorem II.5. *For any φ, $\vDash \varphi$ iff φ has a Herbrand proof.*

For *DF*, the prenexified hereditary $\vee$-expansion

$$\exists \mathsf{x}_1 \forall \mathsf{y}_1 \exists \mathsf{x}_2 \forall \mathsf{y}_2 (\neg\mathsf{P}(\mathsf{x}_1) \vee \mathsf{P}(\mathsf{y}_1)) \vee (\neg\mathsf{P}(\mathsf{x}_2) \vee \mathsf{P}(\mathsf{y}_2))$$

along with the witness $\mathsf{x}_1 := c, \mathsf{x}_2 := \mathsf{y}_1$ form a Herbrand proof.

B. Trees, nets and games for Herbrand's theorem

Herbrand proofs decant the propositional content of a proof, focusing on quantifiers. But their definition is rather indirect; they are very removed from the original formula and are not easily composed. Miller proposes [21] to represent them more geometrically as **expansion trees**.

Expansion trees can be introduced through a game-theoretic metaphor, reminiscent of Coquand's game semantics for classical arithmetic [7]. Two players, ∃loïse and ∀bélard, argue about the validity of a formula. On a formula $\forall \mathsf{x} \varphi$, ∀bélard provides a fresh variable x and we keep playing on φ. On a formula $\exists \mathsf{x} \varphi$, ∃loïse provides a *term* t, possibly containing variables previously introduced by ∀bélard. ∃loïse, though, has a special power: at any time she can *backtrack* to a previous existential position, and propose a new term. Figure 1 (left) shows an expansion tree for *DF*. It may be read from top to bottom, and from left to right: ∃loïse plays c, then ∀bélard introduces x, then ∃loïse *backtracks* (we jump to the right branch) and plays x, and ∀bélard introduces y. It is a win for ∃loïse: the disjunction of the leaves is a tautology.

However, the metaphor has limits: the order between two branches of an expansion tree is not part of the structure, but implicit in the term annotations. In this paper, our chosen representations for proofs (called Σ-*strategies*) will make this causality explicit as a partial order: we also show in Figure 1 the expansion tree represented as Σ-strategy.

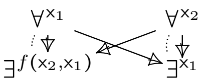

Fig. 2. A partially ordered winning Σ-strategy

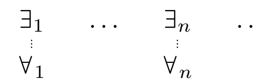

Fig. 3. The arena $[\![DF]\!]^{\exists}$

This causal order (imposed by terms) is not always sequential: we display in Figure 2 what will be a *winning Σ-strategy* on formula $(\forall x_1 \exists y_1 P(x_1, y_1)) \vee (\forall x_2 \exists y_2 \neg P(y_2, f(x_2, y_2)))$ – it is, again, more informative than the corresponding expansion tree where the crossed dependencies would only be implicit.

Although not explicit, these dependencies play an important role in expansion trees. Certainly the tree below should not be valid as the formula it plays on is invalid.

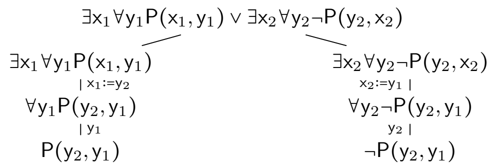

And indeed, the full definition of expansion trees involves an acyclicity correctness criterion that forbids this tree – likewise, it will be impossible to write it as a Σ-strategy as the fact that Σ-strategies are explicitly partially ordered enforces acyclicity.

C. Expansion trees as winning Σ-strategies

We now give our formulation of expansion trees as Σ-*strategies*. Although our definitions look superficially very different from Miller's, the only fundamental difference is the explicit display of the dependency between quantifiers.

Σ-strategies will be certain partial orders, with elements either "$\forall$ events" or "$\exists$ events". Events will carry terms, in a way that respects causal dependency. They will play on *games* representing the formulas. The first component of a game is its *arena*, that specifies the causal ordering between quantifiers.

Definition II.6. *An* ***arena*** *is* $A = (|A|, \leq_A, \mathrm{pol}_A)$ *where* $|A|$ *is a set of* ***events****,* $\leq_A$ *is a partial order that is* forest-shaped*:* (1) *if* $a_1 \leq_A a$ *and* $a_2 \leq_A a$*, then either* $a_1 \leq_A a_2$ *or* $a_2 \leq_A a_1$*, and* (2) *for all* $a \in |A|$*, the branch* $[a]_A = \{a' \in A \mid a' \leq_A a\}$ *is finite. Finally,* $\mathrm{pol}_A : |A| \to \{\forall, \exists\}$ *is a* ***polarity function*** *which expresses if a move belongs to* $\exists$*loïse or* $\forall$*bélard.*

A **configuration** of an arena (or any partial order) is a down-closed set of events. We write $\mathscr{C}^\infty(A)$ for the set of configurations of A, and $\mathscr{C}(A)$ for the set of *finite* configurations.

The arena only describes the moves available to both players; it says nothing about terms or winning. Similarly to expansion trees where only $\exists$loïse can replicate her moves ("*backtrack*", although the terminology is imperfect when strategies are not sequential), our arenas will at first be biased towards $\exists$loïse: each $\exists$ move exists in as many copies as she might desire, whereas $\forall$ events are *a priori* not copied. Figure 3 shows the $\exists$-biased arena $[\![DF]\!]^{\exists}$ for *DF*. The order is drawn from top to bottom, *i.e.* events at the top are minimal. Although only $\exists$loïse can replicate her moves, the universal quantifier is also copied as it depends on the existential quantifier.

Strategies on arena A will be certain *augmentations* of prefixes of A. They carry causal dependency between quantifiers induced by term annotations, but not the terms themselves.

We introduce the notation $\rightarrowtriangle$, already used implicitly in Figures 1 and 2. For A any partial order and $a_1, a_2 \in |A|$, we write $a_1 \rightarrowtriangle_A a_2$ (or $a_1 \rightarrowtriangle a_2$ if A is clear from the context) if $a_1 <_A a_2$ with no other event in between, *i.e.* for any $a \in |A|$ such that $a_1 \leq_A a \leq_A a_2$, then $a_1 = a$ or $a_2 = a$. We call $\rightarrowtriangle$ **immediate causal dependency** in line with event structures where the partial order is that of causal dependency.

Definition II.7. *A* ***strategy*** σ *on arena* A*, written* $\sigma : A$*, is a partial order* $(|\sigma|, \leq_\sigma)$ *with* $|\sigma| \subseteq |A|$*, such that for all* $a \in |\sigma|$*,* $[a]_\sigma$ *is finite (an* elementary event structure*); subject to:*

(1) Arena-respecting. *We have* $\mathscr{C}^\infty(\sigma) \subseteq \mathscr{C}^\infty(A)$*,*
(2) Receptivity. *If* $x \in \mathscr{C}(\sigma)$ *such that* $x \cup \{a^\forall\} \in \mathscr{C}(A)$*, then* $a \in |\sigma|$ *as well (*$a^\forall$ *means that* $\mathrm{pol}_A(a) = \forall$*).*
(3) Courtesy. *If* $a_1 \rightarrowtriangle_\sigma a_2$*, then either* $a_1 \rightarrowtriangle_A a_2$*, or* $\mathrm{pol}_A(a_1) = \forall$ *and* $\mathrm{pol}(a_2) = \exists$*.*

This is a simplification of Rideau and Winskel's *concurrent strategies* [23] permitted by the purely deterministic setting; also equivalent [23] to Melliès and Mimram's earlier *receptive ingenuous strategies* [20] – though the direct handle on the causal order in the definition above is convenient for our purposes. Receptivity means that $\exists$loïse cannot refuse to acknowledge a move by $\forall$bélard, and courtesy that the only new causal constraints that she can enforce with respect to the game is that some existential quantifiers depend on some universal quantifiers. Ignoring terms, Figure 1 (on the right) displays a strategy on the arena of Figure 3 – in Figure 1 we also display via dotted lines the immediate dependency of the arena.

We can now add terms, and define Σ-*strategies*.

Definition II.8. *A* Σ***-strategy*** *on arena* A *is a strategy* $\sigma : A$*, with a* ***labelling function*** $\lambda_\sigma : |\sigma| \to \mathsf{Tm}_\Sigma(|\sigma|)$*, such that:*

$$\begin{array}{rcl} \forall a^\forall \in |\sigma|, \quad \lambda_\sigma(a) & = & a \\ \forall a^\exists \in |\sigma|, \quad \lambda_\sigma(a) & \in & \mathsf{Tm}_\Sigma([a]_\sigma^\forall) \end{array}$$

where $[a]_\sigma^\forall = \{a' \in |\sigma| \mid a' \leq_\sigma a \ \&\ \mathrm{pol}_A(a') = \forall\}$.

Rather than having $\forall$ moves introduce fresh variables, we find it convenient to consider them as *variables themselves*. Hence, the $\exists$ moves are annotated by terms having as free variables the $\forall$ moves in their causal history. For instance the diagram at the right of Figure 1 is meant formally to denote the one on the right (where superscripts are the terms given by λ). In the sequel we omit the (redundant) annotation of $\forall$bélard's events.

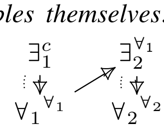

Σ-strategies are more general than expansion trees (besides the fact that they are not assumed finite): they have an explicit causal ordering, which may be more constraining than that given by the terms. A Σ-strategy $\sigma : A$ is **minimal** iff whenever $a_1 \rightarrowtriangle_\sigma a_2$ such that $a_1 \notin \mathrm{fv}(\lambda_\sigma(a_2))$, then $a_1 \rightarrowtriangle_A a_2$ as well. In a minimal Σ-strategy $\sigma : A$, the ordering $\leq_\sigma$ is actually redundant and can be uniquely recovered from λ_σ and $\leq_A$.

Now, we adjoin *winning conditions* to arenas and define *winning Σ-strategies*. As in expansion trees or Herbrand proofs, these amount to the substitution of the expansion of the original formula being a tautology.

Definition II.9. *A **game** $\mathcal{A}$ is an arena A together with **winning conditions**, given as a function:*

$$\mathcal{W}_{\mathcal{A}} : (x \in \mathscr{C}^\infty(A)) \to \mathsf{QF}^\infty_\Sigma(x)$$

*where $\mathsf{QF}^\infty_\Sigma(x)$ is the set of **infinitary quantifier-free formulas** – obtained from $\mathsf{QF}_\Sigma(x)$ by adding infinitary connectives $\bigvee_{i\in I} \varphi_i$ and $\bigwedge_{i\in I}\varphi_i$, where I is some countable set.*

For a game interpreting φ, the winning condition associates configurations of the arena $[\![\varphi]\!]$ with (essentially) the propositional part of matching hereditary $\vee$-expansions. For instance:

$$\begin{array}{rcl}\mathcal{W}_{[\![DF]\!]^\exists}(\{\exists_3, \forall_3, \exists_6, \forall_6\}) &=& (\neg \mathsf{P}(\exists_3) \vee \mathsf{P}(\forall_3)) \vee \\ && (\neg \mathsf{P}(\exists_6) \vee \mathsf{P}(\forall_6)) \\ \mathcal{W}_{[\![DF]\!]^\exists}(\{\exists_3, \forall_3, \exists_6\}) &=& (\neg \mathsf{P}(\exists_3) \vee \mathsf{P}(\forall_3)) \vee \top\end{array}$$

recalling that the arena for *DF* appears in Figure 3. The $\top$ (the true formula) on the second line is due to $\forall$bélard not having played $\forall_6$ yet, yielding victory to $\exists$loïse on that copy. The winning conditions yield syntactic, uninterpreted formulas: we keep the second formula as-is although it is equivalent to $\top$.

Finally, we can define winning strategies.

Definition II.10. *If $\sigma : A$ is a Σ-strategy and $x \in \mathscr{C}^\infty(\sigma)$, we say that x is **tautological** in σ if the formula*

$$\mathcal{W}_{\mathcal{A}}(x)[\lambda_\sigma]$$

corresponding to the substitution of $\mathcal{W}_{\mathcal{A}}(x) \in \mathsf{QF}^\infty_\Sigma(x)$ by $\lambda_\sigma : x \to \mathsf{Tm}_\Sigma(x)$, is a (possibly infinite) tautology.

*A Σ-strategy $\sigma : A$ is **winning** if any $x \in \mathscr{C}^\infty(\sigma)$ that is $\exists$-**maximal** (*i.e.* $x \in \mathscr{C}^\infty(\sigma)$ such that for all $a \in |\sigma|$ with $x \cup \{a\} \in \mathscr{C}^\infty(\sigma)$, $\mathrm{pol}_A(a) = \forall$) is tautological. A Σ-strategy $\sigma : A$ is **top-winning** if $|\sigma| \in \mathscr{C}^\infty(\sigma)$ is tautological.*

D. Constructions on games and Herbrand's theorem

To complete our statement of Herbrand's theorem, it remains to define the interpretation of formulas as games.

a) Arenas: First, we define some operations on arenas. We write $\varnothing$ for the **empty arena**, with no events. If A is an arena, we write $A^\perp$ for the **dual** arena, with the same events and causality but the polarity reversed, *i.e.* $\mathrm{pol}_{A^\perp}(a) = \forall$ iff $\mathrm{pol}_A(a) = \exists$. We review some other important constructions.

Definition II.11. *The **simple parallel composition** $A_1 \parallel A_2$ of A_1 and A_2 has as events the tagged disjoint union $\{1\} \times |A_1| \uplus \{2\} \times |A_2|$, causal order given by $(i, a) \leq_{A_1 \parallel A_2} (j, a')$ iff $i = j$ and $a \leq_{A_i} a'$. Polarity is $\mathrm{pol}_{A_1\parallel A_2}((i,a)) = \mathrm{pol}_{A_i}(a)$.*

Configurations $x \in \mathscr{C}^\infty(A \parallel B)$ have the form $\{1\} \times x_A \cup \{2\} \times x_B$ with $x_A \in \mathscr{C}^\infty(A)$ and $x_B \in \mathscr{C}^\infty(B)$, which we write $x = x_A \parallel x_B$. Binary simple parallel composition has a general counterpart $\parallel_{i\in I} A_i$ with I at most countable, defined likewise. We will use the uniform countably infinite simple parallel composition $\parallel_\omega A$ with ω parallel copies of A.

Another important arena construction is *prefixing*.

