



HAL
open science

The True Concurrency of Herbrand's Theorem

Aurore Alcolei, Pierre Clairambault, Martin Hyland, Glynn Winskel

► **To cite this version:**

Aurore Alcolei, Pierre Clairambault, Martin Hyland, Glynn Winskel. The True Concurrency of Herbrand's Theorem. 27th EACSL Annual Conference on Computer Science Logic (CSL 2018), Sep 2018, Birmingham, United Kingdom. 10.4230/LIPIcs.CSL.2018.5 . hal-01886968

HAL Id: hal-01886968

<https://hal.science/hal-01886968v1>

Submitted on 3 Oct 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

The True Concurrency of Herbrand's Theorem

Aurore Alcolei

Univ Lyon, ENS de Lyon, CNRS, UCB Lyon 1, LIP
Lyon, France
Aurore.Alcolei@ens-lyon.fr

Pierre Clairambault

Univ Lyon, CNRS, ENS de Lyon, UCB Lyon 1, LIP
Lyon, France
Pierre.Clairambault@ens-lyon.fr

Martin Hyland

DPMMS, University of Cambridge
Cambridge, United Kingdom
M.Hyland@dpmms.cam.ac.uk

Glynn Winskel

Computer Laboratory, University of Cambridge
Cambridge, United Kingdom
Glynn.Winskel@cl.cam.ac.uk

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. In the general case, it reduces 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 and proved as a theorem in the framework of concurrent games, a denotational semantics designed to faithfully represent causality and independence in concurrent systems, thereby exposing the concurrency underlying the computational content of classical proofs. The causal structure of concurrent strategies, paired with annotations by first-order terms, is used to specify the dependency between quantifiers implicit in proofs. 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.

2012 ACM Subject Classification Theory of computation → Proof theory, Theory of computation → Denotational semantics

Keywords and phrases Herbrand's theorem, Game semantics, True concurrency

Digital Object Identifier 10.4230/LIPIcs.CSL.2018.5

Acknowledgements We acknowledge support of the French LABEX MILYON (ANR-10-LABX-0070), the ERC Advanced Grant ECSYM and the Collegium de Lyon.



© Aurore Alcolei, Pierre Clairambault, Martin Hyland, and Glynn Winskel;
licensed under Creative Commons License CC-BY

27th EACSL Annual Conference on Computer Science Logic (CSL 2018).

Editors: Dan Ghica and Achim Jung; Article No. 5; pp. 5:1–5:22

Leibniz International Proceedings in Informatics



Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 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 that direction. It is a simple consequence of completeness and compactness in first-order logic. 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 $\vdash A \rightarrow B$ 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 [17] which shows that one cannot hope directly to use collections of Herbrand terms for $\vdash A$ and $\vdash A \rightarrow B$ to give a collection for $\vdash B$. That leaves the possibility of making some richer data compositional, realised indirectly in Gerhardy and Kohlenbach [11] with data provided by Shoenfield's version [30] of Gödel's Dialectica Interpretation [14]. Now functional interpretations make no pretence to be faithful to the structure of proofs as encapsulated in systems like the sequent calculus: they explore in a sequential order terms proposed by a proof as witnesses for existential quantifiers, but this order is certainly not intrinsic to the proof. Thus it is compelling to seek some compositional form of Herbrand's Theorem faithful to the structure of proofs and to the dependency between terms; for cut-free proofs, Miller's *expansion trees* [24] capture precisely this “Herbrand content” (the information pertaining to quantifier instantiations), but not compositionally.

In this paper, we provide such a compositional form of Herbrand's theorem, presented as a game semantics for first-order classical logic. Our games have two players, both playing on the quantifiers of a formula φ . Eloïse, playing the existential quantifiers, defends the validity of φ . \forall bélar, playing the universal quantifiers, attempts to falsify it. This understanding of formulas as games is folklore in mathematical logic and computer science. However, like functional interpretations, such games are usually sequential [7, 19]. In contrast, our model captures the exact dependence and independence between quantifiers. To achieve that we build on *concurrent/asynchronous* games [23, 27, 4], which marry game semantics with the so-called *true concurrency* approach to models of concurrent systems, and avoid interleavings. So in a formal sense, our model highlights a parallelism inherent to classical proofs. In essence, our strategies are close to expansion trees enriched with an explicit acyclicity witness.