Definition II.12. *For $\alpha \in \{\forall, \exists\}$ and A an arena, the **prefixed arena** $\alpha.A$ has events $\{(1,\alpha)\} \cup \{2\} \times |A|$ and causality $(i,a) \leq (j, a')$ iff $i = j = 2$ and $a \leq_A a'$, or $(i,a) = (1,\alpha)$; meaning that $(1,\alpha)$ is the unique minimal event in $\alpha.A$. Its polarity is $\mathrm{pol}_{\alpha.A}((1,\alpha)) = \alpha$ and $\mathrm{pol}_{\alpha.A}((2,a)) = \mathrm{pol}_A(a)$.*

Configurations $x \in \mathscr{C}^\infty(\alpha.A)$ are either empty, or of the form $\{(1,\alpha)\} \cup \{2\} \times x_A$ with $x_A \in \mathscr{C}^\infty(A)$, written $\alpha.x_A$.

b) Winning: To give the inductive interpretation of formulas we have to consider formulas that are not closed, *i.e.* with free variables. For $\mathcal{V}$ a finite set, a **$\mathcal{V}$-game** is defined as a game $\mathcal{A}$ in Definition II.9, except that we have, for $x \in \mathscr{C}^\infty(A)$,

$$\mathcal{W}_{\mathcal{A}}(x) \in \mathsf{QF}^\infty_{\Sigma \uplus \mathcal{V}}(x)\,.$$

We now define all our constructions, on $\mathcal{V}$-games rather than on games. The duality operation on arenas $(-)^\perp$ extends to $\mathcal{V}$-games, simply by negating the winning conditions: for all $x \in \mathscr{C}^\infty(A)$, $\mathcal{W}_{\mathcal{A}^\perp}(x) = \mathcal{W}_{\mathcal{A}}(x)^\perp$. The $\parallel$ of arenas gives rise to *two* constructions on $\mathcal{V}$-games:

Definition II.13. *Let $\mathcal{A}$ and $\mathcal{B}$ be $\mathcal{V}$-games. We define two $\mathcal{V}$-games on arena $A \parallel B$, differing by the winning condition:*

$$\begin{array}{rcl}\mathcal{W}_{\mathcal{A}\otimes\mathcal{B}}(x_A \parallel x_B) &=& \mathcal{W}_{\mathcal{A}}(x_A) \wedge \mathcal{W}_{\mathcal{B}}(x_B) \\ \mathcal{W}_{\mathcal{A}\parr\mathcal{B}}(x_A \parallel x_B) &=& \mathcal{W}_{\mathcal{A}}(x_A) \vee \mathcal{W}_{\mathcal{B}}(x_B)\end{array}$$

Note the implicit renaming so that $\mathcal{W}_{\mathcal{A}}(x_A), \mathcal{W}_{\mathcal{B}}(x_B)$ are in $\mathsf{QF}^\infty_{\Sigma\uplus\mathcal{V}}(x_A \parallel x_B)$ rather than $\mathsf{QF}^\infty_{\Sigma\uplus\mathcal{V}}(x_A), \mathsf{QF}^\infty_{\Sigma\uplus\mathcal{V}}(x_B)$ respectively – here and in the sequel, we will keep such renamings implicit when we believe it helps readability.

Note that $\mathcal{A}$ and $\mathcal{B}$ are De Morgan duals, *i.e.* $(\mathcal{A} \otimes \mathcal{B})^\perp = \mathcal{A}^\perp \parr \mathcal{B}^\perp$. The reader may wonder why these operations are written $\otimes$ and $\parr$ rather than $\wedge$ and $\vee$. This is because, as we will see, these operations by themselves behave more like the connectives of linear logic [11] than those of classical logic; for each $\mathcal{V}$ the $\otimes$ and $\parr$ will form the basis of a $\star$-autonomous structure and hence a model of multiplicative linear logic.

To recover classical logic, we will add *replication* to the interpretation of formulas.

Definition II.14. *Let $\mathcal{A}$ be a $\mathcal{V}$-game. We define two new $\mathcal{V}$-games $!\mathcal{A}$ and $?\mathcal{A}$ with arena $\parallel_\omega A$, and winning conditions:*

$$\begin{array}{rcl}\mathcal{W}_{!\mathcal{A}}(\parallel_{i\in\omega} x_i) &=& \bigwedge_{i\in\omega} \mathcal{W}_{\mathcal{A}}(x_i) \\ \mathcal{W}_{?\mathcal{A}}(\parallel_{i\in\omega} x_i) &=& \bigvee_{i\in\omega} \mathcal{W}_{\mathcal{A}}(x_i)\end{array}$$

Although $\mathcal{W}_{!\mathcal{A}}(x)$ (resp. $\mathcal{W}_{?\mathcal{A}}(x)$) is, syntactically, an infinite conjunction (resp. disjunction), we always implicitly simplify it to a finite one when x visits finitely many copies (as we then have infinitely many occurrences of $\mathcal{W}_{\mathcal{A}}(\varnothing)$).

Next we show how $\mathcal{V}$-games support quantifiers.

$$
\begin{array}{rclrcl}
[\![\top]\!]^{\exists}_{\mathcal{V}} & = & 1 & [\![\mathsf{P}(t_1,\dots,t_n)]\!]^{\exists}_{\mathcal{V}} & = & \mathsf{P}(t_1,\dots,t_n)\\
[\![\bot]\!]^{\exists}_{\mathcal{V}} & = & \bot & [\![\neg\mathsf{P}(t_1,\dots,t_n)]\!]^{\exists}_{\mathcal{V}} & = & \neg\mathsf{P}(t_1,\dots,t_n)\\
[\![\exists \mathsf{x}\varphi]\!]^{\exists}_{\mathcal{V}} & = & ?\exists\mathsf{x}.[\![\varphi]\!]_{\mathcal{V}\uplus\{\mathsf{x}\}} & [\![\varphi_1\vee\varphi_2]\!]^{\exists}_{\mathcal{V}} & = & [\![\varphi_1]\!]^{\exists}_{\mathcal{V}} \parr [\![\varphi_2]\!]^{\exists}_{\mathcal{V}}\\
[\![\forall \mathsf{x}\varphi]\!]^{\exists}_{\mathcal{V}} & = & \forall\mathsf{x}.[\![\varphi]\!]_{\mathcal{V}\uplus\{\mathsf{x}\}} & [\![\varphi_1\wedge\varphi_2]\!]^{\exists}_{\mathcal{V}} & = & [\![\varphi_1]\!]^{\exists}_{\mathcal{V}} \otimes [\![\varphi_2]\!]^{\exists}_{\mathcal{V}}
\end{array}
$$

Fig. 4. $\exists$-biased interpretation of formulas

Definition II.15. *For $\mathcal{A}$ a $(\mathcal{V}\uplus\{\mathsf{x}\})$-game, the $\mathcal{V}$-game $\forall\mathsf{x}.\mathcal{A}$ and its dual $\exists\mathsf{x}.\mathcal{A}$ have arenas $\forall.A$ and $\exists.A$ respectively, and:*

$$
\begin{array}{rclrcl}
\mathcal{W}_{\forall\mathsf{x}.\mathcal{A}}(\varnothing) & = & \top & \mathcal{W}_{\forall\mathsf{x}.\mathcal{A}}(\forall.x_A) & = & \mathcal{W}_{\mathcal{A}}(x_A)[\forall/\mathsf{x}]\\
\mathcal{W}_{\exists\mathsf{x}.\mathcal{A}}(\varnothing) & = & \bot & \mathcal{W}_{\exists\mathsf{x}.\mathcal{A}}(\exists.x_A) & = & \mathcal{W}_{\mathcal{A}}(x_A)[\exists/\mathsf{x}]
\end{array}
$$

Finally, we regard an atomic formula φ (*i.e.* $\mathsf{P}(t_1,\dots,t_n)$ or $\neg\mathsf{P}(t_1,\dots,t_n)$ with $t_i\in\mathsf{Tm}_\Sigma(\mathcal{V})$) as a $\mathcal{V}$-game on arena $\varnothing$, with $\mathcal{W}_\varphi(\varnothing)=\varphi$. We write 1 and $\bot$ for the unit $\mathcal{V}$-games on arena $\varnothing$ with winning conditions respectively $\top$ and $\bot$.

Putting all of these together, we give in Figure 4 the general definition of the $\exists$-biased interpretation of a formula $\varphi\in\mathsf{Form}_\Sigma(\mathcal{V})$ as a $\mathcal{V}$-game. Note the difference between the case of existential and universal formulas, reflecting the bias towards $\exists$loïse in the interpretation. The reader can check that this is indeed compatible with the examples given previously.

We can now state our formulation of Herbrand's theorem.

Theorem II.16. *For any closed formula φ, we have $\vDash\varphi$ iff there exists a finite, top-winning Σ-strategy $\sigma : [\![\varphi]\!]^{\exists}$.*

Although it takes some effort to set up, this is an elegant way of stating Herbrand's theorem, putting the emphasis on the causality between quantifiers. But, besides the game-theoretic language, there is nothing fundamentally new or surprising about this statement. Indeed for now Σ-strategies are static objects, nothing more than alternative bureaucracy-free representations of cut-free proofs. In particular, expansion trees are the *minimal* top-winning Σ-strategies $\sigma : [\![\varphi]\!]^{\exists}$.

E. Compositional Herbrand's theorem

Unlike expansion trees, strategies can be *composed*. Whereas Theorem II.16 above could be deduced via the connection with expansion trees, that proof would intrinsically rely on the admissibility of cut in the sequent calculus. Instead, we will give an alternative proof of Herbrand theorem where the strategy is obtained truly *compositionally* from any sequent proof, without first eliminating cuts. In other words, we want expansion trees to come naturally from the interpretation of the classical sequent calculus in a semantic model.

To compose Σ-strategies, we have to restore the symmetry between $\exists$loïse and $\forall$bélard in the interpretation of formulas. The *non-biased* interpretation $[\![\varphi]\!]_{\mathcal{V}}$ of $\varphi\in\mathsf{Form}_\Sigma(\mathcal{V})$ is defined as for $[\![\varphi]\!]^{\exists}_{\mathcal{V}}$, except for universal formulas, where instead we set $[\![\forall\mathsf{x}\varphi]\!]_{\mathcal{V}} = \,!\forall\mathsf{x}.[\![\varphi]\!]_{\mathcal{V}\uplus\{\mathsf{x}\}}$. This symmetry means that we lose finiteness, since now $\exists$loïse must be reactive to the infinite number of copies potentially opened by $\forall$bélard.

But we can now state:

Theorem II.17. *For φ closed, the following are equivalent:*
(1) $\vDash\varphi$,
(2) *There exists a finite, top-winning Σ-strategy $\sigma : [\![\varphi]\!]^{\exists}$,*
(3) *There exists a* winning Σ-strategy $\sigma : [\![\varphi]\!]$.

We insist that in (3), σ need not be finite. Item (3) is our compositional statement of Herbrand's theorem: these winning strategies will be those computed by our denotational model. The model's construction will show that indeed (1) implies (3), since any proof will yield by interpretation a winning Σ-strategy. But before we go on to that, let us end the discussion here by showing that (3) implies (2), and (2) implies (1).

1) (2) implies (1): We say that a game $\mathcal{A}$ is **a prefix** of $\mathcal{B}$ if $|A|\subseteq|B|$, and all the structure coincides on $|A|$. From any finite, top-winning Σ-strategy $\sigma : [\![\varphi]\!]^{\exists}$, we construct a Herbrand proof (as in Theorem II.5). The $\vee$-expansion follows $\exists$loïse's duplications, the prenexification is any linear ordering of $|\sigma|$ respecting $\leq_\sigma$, and the witness is given by λ_σ.

First, we relate the interpretation of a formula with that of its $\vee$-expansions. The *linear* interpretation $[\![-]\!]^{\ell}_{\mathcal{V}}$ is defined as for $[\![-]\!]^{\exists}_{\mathcal{V}}$ except for the case of existential quantifiers:

$$[\![\exists\mathsf{x}\varphi]\!]^{\ell}_{\mathcal{V}} = \exists\mathsf{x}.[\![\varphi]\!]^{\ell}_{\mathcal{V}}$$

So that both quantifiers are interpreted linearly. Now:

Lemma II.18. *For any closed formula φ and $x\in\mathscr{C}([\![\varphi]\!]^{\exists})$, there is a hereditary $\vee$-expansion φ' of φ and $x'\in\mathscr{C}([\![\varphi']\!]^{\ell})$ order-isomorphic to x (for the order induced by the arena); such that $\mathcal{W}_{[\![\varphi]\!]^{\exists}}(x)$ and $\mathcal{W}_{[\![\varphi']\!]^{\ell}}(x')$ are logically equivalent.*

Proof. The part of $[\![\varphi]\!]^{\exists}$ explored by x (following $\exists$loïse's moves played multiple times) directly informs the syntax tree of a $\vee$-expansion φ' of φ, with an embedding (an injective forest morphism) $[\![\varphi']\!]^{\ell}\hookrightarrow[\![\varphi]\!]^{\exists}$ whose image includes x. Through this embedding x induces $x'\in\mathscr{C}([\![\varphi']\!]^{\ell})$; and $\mathcal{W}_{[\![\varphi]\!]^{\exists}}(x)$ and $\mathcal{W}_{[\![\varphi']\!]^{\ell}}(x')$ are the same up to associativity of $\vee$. □

In particular, for $\sigma : [\![\varphi]\!]^{\exists}$ finite and top-winning, we have $|\sigma|\in\mathscr{C}([\![\varphi]\!]^{\exists})$, so we get a hereditary $\vee$-expansion φ' of φ such that (keeping the renaming silent) $\sigma : [\![\varphi']\!]^{\ell}$ is top-winning. By construction, events of $[\![\varphi']\!]^{\ell}$ are exactly occurrences of quantifiers in φ'. From σ we wish to extract a total order on these, leading to a prenexification. But unlike in Herbrand proofs, σ may refuse to play some existential quantifiers. So we need to *complete* it:

Lemma II.19. *Any finite top-winning $\sigma : [\![\varphi]\!]^{\ell}$ can be extended to a top-winning $\sigma' : [\![\varphi]\!]^{\ell}$ such that $|\sigma'| = |[\![\varphi]\!]^{\ell}|$.*

Proof. A minimal unreached $\exists$ move can be added to σ with the same dependency as in $[\![\varphi]\!]^{\ell}$. The term annotation does not matter; *e.g.* one can use any constant symbol c. Then, we close under receptivity by similarly adding available $\forall$ moves. We obtain a Σ-strategy σ', top-winning as $\mathcal{W}_{[\![\varphi]\!]^{\ell}}(\sigma')[\lambda_{\sigma'}]$ is obtained from $\mathcal{W}_{[\![\varphi]\!]^{\ell}}(\sigma)[\lambda_\sigma]$ by replacing $\bot$ subformulas by something else, preserving its tautological status. As $[\![\varphi]\!]^{\ell}$ is finite, iterating this yields the required Σ-strategy. □

This gives a finite top-winning $\sigma' : [\![\varphi']\!]^{\ell}$ such that $|\sigma'| = |[\![\varphi']\!]^{\ell}|$, *i.e.* a partial order on the quantifiers of the hereditary $\vee$-expansion φ' of φ, and $\lambda_{\sigma'}$ gives a witness. Taking any linear order extending $\leq_{\sigma'}$ yields a prenexification of φ', and $\lambda_{\sigma'}$ completes the data required of a Herbrand proof.

2) (3) implies (2): Take any winning $\sigma : [\![\varphi]\!]$. First we will restrict it to $[\![\varphi]\!]^\exists$ by ignoring ∀bélard's replications; yielding $\sigma^\exists : [\![\varphi]\!]^\exists$ which however is not necessarily finite. However, we will see that it has a finite *top-winning* sub-strategy.

Notice that $[\![\varphi]\!]^\exists$ embeds (subject to renaming) as a prefix of $[\![\varphi]\!]$. Keeping the renaming silent, we have:

Lemma II.20. *For any winning $\sigma : [\![\varphi]\!]$, setting*

$$|\sigma^\exists| = \{a \in |\sigma| \mid [a]_\sigma \subseteq |[\![\varphi]\!]^\exists|\}$$

and inheriting the order, polarity and labelling from σ, we obtain $\sigma^\exists : [\![\varphi]\!]^\exists$ a winning Σ-strategy.

Proof. Most conditions are straightforward. To show $\sigma^\exists : [\![\varphi]\!]^\exists$ winning, we use that for any $\exists$-maximal $x \in \mathscr{C}^\infty(\sigma^\exists)$, we have $x \in \mathscr{C}^\infty(\sigma)$ $\exists$-maximal as well; indeed this follows from $[\![\varphi]\!]^\exists$ being itself $\exists$-maximal in $[\![\varphi]\!]$. □

However, the extracted $\sigma^\exists$ may still not be finite! And indeed it will not always be: there are classical proofs for which our interpretation yields infinite strategies, even after removing ∀bélard's replications (see Appendix A). This reflects the usual issues one has in getting strong normalization in a proof system for classical logic [8] without enforcing too much sequentiality as with a negative translation.

Despite this, the compactness theorem for propositional logic entails that we can always extract a finite top-winning sub-strategy. For $\sigma : [\![\varphi]\!]^\exists$ any Σ-strategy, we write $\mathscr{C}^\forall(\sigma)$ for the set of **∀-maximal** configurations of σ, *i.e.* they can only be extended in σ by ∃loïse moves – inheriting all structure from σ they correspond to its *sub-strategies*, as they are automatically receptive. The proof relies on:

Lemma II.21. *Let X be a directed set of $\forall$-maximal configurations. Then, $\mathcal{W}_{[\![\varphi]\!]^\exists}(\bigcup X)$ is logically equivalent to $\bigvee_{x \in X} \mathcal{W}_{[\![\varphi]\!]^\exists}(x)$.*

Proof. By induction on φ, using simple logical equivalences and that if $x_1 \subseteq x_2$ are $\forall$-maximal configurations, then $\mathcal{W}_{[\![\varphi]\!]^\exists}(x_1)$ implies $\mathcal{W}_{[\![\varphi]\!]^\exists}(x_2)$. □

We complete the proof. For $\sigma : [\![\varphi]\!]^\exists$ winning, by the lemma above the (potentially infinite) disjunction of finite formulas

$$\bigvee_{x \in \mathscr{C}^\forall(\sigma)} \mathcal{W}_{[\![\varphi]\!]^\exists}(x)[\lambda_\sigma]$$

is a tautology. By the compactness theorem there is a finite $X = \{x_1, \dots, x_n\} \subseteq \mathscr{C}^\forall(\sigma)$ such that $\bigvee_{x \in X} \mathcal{W}_{[\![\varphi]\!]^\exists}(x)[\lambda_\sigma]$ is a tautology – *w.l.o.g.* X is directed as $\mathscr{C}^\forall(\sigma)$ is closed under union. By Lemma II.21 again, $\mathcal{W}_{[\![\varphi]\!]^\exists}(\bigcup X)[\lambda_\sigma]$ is a tautology. So, restricting σ to events $\bigcup X$ gives a top-winning finite sub-strategy of σ.

Although the argument above is non-constructive, the extraction of a finite sub-strategy can still be performed in an effective way: Σ-strategies and operations on them can be effectively presented, and the finite top-winning sub-strategy can be effectively obtained by Markov's principle.

$\mathcal{V}$-MLL

$$\textsc{Ax}\ \frac{}{\vdash^{\mathcal{V}} \varphi^\perp, \varphi}\ \mathrm{fv}(\varphi) \subseteq \mathcal{V} \qquad \textsc{Cut}\ \frac{\vdash^{\mathcal{V}} \Gamma, \varphi \qquad \vdash^{\mathcal{V}} \varphi^\perp, \Delta}{\vdash^{\mathcal{V}} \Gamma, \Delta}$$

$$\textsc{Ex}\ \frac{\vdash^{\mathcal{V}} \Gamma, \varphi, \psi, \Delta}{\vdash^{\mathcal{V}} \Gamma, \psi, \varphi, \Delta} \qquad \top\mathrm{I}\ \frac{}{\vdash^{\mathcal{V}} \top} \qquad \bot\mathrm{I}\ \frac{\vdash^{\mathcal{V}} \Gamma}{\vdash^{\mathcal{V}} \Gamma, \bot}$$

$$\wedge\mathrm{I}\ \frac{\vdash^{\mathcal{V}} \Gamma, \varphi \qquad \vdash^{\mathcal{V}} \psi, \Delta}{\vdash^{\mathcal{V}} \Gamma, \varphi \wedge \psi, \Delta} \qquad \vee\mathrm{I}\ \frac{\vdash^{\mathcal{V}} \Gamma, \varphi, \psi, \Delta}{\vdash^{\mathcal{V}} \Gamma, \varphi \vee \psi, \Delta}$$

First-order MLL (MLL$_1$)

$$\forall\mathrm{I}\ \frac{\vdash^{\mathcal{V} \uplus \{\mathsf{x}\}} \Gamma, \varphi}{\vdash^{\mathcal{V}} \Gamma, \forall \mathsf{x}.\, \varphi}\ \mathsf{x} \notin \mathrm{fv}(\Gamma) \qquad \exists\mathrm{I}\ \frac{\vdash^{\mathcal{V}} \Gamma, \varphi[t/\mathsf{x}]}{\vdash^{\mathcal{V}} \Gamma, \exists \mathsf{x}.\, \varphi}\ t \in \mathsf{Tm}_\Sigma(\mathcal{V})$$

LK

$$\mathrm{C}\ \frac{\vdash^{\mathcal{V}} \Gamma, \varphi, \varphi}{\vdash^{\mathcal{V}} \Gamma, \varphi} \qquad \mathrm{W}\ \frac{\vdash^{\mathcal{V}} \Gamma}{\vdash^{\mathcal{V}} \Gamma, \varphi}$$

Fig. 5. Rules for the sequent calculus LK

3) (1) implies (3): This is the key part of our theorem, our main contribution, and the subject of the rest of the paper. It will be proved compositionally, by interpreting proofs from a sequent proof system for first-order classical logic.

Our source sequent calculus is presented in Figure 5. It is a fairly standard one-sided sequent calculus, with rules presented in the multiplicative style; the only notable variation is in how we deal with free variables. We opt for sequents that carry explicitly a set $\mathcal{V}$ of free variables, that may appear freely in the formulas. Then the introduction rule for $\forall$ introduces a fresh variable, whereas the introduction rule for $\exists$ provides a term whose free variables must be in $\mathcal{V}$.

What mathematical structure is required to interpret this sequent calculus? Focusing on the first group of rules and omitting the $\mathcal{V}$ annotations, we get the rules of Multiplicative Linear Logic (MLL). Propositional MLL can be interpreted in a $\star$-autonomous category [2]. Accordingly we will first construct in Section III a $\star$-autonomous category Games of games and winning Σ-strategies. In Section IV, still ignoring contraction and weakening, we construct the structure required for the interpretation of quantifiers. For each set of variables $\mathcal{V}$ we construct a $\star$-autonomous category $\mathcal{V}$-Games, with a fibred structure to link the $\mathcal{V}$-Games together for distinct $\mathcal{V}$s and suitable structure to deal with quantifiers, obtaining a model of first-order MLL. Finally in Section V we complete the interpretation by adding the exponential modalities from linear logic to the interpretation of quantifiers, and get from that an interpretation of contraction and weakening.

III. A $\star$-AUTONOMOUS CATEGORY

The starting point of our discussion will be the following folklore theorem of cut elimination for MLL.

Theorem III.1. *There is a set of reduction rules on MLL sequent proofs, written $\rightsquigarrow_{\mathrm{MLL}}$, such that for any proof π of a sequent $\vdash \Gamma$, there is a cut-free π' of Γ such that $\pi \rightsquigarrow^*_{\mathrm{MLL}} \pi'$.*

The reduction $\rightsquigarrow_{\mathrm{MLL}}$ comprises so-called *logical* reduction steps, reducing a cut on a formula $\varphi/\varphi^\perp$, between two proofs starting with the introduction rule for the main connective of $\varphi/\varphi^\perp$; and *structural* reduction steps, which consist in performing commutations between rules so as to reach the logical steps. We assume in the sequel that the reader has some familiarity with this process.

In this section we aim to give an interpretation of MLL proofs, which should be invariant under cut-elimination. Categorical logic tells us that doing that is essentially the same as producing a *$\ast$-autonomous category* – although we opt here for its reformulation by Cockett and Seely as a *symmetric linearly distributive category with negation* [6].

Definition III.2. *A **symmetric linearly distributive category** is a category $\mathcal{C}$ with two symmetric monoidal structures $(\otimes, 1)$ and $(\parr, \perp)$ which* distribute*: there is a natural transformation*

$$\delta_{A,B,C} : A \otimes (B \parr C) \overset{\mathcal{C}}{\to} (A \otimes B) \parr C$$

the linear distribution*, subject to coherence conditions [6].*

*A symmetric linearly distributive category **with negation** also has a function $(-)^\perp$ on objects and families of maps:*

$$\eta_A : 1 \overset{\mathcal{C}}{\to} A^\perp \parr A \qquad \epsilon_A : A \otimes A^\perp \overset{\mathcal{C}}{\to} \perp$$

such that the canonical composition $A \to A \otimes (A^\perp \parr A) \to (A \otimes A^\perp) \parr A \to A$, and its dual $A^\perp \to A^\perp$, are identities.

Note also the degenerate case of a **compact closed category**, which is a symmetric linearly distributive category where the monoidal structures $(\otimes, 1)$ and $(\parr, \perp)$ coincide.

Abusing terminology, we will refer in the future to *symmetric linearly distributive categories with negation* by the shorter *$\ast$-autonomous categories*. This should not create any confusion in the light of their equivalence [6]. If $\mathcal{C}$ a $\ast$-autonomous category comes with a choice of $[\![\mathsf{P}(t_1, \dots, t_n)]\!]$ (an object of $\mathcal{C}$) for all closed atomic formulas, then this interpretation can be extended to all closed quantifier-free formulas following Figure 4. For all such φ, we have $[\![\varphi^\perp]\!] = [\![\varphi]\!]^\perp$.

For the sake of completeness, we review the (standard) interpretation of MLL proofs in $\mathcal{C}$. A proof π of a *MLL sequent* $\vdash \varphi_1, \dots, \varphi_n$ is interpreted as a morphism

$$[\![\pi]\!] : 1 \overset{\mathcal{C}}{\to} [\![\varphi_1]\!] \parr \cdots \parr [\![\varphi_n]\!]$$

as indicated in Figure 6 (where we omit the semantic brackets on formulas, the unlabeled arrows make use of canonical structural morphisms, and some structural isomorphisms are left silent). This shows *soundness w.r.t.* provability: if φ is provable, then $1 \to_{\mathcal{C}} [\![\varphi]\!]$ is inhabited. But the canonical interpretation of MLL into a $\ast$-autonomous category gives us more: the interpretation is invariant under cut reduction.

Theorem III.3. *If $\pi \rightsquigarrow_{\mathrm{MLL}} \pi'$ are proofs of $\vdash \Gamma$, $[\![\pi]\!] = [\![\pi']\!]$.*

Hence a proof has the same denotation as its cut-free form obtained by Theorem III.1. In the rest of this section we construct a concrete $\ast$-autonomous category of games and winning Σ-strategies; supporting by the reasoning above the interpretation of MLL. This will be done in three stages: first, we will focus on defining composition of Σ-strategies (without winning). Then, we will show that this defines a compact closed category, and finally we will add winning conditions, separating the monoidal product $\parallel$ of the compact closed structure into two monoidal products $\otimes$ and $\parr$.

A. Composition of Σ-strategies

We construct a category Arenas_Σ having arenas as objects, and as morphisms from A to B the Σ-strategies $\sigma : A^\perp \parallel B$, also written $\sigma : A \overset{\mathrm{Ar}_\Sigma}{\to} B$. The composition of $\sigma : A \overset{\mathrm{Ar}_\Sigma}{\to} B$ and $\tau : B \overset{\mathrm{Ar}_\Sigma}{\to} C$ will be computed in two stages: first, the *interaction* $\tau \circledast \sigma$ is obtained as the most general partial-order-with-terms satisfying the constraints given by both σ and τ – Figure 7 displays such an interaction. Then, we will obtain the *composition* $\tau \odot \sigma$ by hiding events in B. In the example of Figure 7 we get the single annotated event $\exists_5^{\mathsf{f}(\mathsf{g}(\mathsf{c}),\mathsf{h}(\mathsf{c}))}$.

First we give a few definitions and notations on terms and substitutions. If $\mathcal{V}_1, \mathcal{V}_2$ are finite sets, a **substitution** $\gamma : \mathcal{V}_1 \overset{\mathcal{S}}{\to} \mathcal{V}_2$ is a function $\gamma : \mathcal{V}_2 \to \mathsf{Tm}_\Sigma(\mathcal{V}_1)$. For $t \in \mathsf{Tm}_\Sigma(\mathcal{V}_2)$, we write $t[\gamma] \in \mathsf{Tm}_\Sigma(\mathcal{V}_1)$ for the substitution operation. Substitutions form a category $\mathcal{S}$, which is *cartesian*: the empty set $\varnothing$ is terminal, and the product of $\mathcal{V}_1$ and $\mathcal{V}_2$ is their disjoint union $\mathcal{V}_1 + \mathcal{V}_2$. From $\gamma : \mathcal{V}_1 \overset{\mathcal{S}}{\to} \mathcal{V}_2$ and $\gamma' : \mathcal{V}_1' \overset{\mathcal{S}}{\to} \mathcal{V}_2$, we say that γ **subsumes** γ', written $\gamma' \preccurlyeq \gamma$, if there is $\alpha : \mathcal{V}_1' \overset{\mathcal{S}}{\to} \mathcal{V}_2$ *s.t.* $\gamma \circ \alpha = \gamma'$ – giving a preorder on substitutions with codomain $\mathcal{V}_2$.