The computational content of classical logic is a longstanding active topic, with a wealth of related works, and it is hard to do it justice in this short introduction. There are, roughly speaking, two families of approaches. On the one hand, some (including the functional interpretations mentioned above) extract from proofs a sequential procedure, *e.g.* via translation to sequential calculi or by annotating a proof to sequentialize or determinize its behaviour under cut reduction [13, 8]. Other than that cited above, influential developments in this “polarized” approach include work by Berardi [2], Coquand [7], Parigot [26], Krivine [18], and others. Polarization yields better-behaved dynamics and a non-degenerate equational theory, but distorts the intent of the proof by an added unintended sequentiality. On the other hand, some works avoid polarization – including, of course, Gentzen's *Hauptsatz* [10]. This causes issues, notably unrestricted cut reduction yields a degenerate equational theory [13] and enjoys only *weak*, rather than *strong*, normalization [8]. Nevertheless, witness extraction

remains possible (though it is non-deterministic). Particularly relevant to our endeavour is a recent activity around the matter of enriching expansion trees so as to support cuts. This includes Heijltjes' *proof forests* [15], McKinley's *Herbrand nets* [21], and Hetzl and Weller's recent *expansion trees with cuts* [16]. 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 and based on rewriting.

Other related works include Laurent's model for the first-order $\lambda\mu$ -calculus [19], whose annotation of moves via first-order terms is similar to ours; and Mimram's categorical presentation of a games model for a linear first-order logic without propositional connectives [25].

Since our model avoids polarization, some phenomena from the proof theory of classical logic reflect in it: our semantics does not preserve cut reduction – if it did, it would be a boolean algebra [13]. Yet it preserves it in a sense for *first-order MLL* [12]. Likewise, just as classical proofs can lead to arbitrary large cut-free proofs [8], our semantics may yield *infinite* strategies, from which *finite* sub-strategies can nonetheless always be extracted. This reflects that non-polarized proof systems for classical logic are often only weakly normalizing.

In Section 2 we recall Herbrand's theorem, and introduce the game-theoretic language leading to our compositional reformulation of it. The rest of the paper describes the interpretation of proofs as winning strategies: in Section 3 we give the interpretation of propositional MLL, in Section 4 we deal with quantifiers, and finally, in Section 5, we add contraction and weakening and complete the interpretation.

2 From Herbrand to winning Σ -strategies


A **signature** is $\Sigma = (\Sigma_f, \Sigma_p)$, with Σ_f a countable set of **function symbols** (f, g, h , etc. range over function symbols), and Σ_p a countable set of **predicate symbols** (P, Q , etc. range over predicate symbols). There is an **arity function** $\text{ar} : \Sigma_f \uplus \Sigma_p \rightarrow \mathbb{N}$ where \uplus is the usual set-theoretic union, where 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 a, b, c, \dots to range over constant symbols.

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


2.1 Herbrand's theorem

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


► **Theorem 1.** *Let \mathcal{T} be a theory finitely axiomatized by universal formulas. Let $\psi = \exists x_1 \dots \exists x_n \varphi(x_1, \dots, x_n)$ be a purely existential formula ($\varphi \in \text{QF}_\Sigma$). Then, $\mathcal{T} \models \psi$ iff there are closed terms $(t_{i,j})_{1 \leq i \leq p, 1 \leq j \leq n}$ such that $\mathcal{T} \models \bigvee_{i=1}^p \varphi(t_{i,1}, \dots, t_{i,n})$.*




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



■ **Figure 2** A partially ordered winning Σ -strategy.



■ **Figure 3** An incorrect expansion tree.



■ **Figure 4** The arena $[DF]^3$.

► **Example 2.** Consider the formula $\psi = \exists x \neg P(x) \vee P(f(x))$ (where $f \in \Sigma_f$). A valid Herbrand disjunction for ψ is $(\neg P(c) \vee P(f(c))) \vee (\neg P(f(c)) \vee P(f(f(c))))$ where c is some constant symbol.

A similar disjunction property holds for general formulas, though it is harder to state. 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* (DF): $\exists x \forall y \neg P(x) \vee P(y)$, yields by Herbrandization the formula ψ of Example 2.

Instead, to avoid prenexification and Skolemization and the corresponding distortion of the formula, one may adopt a representation of proofs that displays the instantiation of existential quantifiers with finitely many witnesses while staying structurally faithful to the original formula. To that end Miller proposes **expansion trees** [24]. They can be introduced via a game-theoretic metaphor, reminiscent of [7]. Two players, \exists loïse and \forall bélar, debate the validity of a formula. On a formula $\forall x \varphi$, \forall bélar provides a fresh variable x and the game keeps going on φ . On $\exists x \varphi$, \exists loïse provides a *term* t , possibly containing variables previously introduced by \forall bélar. \exists 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: \exists loïse plays c , then \forall bélar introduces y , then \exists loïse *backtracks* (we jump to the right branch) and plays y , and finally \forall bélar introduces z . \exists loïse wins: the disjunction of the leaves is a tautology.