1) Interaction: Consider first a *closed* interaction of $\sigma : A$ and $\tau : A^\perp$. As they disagree on the polarities on A we drop them – $\tau \circledast \sigma$ will be a *neutral Σ-strategy* on a *neutral arena*:

Definition III.4. *A **neutral arena** is an arena, without polarities. **Neutral strategies** on A, still written $\sigma : A$, are defined as in Definition II.7 without conditions* (2), (3)*. **Neutral** Σ**-strategies** additionally have $\lambda_\sigma : (s \in |\sigma|) \to \mathsf{Tm}_\Sigma([s]_\sigma)$.*

Forgetting polarities, every Σ-strategy can be silently coerced into a neutral one. Given σ and τ, $\tau \circledast \sigma$ must be a *minimal strengthening* of σ and τ, both in terms of causal structure and term annotations, *i.e.* a *meet* for the partial order:

Definition III.5. *If $\sigma, \tau : A$ are neutral Σ-strategies, we write $\sigma \preccurlyeq \tau$ iff $|\sigma| \subseteq |\tau|$, $\mathscr{C}^\infty(\sigma) \subseteq \mathscr{C}^\infty(\tau)$, and for all $x \in \mathscr{C}(|\sigma|)$, $\lambda_\tau \upharpoonright x$ subsumes $\lambda_\sigma \upharpoonright x$ (both regarded as substitutions $x \overset{\mathcal{S}}{\to} x$).*

Ignoring terms, any two σ and τ have a meet $\sigma \wedge \tau$; indeed this is then a simplification of the *pullback* in the category of event structures [26], exploiting the absence of conflict. The partial order $(|\sigma \wedge \tau|, \leq_{\sigma\wedge\tau})$ has for events all common moves of σ and τ for which there is a consistent causal history compatible with both $\leq_\sigma$ and $\leq_\tau$, and for $\leq_{\sigma\wedge\tau}$ the minimal causal order compatible with both.

However, taking terms into account, two neutral Σ-strategies may not necessarily have a meet for $\preccurlyeq$ (see Appendix B).

$$\left[\!\!\left[\textsc{Ax}\,\frac{}{\vdash \varphi^\perp, \varphi}\right]\!\!\right] = 1 \xrightarrow[\mathcal{C}]{\eta_\varphi} \varphi^\perp \parr \varphi \qquad \left[\!\!\left[\top\mathrm{I}\,\frac{}{\vdash \top}\right]\!\!\right] = 1 \xrightarrow{\mathcal{C}} 1 \qquad \left[\!\!\left[\bot\mathrm{I}\,\frac{\dfrac{\pi}{\vdash \Gamma}}{\vdash \Gamma, \bot}\right]\!\!\right] = 1 \xrightarrow[\mathcal{C}]{[\![\pi]\!]} \Gamma \cong \Gamma \parr \bot \qquad \left[\!\!\left[\textsc{Ex}\,\frac{\dfrac{\pi}{\vdash \Gamma, \varphi, \psi, \Delta}}{\vdash \Gamma, \psi, \varphi, \Delta}\right]\!\!\right] = 1 \xrightarrow[\mathcal{C}]{[\![\pi]\!]} \Gamma \parr \varphi \parr \psi \parr \Delta \cong \Gamma \parr \psi \parr \varphi \parr \Delta$$

$$\left[\!\!\left[\textsc{Cut}\,\frac{\dfrac{\pi_1}{\vdash \Gamma, \varphi} \quad \dfrac{\pi_2}{\vdash \varphi^\perp, \Delta}}{\vdash \Gamma, \Delta}\right]\!\!\right] = 1 \xrightarrow[\mathcal{C}]{[\![\pi_1]\!] \otimes [\![\pi_2]\!]} (\Gamma \parr \varphi) \otimes (\varphi^\perp \parr \Delta) \xrightarrow{\mathcal{C}} \Gamma \parr (\varphi \otimes \varphi^\perp) \parr \Delta \xrightarrow[\mathcal{C}]{\Gamma \otimes \epsilon_\varphi \otimes \Delta} \Gamma \parr \bot \parr \Delta \cong_{\mathcal{C}} \Gamma \parr \Delta$$

$$\left[\!\!\left[\wedge\mathrm{I}\,\frac{\dfrac{\pi_1}{\vdash \Gamma, \varphi} \quad \dfrac{\pi_2}{\vdash \psi, \Delta}}{\vdash \Gamma, \varphi \wedge \psi, \Delta}\right]\!\!\right] = 1 \xrightarrow[\mathcal{C}]{[\![\pi_1]\!] \otimes [\![\pi_2]\!]} (\Gamma \parr \varphi) \otimes (\psi \parr \Delta) \xrightarrow{\mathcal{C}} \Gamma \parr (\varphi \otimes \psi) \parr \Delta \qquad \left[\!\!\left[\vee\mathrm{I}\,\frac{\dfrac{\pi}{\vdash \Gamma, \varphi, \psi, \Delta}}{\vdash \Gamma, \varphi \vee \psi, \Delta}\right]\!\!\right] = 1 \xrightarrow[\mathcal{C}]{[\![\pi]\!]} \Gamma \parr \varphi \parr \psi \parr \Delta \cong \Gamma \parr (\varphi \parr \psi) \parr \Delta$$

Fig. 6. Interpretation of MLL in a $\ast$-autonomous category

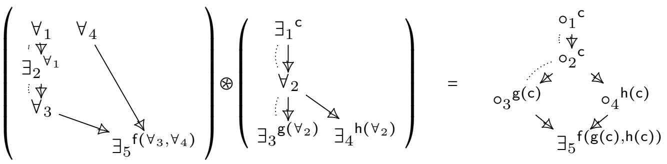

Fig. 7. Interaction of $\sigma : 1^\perp \parallel (\exists_1 \forall_2 \exists_3 \parallel \exists_4)$ and $\tau : (\exists_1 \forall_2 \exists_3 \parallel \exists_4)^\perp \parallel \exists_5$

Hence, we focus on the meets occurring from compositions of Σ-strategies. For $\sigma : A$ and$\tau : A^\perp$ dual Σ-strategies we will see that the meet *does* exists. However this is not sufficient since for composable $\sigma : A^\perp \parallel B$ and $\tau : B^\perp \parallel C$, the games are not purely dual: we need to "pad out" σ and τ and compute instead the meet $(\sigma \parallel C^\perp) \wedge (A \parallel \tau)$, where the parallel composition of Definition II.11 is extended with terms in the obvious way, and we set $\lambda_A(a) = a$ for all $a \in |A|$. Now $\sigma \parallel C^\perp : A^\perp \parallel B \parallel C^\perp$ and $A \parallel \tau : A \parallel B^\perp \parallel C$ are dual, albeit only *pre-Σ-strategies*.

Lemma III.6. *If A is an arena, $\sigma : A$ is a **pre-Σ-strategy** if $\sigma : A$ is a strategy, $\lambda_\sigma(a) = a$ for all $a^\forall \in |\sigma|$, and λ_σ is **idempotent**: for all $a \in |\sigma|$, $\lambda_\sigma(a)[\lambda_\sigma] = \lambda_\sigma(a)$.*

Any $\sigma : A$ and $\tau : A^\perp$ pre-Σ-strategies have a meet $\sigma \wedge \tau$.

Proof. We start with the causal meet $\sigma \wedge \tau$ mentioned above. Then $\lambda_{\sigma\wedge\tau}$ is the *most general unifier* of $\lambda_\sigma \upharpoonright |\sigma \wedge \tau|$ and $\lambda_\tau \upharpoonright |\sigma \wedge \tau|$, obtained by well-founded induction on $\leq_{\sigma\wedge\tau}$:

$$\lambda_{\sigma\wedge\tau}(a) = \begin{cases} \lambda_\sigma(a)[\lambda_{\sigma\wedge\tau} \upharpoonright [a)] & \text{if } \mathrm{pol}_A(a) = \exists \\ \lambda_\tau(a)[\lambda_{\sigma\wedge\tau} \upharpoonright [a)] & \text{if } \mathrm{pol}_A(a) = \forall \end{cases}$$

where $[a) = \{a' \in A \mid a' <_{\sigma\wedge\tau} a\}$. It follows that this is indeed the *m.g.u.* – in particular, we exploit that from idempotence, if $a^\exists \in |\sigma|$ then either $\lambda_\sigma(a) \in \mathsf{Tm}_\Sigma([a)_\sigma)$ or $\lambda_\sigma(a) = a$. ☐

From Σ-strategies $\sigma : A^\perp \parallel B$ and $\tau : B^\perp \parallel C$ we define $\tau \circledast \sigma = (\sigma \parallel C^\perp) \wedge (A \parallel \tau) : A \parallel B \parallel C$. We observe that variables appearing in $\lambda_{\tau\circledast\sigma}$ cannot be events in B – in fact they must be negative in $A^\perp \parallel C$. So we can define

$$\tau \odot \sigma = (\tau \circledast \sigma) \cap (A \parallel C)$$

the restriction of $\tau \circledast \sigma$ to $A \parallel C$, with the same causal order and term annotation. The causal structure $(|\tau \odot \sigma|, \leq_{\tau\odot\sigma})$ is receptive and courteous as an instance of the constructions in [4], and the terms satisfy our conditions. Therefore $\tau \odot \sigma : A^\perp \parallel C$ is a Σ-strategy, the **composition** of σ and τ.

Because interaction is defined as a meet for $\preccurlyeq$, it follows that it is compatible with it, *i.e.* if $\sigma \preccurlyeq \sigma'$, then $\tau \circledast \sigma \preccurlyeq \tau \circledast \sigma'$. This is preserved by projection, and hence $\tau \odot \sigma \preccurlyeq \tau \odot \sigma'$ as well. We will use later on this compatibility of composition with $\preccurlyeq$, and the fact that $\preccurlyeq$ is more constrained on Σ-strategies:

Lemma III.7. *If $\sigma, \sigma' : A$ are Σ-strategies, then if $\sigma \preccurlyeq \sigma'$, it automatically holds that $\lambda_\sigma(s) = \lambda_{\sigma'}(s)$, for all $s \in |\sigma|$.*

This follows from the definition of $\preccurlyeq$, and the constraints on term labelling for Σ-strategies.

To get a category, we also define the *copycat strategy*.

Definition III.8. *For an arena A, the **copycat Σ-strategy** $\text{\rm c\!\!c}_A : A^\perp \parallel A$ has events $|\text{\rm c\!\!c}_A| = A^\perp \parallel A$. Writing $\overline{(i,a)} = (3 - i, a)$, its partial order $\leq_{\text{\rm c\!\!c}_A}$ is the transitive closure of*

$$\leq_{A^\perp \parallel A} \cup \{(c, \overline{c}) \mid c^\forall \in |A^\perp \parallel A|\}$$

and its labelling function is $\lambda_{\text{\rm c\!\!c}_A}(c^\forall) = c$, $\lambda_{\text{\rm c\!\!c}_A}(c^\exists) = \overline{c}$.

We have now given all the data of the category Arenas_Σ.

Proposition III.9. *There is a poset-enriched category Arenas_Σ with objects arenas, and Σ-strategies as morphisms.*

The proof of associativity of composition and neutrality of copycat are elaborations on the construction of the bicategory in [4], and are omitted for lack of space.

B. Compact closed structure

We show that Arenas_Σ is compact closed. The **tensor product** of arenas A and B is $A \parallel B$. For Σ-strategies $\sigma_1 : A_1^\perp \parallel B_1$ and $\sigma_2 : A_2^\perp \parallel B_2$, we have $\sigma_1 \parallel \sigma_2 : (A_1^\perp \parallel B_1) \parallel (A_2^\perp \parallel B_2)$, which is isomorphic to $(A_1 \parallel A_2)^\perp \parallel (B_1 \parallel B_2)$ – overloading notation, we keep $\sigma_1 \parallel \sigma_2$ for the renaming:

$$\sigma_1 \parallel \sigma_2 : (A_1 \parallel A_2)^\perp \parallel (B_1 \parallel B_2)$$

When writing $\sigma_1 \parallel \sigma_2$, the context should always make clear which parallel operation we are referring to. We get:

Proposition III.10. *Simple parallel composition extends to an enriched functor* $\parallel : \mathrm{Arenas}_\Sigma \times \mathrm{Arenas}_\Sigma \to \mathrm{Arenas}_\Sigma$.

For the compact closed structure, we elaborate the renaming operation used above. We write $f : A \cong B$ for an **isomorphism of arenas**, preserving and reflecting all structure.

Definition III.11. *For $f : A \cong B$ and $\sigma : A$ a Σ-strategy, the* ***renaming*** *$f_*(\sigma) : B$ has components $|f_*(\sigma)| = f|\sigma|$, $\leq_{f_*(\sigma)} = \{(f\,a_1, f\,a_2) \mid a_1 \leq_\sigma a_2\}$ and $\lambda_{f_*(\sigma)}(f\,a) = \lambda_\sigma(a)[f]$.*

In particular, if $f : A \cong B$, then its **lifting** is $\overline{f} = (A^\perp \parallel f)_*(c\!\!c_A) : A^\perp \parallel B$. Lifting gives us an easy way to define the structural morphisms for the symmetric monoidal closed structure of Arenas_Σ. For instance, we have an iso $\alpha_{A,B,C} : (A \parallel B) \parallel C \cong A \parallel (B \parallel C)$, which by lifting gives us:

$$\overline{\alpha_{A,B,C}} : (A \parallel B) \parallel C \xrightarrow{\mathrm{Ar}_\Sigma} A \parallel (B \parallel C)$$

The other structural morphisms arise similarly. In order to prove coherence and naturality, we observe:

Lemma III.12. *Let $f : A \cong A'$, $g : B \cong B'$, $h : C \cong C'$, and $\sigma : A^\perp \parallel B, \tau : B^\perp \parallel C$. Then,*

$$(f^\perp \parallel h)_*(\tau \odot \sigma) = (g^\perp \parallel h)_*(\tau) \odot (f^\perp \parallel g)_*(\sigma)$$

Indeed, the interaction-as-meet easily transfers through isomorphisms so the equality above holds for interactions, and hence for composition. From that and neutrality of copycat for composition we immediately deduce the *lifting lemma*:

Lemma III.13. *For $\sigma : A^\perp \parallel B$ a Σ-strategy and $f : B \cong C$,*

$$\overline{f} \odot \sigma = (A^\perp \parallel f)_*(\sigma) : A^\perp \parallel C\,.$$

As a corollary we get coherence for the structural morphisms (following from those on isomorphisms), and naturality. For all A we get by the obvious renaming from $c\!\!c_A$:

$$\eta_A : \varnothing \xrightarrow{\mathrm{Ar}_\Sigma} A^\perp \parallel A \qquad \epsilon_A : A \parallel A^\perp \xrightarrow{\mathrm{Ar}_\Sigma} \varnothing$$

checking the law for compact closed categories is a variation of the idempotence of copycat. Overall we have:

Proposition III.14. Arenas_Σ *is a poset-enriched compact closed category.*

C. A linearly distributive category with negation

Finally, we reinstate winning conditions. We first note:

Proposition III.15. *There is a (poset-enriched) category* Games_Σ *with objects the games (Definition II.9) on Σ, and morphisms Σ-strategies $\sigma : \mathcal{A}^\perp \parr \mathcal{B}$, also written $\sigma : \mathcal{A} \xrightarrow{\mathrm{Ga}_\Sigma} \mathcal{B}$.*

That copycat is winning boils down to the excluded middle. That $\tau \odot \sigma : \mathcal{A}^\perp \parr \mathcal{C}$ is winning if $\sigma : \mathcal{A}^\perp \parr \mathcal{B}$ and $\tau : \mathcal{B}^\perp \parr \mathcal{C}$ are, is as in [5]: for $x \in \mathscr{C}(\tau \odot \sigma)$ $\exists$-maximal we find a witness $y \in \mathscr{C}(\tau \circledast \sigma)$ (*i.e.* $y \cap (A \parallel C) = x$) *s.t.* $y \cap (A \parallel B) \in \sigma$, $y \cap (B \parallel C) \in \tau$ are $\exists$-maximal; the property follows by transitivity of implication. The equations follow from Arenas_Σ. Likewise:

Proposition III.16. *The functor* $\parallel : \mathrm{Arenas}_\Sigma \times \mathrm{Arenas}_\Sigma \to \mathrm{Arenas}_\Sigma$ *splits into* $\otimes, \parr : \mathrm{Games}_\Sigma \times \mathrm{Games}_\Sigma \to \mathrm{Games}_\Sigma$.

It suffices to check winning, which is straightforward.

For the rest of the structure of a linearly distributive category with negation, we only need to check that the structural morphisms from Arenas_Σ satisfy the required winning conditions. As they are obtained from lifting, we give conditions for a lifted map to give a winning Σ-strategy. For $\mathcal{A}, \mathcal{B}$ games, a **win-iso** $f : \mathcal{A} \to \mathcal{B}$ is an iso $f : A \cong B$ such that $(\mathcal{W}_\mathcal{A}(x))^\perp \vee \mathcal{W}_B(f\,x)$ is a tautology, for all $x \in \mathscr{C}^\infty(A)$.

Lemma III.17. *If $f : \mathcal{A} \to \mathcal{B}$ is a win-iso, then $\overline{f} : \mathcal{A}^\perp \parr \mathcal{B}$ is a winning Σ-strategy.*

This easily entails that all structural morphisms (including linear distributivity) are winning. Finally $\eta_\mathcal{A} : 1 \xrightarrow{\mathrm{Ga}_\Sigma} \mathcal{A}^\perp \parr \mathcal{A}$ and $\epsilon_\mathcal{A} : \mathcal{A} \otimes \mathcal{A}^\perp \xrightarrow{\mathrm{Ga}_\Sigma} \bot$ are winning, which concludes:

Proposition III.18. Games_Σ *is a poset-enriched $*$-autonomous category.*

IV. A MODEL OF FIRST-ORDER MLL

$$\textsc{Cut}\ \dfrac{\forall\mathrm{I}\ \dfrac{\dfrac{\pi_1}{\vdash^{\mathcal{V}\uplus\{\mathsf{x}\}} \Gamma, \varphi}}{\vdash^{\mathcal{V}} \Gamma, \forall\mathsf{x}.\,\varphi} \qquad \exists\mathrm{I}\ \dfrac{\dfrac{\pi_2}{\vdash^{\mathcal{V}} \varphi^\perp[t/\mathsf{x}], \Delta}}{\vdash^{\mathcal{V}} \exists\mathsf{x}.\,\varphi^\perp, \Delta}}{\vdash^{V} \Gamma, \Delta} \ \rightsquigarrow_{\forall/\exists}\ \textsc{Cut}\ \dfrac{\dfrac{\pi_1[t/\mathsf{x}]}{\vdash^{\mathcal{V}} \Gamma, \varphi[t/\mathsf{x}]} \qquad \dfrac{\pi_2}{\vdash^{\mathcal{V}} \varphi^\perp[t/\mathsf{x}], \Delta}}{\vdash^{\mathcal{V}} \Gamma, \Delta}$$

$$\textsc{Cut}\ \dfrac{\dfrac{\pi_1}{\vdash^{\mathcal{V}} \Gamma, \psi} \qquad \forall\mathrm{I}\ \dfrac{\dfrac{\pi_2}{\vdash^{\mathcal{V}\uplus\{\mathsf{x}\}} \psi^\perp, \Delta, \varphi}}{\vdash^{\mathcal{V}} \psi^\perp, \Delta, \forall\mathsf{x}.\,\varphi}}{\vdash^{\mathcal{V}} \Gamma, \Delta, \forall\mathsf{x}.\,\varphi} \ \rightsquigarrow_{\textsc{Cut}/\forall}\ \forall\mathrm{I}\ \dfrac{\textsc{Cut}\ \dfrac{\dfrac{\pi_1}{\vdash^{\mathcal{V}\uplus\{\mathsf{x}\}} \Gamma, \psi} \qquad \dfrac{\pi_2}{\vdash^{\mathcal{V}\uplus\{\mathsf{x}\}} \psi^\perp, \Delta, \varphi}}{\vdash^{\mathcal{V}\uplus\{\mathsf{x}\}} \Gamma, \Delta, \varphi}}{\vdash^{\mathcal{V}} \Gamma, \Delta, \forall\mathsf{x}.\,\varphi}$$

$$\textsc{Cut}\ \dfrac{\dfrac{\pi_1}{\vdash^{\mathcal{V}} \Gamma, \psi} \qquad \exists\mathrm{I}\ \dfrac{\dfrac{\pi_2}{\vdash^{\mathcal{V}} \psi^\perp, \Delta, \varphi[t/\mathsf{x}]}}{\vdash^{\mathcal{V}} \psi^\perp, \Delta, \exists\mathsf{x}.\,\varphi}}{\vdash^{\mathcal{V}} \Gamma, \Delta, \exists\mathsf{x}.\,\varphi} \ \rightsquigarrow_{\textsc{Cut}/\exists}\ \exists\mathrm{I}\ \dfrac{\textsc{Cut}\ \dfrac{\dfrac{\pi_1}{\vdash^{\mathcal{V}} \Gamma, \psi} \qquad \dfrac{\pi_2}{\vdash^{\mathcal{V}} \psi^\perp, \Delta, \varphi[t/\mathsf{x}]}}{\vdash^{\mathcal{V}} \Gamma, \Delta, \varphi[t/\mathsf{x}]}}{\vdash^{\mathcal{V}} \Gamma, \Delta, \exists\mathsf{x}.\,\varphi}$$

Fig. 8. Additional cut elimination rules for MLL_1

We move on to MLL_1, *i.e.* the MLL rules (taking into account the annotation with $\mathcal{V}$) plus the two rules for quantifiers. Before developing the interpretation of these, let us discuss cut elimination for them. The new cut reduction rules are displayed in Figure 8. There are only three of them: one for the unique new *logical* cut reduction rule ($\forall/\exists$), and two which explain the propagation of cuts past introduction rules for $\forall$ and $\exists$. Writing $\pi \rightsquigarrow_{\mathrm{MLL}_1} \pi'$ for the reduction obtained with these new rules together with those for MLL, we have:

Proposition IV.1. *Let π be any MLL_1 proof of $\vdash^{\mathcal{V}} \Gamma$. Then, there exists a cut-free proof π' of $\vdash^{\mathcal{V}} \Gamma$ such that $\pi \rightsquigarrow^*_{\mathrm{MLL}_1} \pi'$.*

The first rule of Figure 8 requires the introduction of *substitution* as an operation on proofs. In general, given a proof π of $\vdash^{\mathcal{V}_2} \Gamma$ and substitution $\gamma : \mathcal{V}_1 \to \mathcal{V}_2$ we obtain $\pi[\gamma]$ a proof of $\vdash^{\mathcal{V}_1} \Gamma[\gamma]$ by propagating γ through π, performing the substitution on formulas and terms. A degenerate case of this is the substitution of a proof π of $\vdash^{\mathcal{V}} \Gamma$ by the *weakening* substitution $\mathrm{w}_{\mathcal{V},\mathsf{x}} : \mathcal{V} \uplus \{\mathsf{x}\} \to \mathcal{V}$, obtaining $\pi_1[\mathrm{w}_{\mathcal{V},\mathsf{x}}]$, a proof of $\vdash^{\mathcal{V}\uplus\mathsf{x}} \Gamma$. As this actually leaves the formulas and terms unchanged we leave it implicit in the cut reduction rules – it is used for instance implicitly in the commutation $\textsc{Cut}/\forall$.

Substitution is key in the cut reduction of quantifiers. However it is best studied independently of quantifiers, in a model of $\mathcal{V}$-MLL. This is the topic of the next subsection, prior to the interpretation of the introduction rules for quantifiers.

A. A fibred model of $\mathcal{V}$-MLL

1) Categorical structure: Following categorical logic [18], [24], $\mathcal{V}$-MLL together with substitution is modeled in:

Definition IV.2. *Let $\ast$-Aut be the category of $\ast$-autonomous categories and functors preserving the structure on the nose. A **strict $\mathcal{S}$-indexed $\ast$-autonomous category** is a functor:*

$$\mathcal{T} : \mathcal{S}^{\mathrm{op}} \to \ast\text{-}Aut$$

Usually (*e.g. hyperdoctrines* [24]), similar definitions are more general: $\ast$-autonomous functors only preserve the structure up to coherent isomorphism, and the category of substitutions $\mathcal{S}$ is an abstract cartesian category. We opt here for this simpler definition to focus on our concrete interpretation.

Writing $\mathcal{V}_n = \{\mathsf{x}_1, \dots, \mathsf{x}_n\}$, we say that $\mathcal{T}$ **supports** Σ if for every predicate symbol P of arity n there is is $[\![\mathsf{P}]\!]_{\mathcal{V}_n}$ a chosen object of $\mathcal{T}(\mathcal{V}_n)$. For $t_1, \dots, t_n \in \mathsf{Tm}_\Sigma(\mathcal{V})$ we can then set

$$[\![\mathsf{P}(t_1, \dots, t_n)]\!] = \mathcal{T}([t_1/\mathsf{x}_1, \dots, t_n/\mathsf{x}_n])([\![\mathsf{P}]\!]_{\mathcal{V}_n})\,,$$

an object of $\mathcal{T}(\mathcal{V})$, also written $[\![\mathsf{P}]\!]_{\mathcal{V}_n}[t_1/\mathsf{x}_1, \dots, t_n/\mathsf{x}_n]$.

For any finite $\mathcal{V}$, this lets us interpret $\mathcal{V}$-MLL in $\mathcal{T}(\mathcal{V})$ by following Section III. But besides interpreting $\mathcal{V}$-MLL in isolation, this structure also models the action of substitutions. Indeed, the functorial action of $\mathcal{T}$ on $\gamma : \mathcal{V}_1 \to \mathcal{V}_2$ induces substitution operations, which to any object A of $\mathcal{T}(\mathcal{V}_2)$ associate $A[\gamma]$ in $\mathcal{T}(\mathcal{V}_1)$, and which to any morphism $\sigma : A \overset{\mathcal{T}(\mathcal{V}_2)}{\to} B$ associates $\sigma[\gamma] : A[\gamma] \overset{\mathcal{T}(\mathcal{V}_1)}{\to} B[\gamma]$. This matches syntactic substitution, as $\mathcal{T}(\gamma)$ preserves the $\ast$-autonomous structure.

2) Concrete structure: For any finite set $\mathcal{V}$, the *fibre* $\mathcal{T}(\mathcal{V})$ is the $\ast$-autonomous category $\mathrm{Games}_{\Sigma\uplus\mathcal{V}}$ built in Section III, on the extended signature $\Sigma \uplus \mathcal{V}$.

We comment on this definition of $\mathcal{T}(\mathcal{V})$. Its *objects* are games on the signature $\Sigma\uplus\mathcal{V}$, *i.e.* the $\mathcal{V}$*-games* of Section II-D. *Morphisms* between $\mathcal{V}$-games $\mathcal{A}$ and $\mathcal{B}$ are winning $(\Sigma \uplus \mathcal{V})$-strategies on $\mathcal{A}^\perp \parr \mathcal{B}$ regarded as a game on signature $\Sigma \uplus \mathcal{V}$ – also called **winning Σ-strategies on the $\mathcal{V}$-game** $\mathcal{A}^\perp \parr \mathcal{B}$.

We introduce the functorial action of $\mathcal{T}$. For $\mathcal{A}$ a $\mathcal{V}_2$-game and $\gamma : \mathcal{V}_1 \to \mathcal{V}_2$ a substitution, the game $\mathcal{T}(\gamma)(\mathcal{A}) = \mathcal{A}[\gamma]$ is defined as having arena A, and, for $x \in \mathscr{C}^\infty(A)$:

$$\mathcal{W}_{\mathcal{A}[\gamma]}(x) = \mathcal{W}_{\mathcal{A}}(x)[\gamma] \in \mathsf{QF}^\infty_{\Sigma\uplus\mathcal{V}_1}(x)$$

Likewise, given $\mathcal{A}$ and $\mathcal{B}$ two $\mathcal{V}$-games and $\sigma : \mathcal{A}^\perp \parr \mathcal{B}$ a winning strategy, we set $\sigma[\gamma]$ to have the same components as σ, but with the term annotations substituted with γ:

$$\lambda_{\sigma[\gamma]}(s) = \lambda(s)[\gamma] \in \mathsf{Tm}_{\Sigma\uplus\mathcal{V}_1}(x)$$

It is a simple verification to prove:

Proposition IV.3. *For any substitution $\gamma : \mathcal{V}_1 \to \mathcal{V}_2$, $\mathcal{T}(\gamma) : \mathcal{T}(\mathcal{V}_2) \to \mathcal{T}(\mathcal{V}_1)$ is a strict $\ast$-autonomous functor.*

B. Quantifiers

To complete the interpretation of MLL_1, we give the interpretation of $\forall$I and $\exists$I. Recall that for now, we use the *linear* interpretation $[\![-]\!]^\ell$ of formulas from Section II-E.

Besides preserving the $\ast$-autonomous structure, substitution also propagates through quantifiers, from which we have:

Lemma IV.4. *Let $\varphi \in \mathsf{Form}_\Sigma(\mathcal{V}_2)$ and $\gamma : \mathcal{V}_1 \to \mathcal{V}_2$ a substitution, then $[\![\varphi[\gamma]]\!]^\ell_{\mathcal{V}_1} = [\![\varphi]\!]^\ell_{\mathcal{V}_2}[\gamma]$.*

This fact will be used implicitly from now on.

First of all, we note that the definition of quantifiers on games of Definition II.15 extends to functors:

$$\forall_{\mathcal{V},\mathsf{x}}, \exists_{\mathcal{V},\mathsf{x}} : \mathcal{T}(\mathcal{V} \uplus \{\mathsf{x}\}) \to \mathcal{T}(\mathcal{V})$$

where from $\sigma : \mathcal{A}^\perp \parr \mathcal{B}$, $\forall_{\mathcal{V},\mathsf{x}}(\sigma) : (\forall \mathsf{x}.\, \mathcal{A})^\perp \parr \forall \mathsf{x}.\, \mathcal{B}$ plays copycat on the initial $\forall$, then proceeds as σ (and similarly for $\exists_{\mathcal{V},\mathsf{x}}(\sigma)$). Following Lawvere [18], one would expect adjunctions $\exists_{\mathcal{V},\mathsf{x}} \dashv \mathcal{T}(\mathrm{w}_{\mathcal{V},\mathsf{x}}) \dashv \forall_{\mathcal{V},\mathsf{x}}$. We will, however, not quite get this here (we present this failure later as the non-preservation of $\rightsquigarrow_{\mathrm{CUT}/\forall}$). We now interpret $\forall$I and $\exists$I.

1) Interpretation of $\exists$I: For any arena A and term $t \in \mathsf{Tm}_\Sigma(\mathcal{V})$, we first define the $(\Sigma \uplus \mathcal{V})$-strategy below.

Definition IV.5. *The $(\Sigma\uplus\mathcal{V})$-strategy $\exists^t_A : A^\perp \parallel \exists.A$ is $(|A^\perp \parallel \exists.\, A|, \leq_{\exists^t_A}, \lambda_{\exists^t_A})$ where $\leq_{\exists^t_A}$ includes $\leq_{\text{\c{c}}_A}$, plus dependencies*

$$\{((2,\exists),(2,a)) \mid a \in A\} \uplus \{((2,\exists),(1,a)) \mid \exists a_0^\forall \in A.\; a_0 \leq_A a\}$$

and term assignment that of $\text{\c{c}}_A$ plus $\lambda_{\exists^t_A}((2,\exists)) = t$.

In other words, $\exists^t_A$ plays the existential quantifier annotated with t, then proceeds as copycat on A. We have:

Proposition IV.6. *Let $\mathcal{A}$ be a $\mathcal{V}$-game, and $t \in \mathsf{Tm}_\Sigma(\mathcal{V})$. Then,*

$$\exists^t_A : \mathcal{A}[t/\mathsf{x}] \overset{\mathcal{V}\text{-}\mathrm{Ga}_\Sigma}{\twoheadrightarrow} \exists \mathsf{x}.\, \mathcal{A}$$

Indeed, any $\exists$-maximal $x_A \parallel \exists.x_A \in \mathscr{C}^\infty(\exists^t_A)$ corresponds to a tautology $\mathcal{W}_{\mathcal{A}[t/\mathsf{x}]}(x_A)^\perp \vee \mathcal{W}_{\mathcal{A}}(x_A)[t/\mathsf{x}]$. We interpret $\exists$I by post-composing with $\exists^t_A$ (as in Figure 10 without the last step). This validates $\rightsquigarrow_{\mathrm{CUT}/\exists}$, by associativity of composition.

2) Interpretation of $\forall$I: We define the following operation.

Definition IV.7. *Let σ be a $(\Sigma\uplus\mathcal{V}\uplus\{\mathsf{x}\})$-strategy on an arena $A^\perp \parallel B$. The $(\Sigma\uplus\mathcal{V})$-strategy $\forall I^{\mathsf{x}}_{A,B}(\sigma) : A^\perp \parallel \forall.B$ has events $|\sigma| \uplus \{(2,\forall)\}$, and dependency $\leq_\sigma$, plus:*

$$\{((2,\forall), s) \mid s \in \forall.B \quad \vee \quad \exists s' \leq_\sigma s.\, \mathsf{x} \in \mathrm{fv}(\lambda_\sigma(s'))\}$$

with, finally, $\lambda_{\forall I^{\mathsf{x}}_{A,B}(\sigma)}((2,\forall)) = (2,\forall)$ and $\lambda_{\forall I^{\mathsf{x}}_{A,B}(\sigma)}(s) = \lambda_\sigma(s)[(2,\forall)/\mathsf{x}]$ for $s \in |\sigma|$.

In other words we add $\forall$ as a new minimal event, which we set as a dependency for all $\exists$loïse moves with annotation comprising x. Finally, we rename x to $\forall$ in those annotations. We have the following proposition.

Proposition IV.8. *If σ is winning on a $(\mathcal{V} \uplus \{\mathsf{x}\})$-game $\mathcal{A}[\mathrm{w}_{\mathcal{V},\mathsf{x}}] \parr \mathcal{B}$, $\forall I^{\mathsf{x}}_{A,B}(\sigma)$ is winning on the $\mathcal{V}$-game $\mathcal{A} \parr \forall \mathsf{x}.\, \mathcal{B}$.*

Indeed, if $\forall$bélard does not play $(2,\forall)$ we get a tautology, otherwise the remaining configuration is one of σ and yields a tautology. This gives an interpretation for $\forall$I, which completes the interpretation of MLL_1. Just like $\rightsquigarrow_{\mathrm{CUT}/\exists}$, it validates $\rightsquigarrow_{\forall/\exists}$. However, it fails $\rightsquigarrow_{\mathrm{CUT}/\forall}$. This stems from the fact that the *minimal* Σ-strategies are not stable under composition.

Example IV.9. *Consider the* Σ*-strategies* $\sigma : \forall_1.1 \overset{\mathrm{Ar}_\Sigma}{\twoheadrightarrow} \forall_2.\forall_3.1$ *and* $\tau : \forall_2.\forall_3.1 \overset{\mathrm{Ar}_\Sigma}{\twoheadrightarrow} \forall_4.1$ *depicted below.*

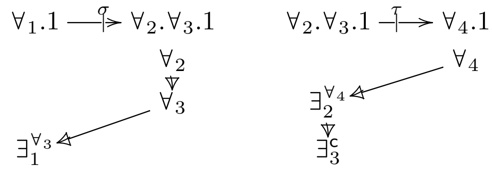

Although both Σ*-strategies are minimal, their composition* $\tau \odot \sigma$ *is* $\forall_4 \rightarrow \exists_1^{\mathsf{c}}$ *and therefore not minimal.*

In the example above, τ is some $\forall\mathrm{I}^{\mathsf{x}}_{\forall_2.\forall_3.1,1}(\tau')$ for τ' playing $\exists_2^{\mathsf{x}} \rightarrow \exists_3^{\mathsf{c}}$. But $\tau' \odot \sigma$ has for only move $\exists_1^{\mathsf{c}}$, so the dependency $\forall_4 \rightarrow \exists_1^{\mathsf{c}}$ present in $\forall\mathrm{I}^{\mathsf{x}}_{\forall_2.\forall_3.1,1}(\tau') \odot \sigma$ is not there in $\forall\mathrm{I}^{\mathsf{x}}_{\forall_1.1,1}(\tau' \odot \sigma)$, as would be required for $\rightsquigarrow_{\mathrm{CUT}/\forall}$.

The interpretation of cut-free proofs yield minimal Σ-strategies, whereas compositions interpreting cuts may create non-minimal causal dependencies as the dependency flows through the syntax tree of the cut formula. Hence, cut reduction has the effect of weakening the causal structure of the interpretation. This can be captured by:

Lemma IV.10. *Let* $\sigma : A \overset{\mathrm{Ar}_\Sigma}{\twoheadrightarrow} B$ *be a* Σ*-strategy and* $\tau : B \overset{\mathrm{Ar}_\Sigma}{\twoheadrightarrow} C$ *a* $(\Sigma \uplus \{\mathsf{x}\})$*-strategy. Then,*

$$\forall I^{\mathsf{x}}_{A,C}(\tau \odot \sigma) \preccurlyeq \forall I^{\mathsf{x}}_{B,C}(\tau) \odot \sigma$$

Recall that by Lemma III.7, the two Σ-strategies have the same term annotations on their common events. In fact, $\forall\mathrm{I}^{\mathsf{x}}_{A,C}(\tau \odot \sigma)$ and $\forall\mathrm{I}^{\mathsf{x}}_{B,C}(\tau) \odot \sigma$ also have the same *events* – they correspond to the same *expansion tree*, only the acyclicity witness differs. But a variant of $\preccurlyeq$ with $|\sigma_1| = |\sigma_2|$ would fail to be a congruence: relaxing causality of σ in $\tau \odot \sigma$ may unlock new events in a composition, previously part of causal loops.

As $\preccurlyeq$ is preserved by all operations on Σ-strategies, we deduce the main theorem of this section.

Theorem IV.11. *If* $\pi \rightsquigarrow_{\mathrm{MLL}_1} \pi'$*, then* $[\![\pi']\!] \preccurlyeq [\![\pi]\!]$.

We conjecture that this can be strengthened to $[\![\pi]\!]$ and $[\![\pi']\!]$ having the same events, by showing that compositions occuring in the interpretation of MLL_1 proofs are deadlock-free and do not create any causal cycle; then, the equivalence relation "having the same expansion tree", *i.e.* the same data apart from causal dependency, would be preserved under composition, and we would have a $\ast$*-autonomous hyperdoctrine* of Σ-strategies up to it. However, as these good properties would not hold in the presence of contraction and weakening, we leave this out as beyond the scope of the present paper.

V. Contraction and weakening

We reinstate ! and ? in the interpretation of quantifiers, *i.e.* $[\![\forall\mathsf{x}.\,\varphi]\!]_{\mathcal{V}} = \,!\forall\mathsf{x}.\,[\![\varphi]\!]_{\mathcal{V}\uplus\{\mathsf{x}\}}$ and $[\![\exists\mathsf{x}\,\varphi]\!]_{\mathcal{V}} = \,?\exists\mathsf{x}\,[\![\varphi]\!]_{\mathcal{V}\uplus\mathsf{x}}$.

Unlike for MLL_1, now we only aim to associate to any proof a Σ-winning strategy on the appropriate game, with no preservation of cut elimination. We need to provide interpretations for contraction and weakening, but also to revisit the interpretation of the introduction rules for quantifiers, as now the interpretation of formulas has changed.

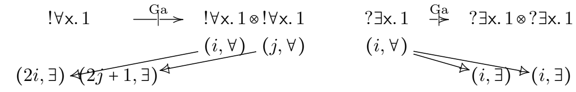

Fig. 9. Two examples of contraction

The interpretation of *weakening* is easy. For any game $\mathcal{A}$, any Σ-strategy $\sigma : \mathcal{A} \twoheadrightarrow 1$ is winning; for definiteness, we use the minimal $e_{\mathcal{A}} : \mathcal{A} \twoheadrightarrow 1$, only closed under receptivity. *Contraction* is much more subtle. To illustrate the difficulty, we present in Figure 9 two simple instances of the contraction Σ-strategy (term annotations are omitted). The first looks like the usual contraction strategy in the presence of copy indices (see *e.g.* AJM games [1]). It can be used to interpret the contraction rule on existential formulas, where it has the effect of taking the union of the different witnesses proposed. But in LK, one can also use contraction on a universal formula, which will appeal to a strategy like the second. Any witness proposed by $\forall$bélard will then have to be propagated to both branches to ensure that we are winning.

In order to define this contraction Σ-strategy along with the right tools to revisit the introduction rules for quantifiers, we will first study some properties of the exponential modalities.

1) The modalities ! *and* ?*:* Recall ! and ? from Definition II.14, both based on arena $\parallel_\omega A$. First, we examine their functorial action. Let $\sigma : A \overset{\mathrm{Ar}_\Sigma}{\twoheadrightarrow} B$. Then, $\parallel_\omega \sigma : \parallel_\omega (A^\perp \parallel B)$ which is isomorphic to $(\parallel_\omega A)^\perp \parallel (\parallel_\omega B)$; by abuse of notation (implicitly using the iso) we write $\parallel_\omega \sigma : \parallel_\omega A \overset{\mathrm{Ar}_\Sigma}{\twoheadrightarrow} \parallel_\omega B$.

Lemma V.1. *Let* $\sigma : \mathcal{A} \overset{\mathrm{Ga}_\Sigma}{\twoheadrightarrow} \mathcal{B}$*. Then, we have*

$$!\sigma = \parallel_\omega \sigma : !\mathcal{A} \overset{\mathrm{Ga}_\Sigma}{\twoheadrightarrow} !\mathcal{B} \qquad ?\sigma = \parallel_\omega \sigma : ?\mathcal{A} \overset{\mathrm{Ga}_\Sigma}{\twoheadrightarrow} ?\mathcal{B}$$

Besides, Figure 11 shows win-isos whose liftings are used in the interpretation. We also have a winning $!\mathcal{A} \twoheadrightarrow \mathcal{A}$ (though not obtained by lifting) playing copycat between $\mathcal{A}$ and the 0th copy on the left hand side, and closed under receptivity.

2) Perennialisation: Rather than defining directly the contraction $[\![\varphi]\!]_{\mathcal{V}} \twoheadrightarrow [\![\varphi]\!]_{\mathcal{V}} \otimes [\![\varphi]\!]_{\mathcal{V}}$, we will build

$$co_\varphi : [\![\varphi]\!]_{\mathcal{V}} \overset{\mathrm{Ga}_{\Sigma\uplus\mathcal{V}}}{\twoheadrightarrow} ![\![\varphi]\!]_{\mathcal{V}}$$

for every $\varphi \in \mathsf{Form}_\Sigma(\mathcal{V})$. If φ is quantifier-free, the empty Σ-strategy $co_\varphi : [\![\varphi]\!]_{\mathcal{V}} \twoheadrightarrow ![\![\varphi]\!]_{\mathcal{V}}$ is winning as $\wedge$ is idempotent. For a universal formula, we use a particular case of $!\mathcal{A} \twoheadrightarrow !!\mathcal{A}$

$$co_{\forall\mathsf{x}.\,\varphi} : !\forall\mathsf{x}.\,[\![\varphi]\!]_{\mathcal{V}} \twoheadrightarrow !!\forall\mathsf{x}.\,[\![\varphi]\!]_{\mathcal{V}}$$

from Figure 11. We get $co_{\varphi\wedge\psi}$ and $co_{\varphi\vee\psi}$ by induction and composition with $!\mathcal{A} \otimes !\mathcal{B} \twoheadrightarrow !(\mathcal{A} \otimes \mathcal{B})$, $!\mathcal{A} \parr !\mathcal{B} \twoheadrightarrow !(\mathcal{A} \parr \mathcal{B})$.

The only case left is existential quantification; it is also the most subtle, as it is analogous to the contraction of existential formulas presented on the right hand side of Figure 9.

Lemma V.2. *For any* $(\mathcal{V} \uplus \{\mathsf{x}\})$*-game* $\mathcal{A}$*, there is a winning*

$$\mu_{\mathcal{A},\mathsf{x}} : \exists\mathsf{x}.\,!\mathcal{A} \overset{\mathcal{V}\text{-Games}}{\twoheadrightarrow} !\exists\mathsf{x}.\,\mathcal{A}\,.$$

$$\left[\!\!\left[\mathrm{C}\ \frac{\dfrac{\pi}{\vdash^{\mathcal{V}} \Gamma, \varphi, \varphi}}{\vdash^{\mathcal{V}} \Gamma, \varphi}\right]\!\!\right] = \Gamma^{\perp} \overset{[\![\pi]\!]}{\underset{\mathcal{T}(\mathcal{V})}{\to}} \varphi \parr \varphi \overset{\delta^{\perp}_{\varphi^{\perp}}}{\underset{\mathcal{T}(\mathcal{V})}{\to}} \varphi \qquad \left[\!\!\left[\forall\mathrm{I}\ \frac{\dfrac{\pi}{\vdash^{\mathcal{V} \uplus \{\mathsf{x}\}} \Gamma, \varphi}}{\vdash^{\mathcal{V}} \Gamma, \forall \mathsf{x}.\, \varphi}\right]\!\!\right] = \Gamma^{\perp} \overset{co_{\Gamma^{\perp}}}{\underset{\mathcal{T}(\mathcal{V})}{\to}} !\Gamma^{\perp} \overset{!(\forall\mathrm{I}\,([\![\pi]\!]))}{\underset{\mathcal{T}(\mathcal{V})}{\to}} !\forall \mathsf{x}.\, \varphi \qquad \left[\!\!\left[\forall\mathrm{I}\ \frac{\dfrac{\pi}{\vdash^{\mathcal{V}} \Gamma, \varphi[t/\mathsf{x}]}}{\vdash^{\mathcal{V}} \Gamma, \exists \mathsf{x}.\, \varphi}\right]\!\!\right] = \Gamma^{\perp} \overset{[\![\pi]\!]}{\underset{\mathcal{T}(\mathcal{V})}{\to}} \varphi[t/\mathsf{x}] \overset{\exists^{t}_{\varphi}}{\underset{\mathcal{T}(\mathcal{V})}{\to}} \exists \mathsf{x}.\varphi \overset{\mathcal{T}(\mathcal{V})}{\to} ?\exists \mathsf{x}.\varphi$$

Fig. 10. Interpretation of the remaining rules of LK

$$\begin{array}{rcl@{\quad}rcl@{\quad}rcl}
!\mathcal{A} & \to & !!\mathcal{A} & !\mathcal{A} & \to & !\mathcal{A} \otimes !\mathcal{A} & ?!\mathcal{A} & \to & !?\mathcal{A} \\
(\langle i, j \rangle, a) & \mapsto & (i, (j, a)) & (2i, a) & \mapsto & (1, (i, a)) & (i, (j, a)) & \mapsto & (j, (i, a)) \\
!\mathcal{A} \otimes !\mathcal{B} & \to & !(\mathcal{A} \otimes \mathcal{B}) & (2i+1, a) & \mapsto & (2, (i, a)) & !\mathcal{A} \parr !\mathcal{B} & \to & !(\mathcal{A} \parr \mathcal{B}) \\
(j, (i, a)) & \mapsto & (i, (j, a)) & & & & (j, (i, a)) & \mapsto & (i, (j, a))
\end{array}$$

Fig. 11. Some win-isos with exponentials

Proof. After the unique minimal $\forall$ move (on the left hand side), the strategy simultaneously plays all the $(i, \exists)$ (on the right hand side) with annotation $\forall$; then proceeds as $\mathrm{c\!c}_{!\mathcal{A}}$. □

Using that, we can define $co_{?\exists \mathsf{x}.\, [\![\varphi]\!]_{\mathsf{x}}}$ as the composition

$$?\exists \mathsf{x}.\, [\![\varphi]\!] \overset{?\exists_{\mathcal{V},\mathsf{x}} co_{[\![\varphi]\!]}}{\twoheadrightarrow} ?\exists \mathsf{x}.\, ![\![\varphi]\!] \overset{?\mu_{[\![\varphi]\!],\mathsf{x}}}{\twoheadrightarrow} ?!\exists \mathsf{x}.\, [\![\varphi]\!] \twoheadrightarrow !?\exists \mathsf{x}.\, [\![\varphi]\!]\,.$$

Summing up, we have proved:

Proposition V.3. *For any $\varphi \in \mathsf{Form}_{\Sigma}(\mathcal{V})$, there is a winning*

$$co_{[\![\varphi]\!]_{\mathcal{V}}} : [\![\varphi]\!]_{\mathcal{V}} \overset{\mathcal{V}\text{-Games}}{\twoheadrightarrow} ![\![\varphi]\!]_{\mathcal{V}}\,.$$

3) Completing the interpretation: Using Proposition V.3 along with the basic Σ-strategies defined above, we build $\delta_{[\![\varphi]\!]_{\mathcal{V}}} : [\![\varphi]\!]_{\mathcal{V}} \twoheadrightarrow [\![\varphi]\!]_{\mathcal{V}} \otimes [\![\varphi]\!]_{\mathcal{V}}$ for every $\varphi \in \mathsf{Form}_{\Sigma}(\mathcal{V})$. Using it we give in Figure 10 the final clauses of the interpretation of LK (omitting W, which is by post-composition with $e_{\mathcal{A}}$), completing Figure 6 – again, for readability we omit the semantic brackets on formulas and silently use the isomorphism between winning Σ-strategies from 1 to $\Gamma \parr \mathcal{A}$ and from $\Gamma^{\perp}$ to $\mathcal{A}$ (due to the $\star$-autonomous structure of each fibre). This concludes the interpretation, and the proof of Theorem II.17.

VI. Conclusion

For LK there is no hope of preserving unrestricted cut reduction without collapsing the model to a boolean algebra [12]. There are non-degenerate categorical models for classical logic with an involutive negation, *e.g.* Führman and Pym's *classical categories* [9] where cut reduction is only preserved in a lax sense; but our model does not preserve cut reduction even in this weaker sense. Besides non-preservation of cut elimination, the interpretation is infinitary: from the example in [8] of a LK proof with arbitrarily large cut-free forms, one can construct a proof of some $\exists \mathsf{x}.\, \varphi$ with φ quantifier-free (*i.e.*, no $\forall$bélard moves at all) yielding an infinite Σ-strategy.

Both phenomena are due to our commitment to preserve the symmetry and involutive negation of classical logic – they do not appear in *e.g.* Laurent's games model for the first-order $\lambda\mu$-calculus. It is a fascinating open question whether one can find a non-degenerate finitary interpretation which avoids the sequentiality resulting from the negative translation.

Acknowledgments

We gratefully acknowledge the support of the ERC Advanced Grant ECSYM.

References

[1] Samson Abramsky, Radha Jagadeesan, and Pasquale Malacaria. Full abstraction for PCF. *Inf. Comput.*, 163(2):409–470, 2000.

[2] Michael Barr. *-autonomous categories and linear logic. *Mathematical Structures in Computer Science*, 1(2):159–178, 1991.

[3] Samuel R. Buss. On Herbrand's theorem. In Daniel Leivant, editor, *Logical and Computational Complexity. Selected Papers. Logic and Computational Complexity, International Workshop LCC '94, Indianapolis, Indiana, USA, 13-16 October 1994*, volume 960 of *Lecture Notes in Computer Science*, pages 195–209. Springer, 1994.

[4] Simon Castellan, Pierre Clairambault, Silvain Rideau, and Glynn Winskel. Games and strategies as event structures. 2016. Submitted, https://hal.inria.fr/hal-01302713.

[5] Pierre Clairambault, Julian Gutierrez, and Glynn Winskel. The winning ways of concurrent games. In *Proceedings of the 27th Annual IEEE Symposium on Logic in Computer Science, LICS 2012, Dubrovnik, Croatia, June 25-28, 2012*, pages 235–244. IEEE, 2012.

[6] J.R.B. Cockett and R.A.G. Seely. Weakly distributive categories. *Journal of Pure and Applied Algebra*, 114(2):133–173, 1997.

[7] Thierry Coquand. A semantics of evidence for classical arithmetic. *J. Symb. Log.*, 60(1):325–337, 1995.

[8] Vincent Danos, Jean-Baptiste Joinet, and Harold Schellinx. A new deconstructive logic: Linear logic. *J. Symb. Log.*, 62(3):755–807, 1997.

[9] Carsten Führmann and David J. Pym. On categorical models of classical logic and the geometry of interaction. *Mathematical Structures in Computer Science*, 17(5):957–1027, 2007.

[10] Philipp Gerhardy and Ulrich Kohlenbach. Extracting herbrand disjunctions by functional interpretation. *Arch. Math. Log.*, 44(5):633–644, 2005.

[11] Jean-Yves Girard. Linear logic. *Theor. Comput. Sci.*, 50:1–102, 1987.

[12] Jean-Yves Girard. A new constructive logic: Classical logic. *Mathematical Structures in Computer Science*, 1(3):255–296, 1991.

[13] Kurt Gödel. Über eine bisher noch nicht benützte erweiterung des finiten standpunktes. *dialectica*, 12(3-4):280–287, 1958.

[14] Willem Heijltjes. Classical proof forestry. *Ann. Pure Appl. Logic*, 161(11):1346–1366, 2010.

[15] Stefan Hetzl and Daniel Weller. Expansion trees with cut. *CoRR*, abs/1308.0428, 2013.

[16] Ulrich Kohlenbach. On the no-counterexample interpretation. *J. Symb. Log.*, 64(4):1491–1511, 1999.

[17] Olivier Laurent. Game semantics for first-order logic. *Logical Methods in Computer Science*, 6(4), 2010.

[18] F William Lawvere. Adjointness in foundations. *Dialectica*, 23(3-4):281–296, 1969.

[19] Richard McKinley. Proof nets for Herbrand's theorem. *ACM Trans. Comput. Log.*, 14(1):5, 2013.

[20] Paul-André Melliès and Samuel Mimram. Asynchronous games: Innocence without alternation. In Luís Caires and Vasco Thudichum Vasconcelos, editors, *CONCUR*, volume 4703 of *LNCS*, pages 395–411. Springer, 2007.

[21] Dale A Miller. A compact representation of proofs. *Studia Logica*, 46(4):347–370, 1987.

[22] Samuel Mimram. The structure of first-order causality. *Mathematical Structures in Computer Science*, 21(1):65–110, 2011.

[23] Silvain Rideau and Glynn Winskel. Concurrent strategies. In *Proceedings of the 26th Annual IEEE Symposium on Logic in Computer Science, LICS 2011, June 21-24, 2011, Toronto, Ontario, Canada*, pages 409–418. IEEE Computer Society, 2011.

[24] Robert A. G. Seely. Hyperdoctrines, natural deduction and the Beck condition. *Math. Log. Q.*, 29(10):505–542, 1983.

[25] J.R. Shoenfield. *Mathematical Logic*. Addison-Wesley, 1967.

[26] Glynn Winskel. Event structures. In Wilfried Brauer, Wolfgang Reisig, and Grzegorz Rozenberg, editors, *Advances in Petri Nets*, volume 255 of *LNCS*, pages 325–392. Springer, 1986.

APPENDIX

A. The interpretation is infinitary

In this appendix, we construct an LK proof of the formula $\exists x.\top$ whose interpretation is infinite, despite the fact that there is no move by ∀bélard in the game. Besides showing that the interpretation is infinitary, we also take advantage of the presentation of the example to detail as much as reasonable the interpretation, so that the interested reader can see it at play in a non-trivial case.

Our starting point is the following proof:

$$\varpi_1 = \dfrac{\dfrac{\dfrac{\dfrac{}{\vdash \varphi,\varphi^\perp}\textsc{Ax} \quad \dfrac{}{\vdash \varphi,\varphi^\perp}\textsc{Ax}}{\vdash \varphi\wedge\varphi,\varphi^\perp,\varphi^\perp}\wedge\mathrm{I}}{\vdash \varphi\wedge\varphi,\varphi^\perp}\mathrm{C} \qquad \dfrac{\dfrac{\dfrac{}{\vdash \varphi,\varphi^\perp}\textsc{Ax} \quad \dfrac{}{\vdash \varphi,\varphi^\perp}\textsc{Ax}}{\vdash \varphi,\varphi,\varphi^\perp\wedge\varphi^\perp}\wedge\mathrm{I}}{\vdash \varphi,\varphi^\perp\wedge\varphi^\perp}\mathrm{C}}{\vdash \varphi\wedge\varphi,\varphi^\perp\wedge\varphi^\perp}\textsc{Cut}$$

This proof is referred to in [8] as a *structural dilemma*. There are two ways to push the CUT beyond contraction, as the two proofs interact, and try to duplicate one another. This is often used as an example of a proof where unrestricted cut reduction does not necessarily terminate; and which has infinitely large cut-free forms.

In order to construct a proof with an infinite interpretation, we will start with this proof, with $\varphi = \forall x.\bot \vee \exists y.\top$, which to shorten notations we will just write as $\forall \vee \exists$.

1) Interpretation of ϖ_1: We detail the interpretation of ϖ_1. We start from the axioms on the left branch:

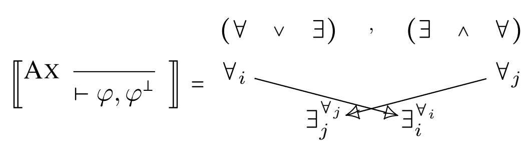

The indices i, j are the copy indices for the ! and ? arising from the interpretation of formulas, and we only display the term annotations for ∃loïse's moves. The Σ-strategy above is the copycat Σ-strategy as defined in Definition III.8.

Interpreting the introduction rule for ∧ simply has the effect of tensoring two copies of copycat together, obtaining:

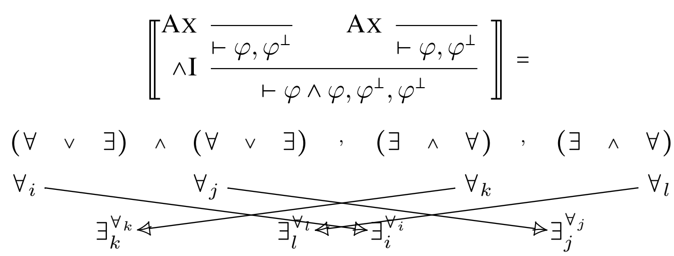

i.e. again copycat, in accordance with the functoriality of ⊗.

Now, to interpret contraction, we need to compose with $\delta^\perp_{\forall\vee\exists} : (\exists\wedge\forall)\vee(\exists\wedge\forall) \twoheadrightarrow \exists\wedge\forall$, where

$$\delta_{\forall\vee\exists} : (!\forall \parr ?\exists) \twoheadrightarrow (!\forall \parr ?\exists) \otimes (!\forall \parr ?\exists)$$

is the contraction on φ. Note that this time, make explicit the exponential modalities. Recall also that this strategy is derived from $co_{\forall\vee\exists} : (!\forall\parr?\exists) \twoheadrightarrow !(!\forall\parr?\exists)$, which we display below. To display it best we deviate from the representation below by showing exactly the correspondence between copy indices and occurrences of ! and ?, and we omit the terms, which are trivial and always correspond with the unique predecessor for ∃loïse's events. We display the Σ-strategy separating two sub-configurations for clarity; the full Σ-strategy is obtained by taking their union.

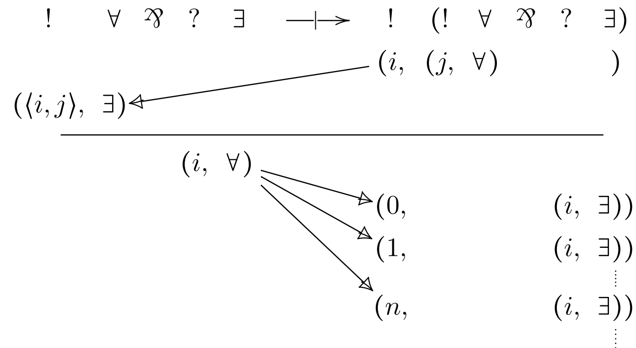

We do not detail the construction of this Σ-strategy, but it is easy to get from the definitions. This Σ-strategy $co_{\forall\vee\exists}$ obviously performs an infinitary duplication, however it does not show by itself that the interpretation is infinitary, as $co_{\forall\vee\exists}$ is just an auxiliary device in the definition of the interpretation, rather than itself the interpretation of a proof.

To get contraction on φ from $co_{\forall\vee\exists}$, we compose it with the derelicted version of contraction on $!\varphi$:

$$!(!\forall\parr?\exists) \twoheadrightarrow !(!\forall\parr?\exists)\otimes!(!\forall\parr?\exists) \twoheadrightarrow (!\forall\parr?\exists)\otimes(!\forall\parr?\exists)$$

which we display here:

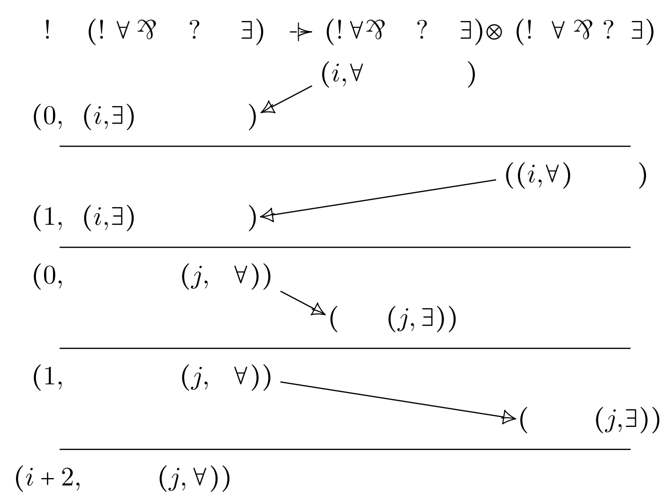

where the final case is just closure under receptivity. Performing the composition, we get the contraction Σ-strategy $\delta_{\forall\vee\exists}$:

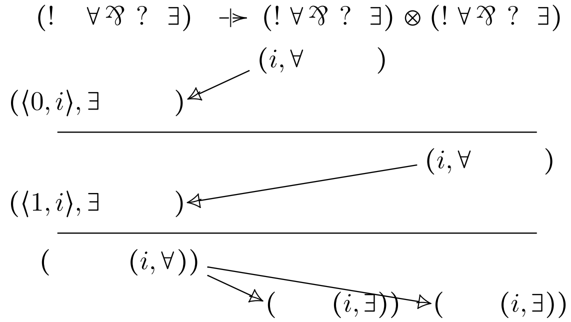

With that in place, we can finally obtain by composition (where we adopt again the simplified annotation for copy indices, since in this games ! and ? are again always attached to quantifiers – we still omit the trivial term annotations):

$$\left[\!\!\left[\dfrac{\dfrac{\dfrac{}{\vdash \varphi, \varphi^{\perp}}\,\text{Ax} \qquad \dfrac{}{\vdash \varphi, \varphi^{\perp}}\,\text{Ax}}{\vdash \varphi \wedge \varphi, \varphi^{\perp}, \varphi^{\perp}}\,\wedge\text{I}}{\vdash \varphi \wedge \varphi, \varphi^{\perp}}\,\text{C} \right]\!\!\right] =$$

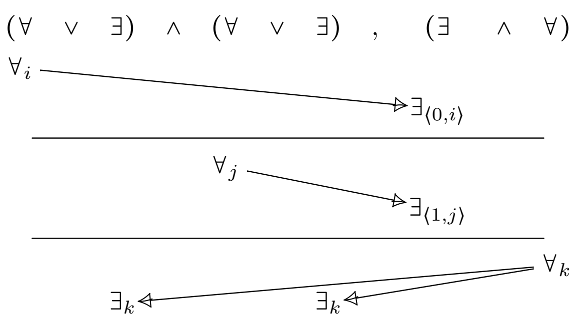

The second branch of ϖ_1 is symmetric, so we do not make it explicit. Now, we interpret the CUT rule and the composition yields $[\![\varpi_1]\!]$ below (again, we omit term annotations which coincide with the unique predecessor for ∃loïse's moves).

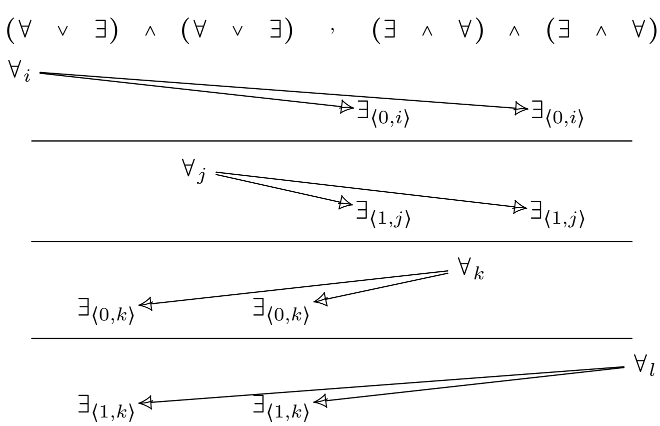

It is interesting to note that although ϖ_1 has arbitrarily large cut-free forms, the corresponding strategy only plays finitely many ∃loïse moves for every ∀bélard move. However, we are on the right path to finding a truly infinitary Σ-strategy.

2) An infinitary proof: The next step is to set (with s some unary function symbol):

$$\varpi_2 = \dfrac{\dfrac{\dfrac{\dfrac{\dfrac{}{\vdash^{\mathsf{x}} \top[\mathsf{s}(\mathsf{x})/\mathsf{y}], \bot}\,\text{Ax}}{\vdash^{\mathsf{x}} \exists \mathsf{y}.\, \top, \bot}\,\exists\text{I}}{\vdash \exists \mathsf{y}.\, \top, \forall \mathsf{x}.\, \bot}\,\forall\text{I}}{\vdash \forall \mathsf{x}.\, \bot, \exists \mathsf{y}.\, \top, \forall \mathsf{x}.\, \bot, \exists \mathsf{x}.\, \top}\,\text{W}}{\vdash (\forall \mathsf{x}.\, \bot \vee \exists \mathsf{y}.\, \top) \vee (\forall \mathsf{x}.\, \bot \vee \exists \mathsf{x}.\, \top)}\,\vee\text{I}$$

Leaving to the reader the details of the interpretation, we have by design that $[\![\varpi_2]\!]$ is:

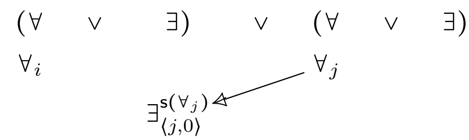

We now use these to compute the interpretation of:

$$\varpi_3 = \dfrac{\dfrac{\varpi_1}{\vdash \varphi \wedge \varphi, \varphi^{\perp} \wedge \varphi^{\perp}} \qquad \dfrac{\varpi_2}{\vdash (\forall \vee \exists) \vee (\forall \vee \exists)}}{\vdash \varphi \wedge \varphi}\,\text{CUT}$$

The associated composition reveals $[\![\varpi_3]\!]$ to be:

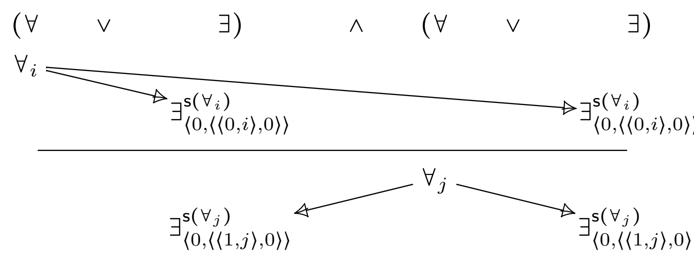

We are almost there. It suffices now to note that ϖ_3 provides a proof of

$$(\exists \mathsf{x}.\, \top \implies \exists \mathsf{x}.\, \top) \wedge (\exists \mathsf{x}.\, \top \implies \exists \mathsf{x}.\, \top)$$

These two implications can be *composed* by cutting ϖ_3 against the proof ϖ_4 or $(\exists \Rightarrow \exists) \wedge (\exists \Rightarrow \exists) \Rightarrow (\exists \Rightarrow \exists)$ performing the composition:

$$\varpi_4 = \dfrac{\dfrac{\dfrac{\dfrac{\dfrac{}{\vdash \forall, \exists}\,\text{Ax} \qquad \dfrac{}{\vdash \forall, \exists}\,\text{Ax}}{\vdash \forall, \exists \wedge \forall, \exists}\,\wedge\text{I} \qquad \dfrac{}{\vdash \forall, \exists}\,\text{Ax}}{\vdash \forall, \exists \wedge \forall, \exists \wedge \forall, \exists}\,\wedge\text{I}}{\vdash \exists \wedge \forall, \exists \wedge \forall, \exists, \forall}\,\text{Ex}}{\vdash (\exists \wedge \forall) \vee (\exists \wedge \forall), \exists \vee \forall}\,\vee\text{I}$$

with interpretation:

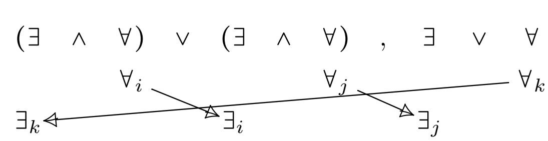

Write ϖ_5 for the proof of $\exists \mathsf{x}.\, \top \vee \forall \mathsf{y}.\, \bot$ obtained by cutting ϖ_3 and ϖ_4 in the obvious way. The interpretation of ϖ_5 is the composition of $[\![\varpi_3]\!]$ and $[\![\varpi_4]\!]$, which triggers the feedback loop causing the infiniteness phenomenon. We display below the corresponding interaction. For the "synchronised" part of formulas, we will use 0 for components resulting from matching dual quantifiers, and $\|$ for components resulting for matching dual propositional connectives. We write $\circ$ for synchronized events (*i.e.* of neutral polarity), and omit copy indices, which get very unwieldy. For readability, we also annotate the immediate causal links with the sub-proof that they originate from, *i.e.* ϖ_3 or ϖ_4.

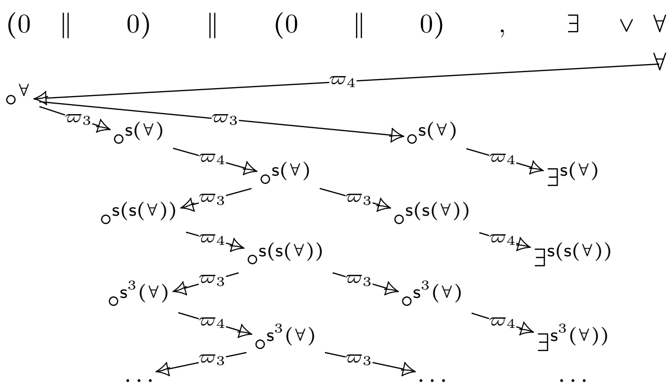

Therefore, after hiding, ∃loïse responds to an initial ∀bélard move ∀ by playing simultaneously all $\exists^{\mathsf{s}^n(\forall)}$, for $n \geq 1$. Finally, cutting ϖ_5 against a proof of $\exists \mathsf{x}.\top$ playing a constant symbol 0, we get a proof ϖ_6 of $\vdash \exists \mathsf{x}.\top$ whose interpretation plays simultaneously all $\exists^{\mathsf{s}^n(0)}$ for $n \geq 1$.

B. Neutral Σ-strategies have no meet

We noted in Section III-A1 that although neutral strategies (without taking terms into account) always had a meet (this is, in essence, the *pullback* construction in the category of event structures), neutral Σ-*strategies* in general do not have all meets.

Indeed, consider the following example:

Example B.1. *If A has two incomparable a_1 and a_2, then the neutral Σ-strategies $a_1^{a_1} \twoheadrightarrow a_2^{a_1}$ and $a_1^{a_1} \quad a_2^{\mathsf{c}}$ have no meet.*

Let us assume it has a meet $q = (|q|, \leq_q, \lambda_q)$. First, let us find out what its events would be. Since it is a lower bound, we have $\mathscr{C}(q) \subseteq \mathscr{C}(a_1^{a_1} \twoheadrightarrow a_2^{a_1})$ and $\mathscr{C}(q) \subseteq \mathscr{C}(a_1^{a_1} \quad a_2^{\mathsf{c}})$.

But we also note that:

$$a_1^{\mathsf{c}} \twoheadrightarrow a_2^{\mathsf{c}} \preccurlyeq a_1^{a_1} \twoheadrightarrow a_2^{a_1} \qquad a_1^{\mathsf{c}} \twoheadrightarrow a_2^{\mathsf{c}} \preccurlyeq a_1^{a_1} \quad a_2^{\mathsf{c}}$$

So $a_1^{\mathsf{c}} \twoheadrightarrow a_2^{\mathsf{c}}$ must be below the meet, whose events must therefore include a_1 and a_2. Hence, we have $|q| = \{a_1, a_2\}$.

Let us now examine the labelling. Because $q \preccurlyeq a_1^{a_1} \quad a_2^{\mathsf{c}}$, we must have $\lambda_q(a_2) = \mathsf{c}$. But then because $q \preccurlyeq a_1^{a_1} \twoheadrightarrow a_2^{a_1}$, we must have $\lambda_q(a_1) = \lambda_q(a_2) = \mathsf{c}$ as well.

But then, we observe that $a_1^{a_1} \preccurlyeq a_1^{a_1} \twoheadrightarrow a_2^{a_1}$ and $a_1^{a_1} \preccurlyeq a_1^{a_1} \quad a_2^{\mathsf{c}}$, so we should have $a_1^{a_1} \preccurlyeq q$ as well if q is to be the meet. But that not the case, since c does not subsume the variable a_1. Therefore, $a_1^{a_1} \twoheadrightarrow a_2^{a_1}$ and $a_1^{a_1} \quad a_2^{\mathsf{c}}$ have no meet.