However the metaphor has limits, it suggests a sequential ordering between branches, which expansion trees do not have in reality: the order is only implicit in the term annotations. Besides, the natural ordering between quantifiers induced by terms is not always sequential. It is, of course, always acyclic – on expansion trees this is ensured by an *acyclicity correctness criterion*, whose necessity is made obvious by the (incorrect) expansion tree of Figure 3 “proving” a falsehood. This acyclicity entails the existence of a sequentialization, but committing to one is an arbitrary choice not forced by the proof.

A partial order is much more faithful to the proof. In this paper, we show that expansion trees can be made compositional modulo a change of perspective: rather than derived we consider this order primitive, and only later decorate it with term annotations. For instance,

we display in Figure 2 the formal object, called a (sequential) *winning Σ -strategy*, matching in our framework the expansion tree for DF. Another winning Σ -strategy, displayed in Figure 2, illustrates that this order is not always naturally sequential. By lack of space we do not define expansion trees here, though they are captured in essence by our strategies.

2.2 Expansion trees as winning Σ -strategies

We now introduce 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. Σ -strategies will play on *games* representing the formulas. The first component of a game is its *arena*, that specifies the causal ordering between quantifiers.

► **Definition 3.** An **arena** is $A = (|A|, \leq_A, \text{pol}_A)$ 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, $\text{pol}_A : |A| \rightarrow \{\forall, \exists\}$ is a **polarity function** which expresses if a move belongs to Eloïse or \forall bélar.

A **configuration** of an arena (or any partial order) is a down-closed set of events. We write $\mathcal{C}^\infty(A)$ for the set of configurations of A , and $\mathcal{C}(A)$ 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, 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 4 shows the \exists -biased arena $\llbracket DF \rrbracket^\exists$ for DF. The order is drawn from top to bottom. Although only \exists loïse can replicate her moves, the universal quantifier is also copied as it depends on the existential quantifier.

Strategies on an 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.

For any partial order A and $a_1, a_2 \in |A|$, we write $a_1 \rightarrow_A a_2$ (or $a_1 \rightarrow a_2$ if A is clear from the context) if $a_1 <_A a_2$ with no other event in between – this notation was used implicitly in Figures 1 and 2. We call \rightarrow **immediate causal dependency**.

► **Definition 4.** 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 $\mathcal{C}^\infty(\sigma) \subseteq \mathcal{C}^\infty(A)$,
- (2) *Receptivity.* If $x \in \mathcal{C}(\sigma)$ s.t. $x \cup \{a^\forall\} \in \mathcal{C}(A)$, then $a \in |\sigma|$,
- (3) *Courtesy.* If $a_1 \rightarrow_\sigma a_2$ and $(\text{pol}(a_1) = \exists \text{ or } \text{pol}(a_2) = \forall)$, then $a_1 \rightarrow_A a_2$.

These strategies are essentially the *receptive ingenuous strategies* of Melliès and Mimram [23], though their formulation, with a direct handle on causality, is closer to Rideau and Winskel's later *concurrent strategies* [27]. Receptivity means that \exists loïse cannot refuse to acknowledge a move by \forall bélar, 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 2 (right) displays a strategy on the arena of Figure 4 – in Figure 2 we also show via dotted lines the immediate dependency of the arena.


Let us now add terms, and define Σ -strategies.

► **Definition 5.** A Σ -strategy on arena A is a strategy $\sigma : A$, with a **labelling function** $\lambda_\sigma : |\sigma| \rightarrow \text{Tm}_\Sigma(|\sigma|)$, satisfying (with $[a]_\sigma^\forall = \{a' \in |\sigma| \mid a' \leq_\sigma a \ \& \ \text{pol}_A(a') = \forall\}$):

- (1) Σ -receptivity: $\forall a^\forall \in |\sigma|, \lambda_\sigma(a) = a$,
- (2) Σ -courtesy: $\forall a^\exists \in |\sigma|, \lambda_\sigma(a) \in \text{Tm}_\Sigma([a]_\sigma^\forall)$.

Rather than having \forall moves introduce fresh variables, we consider them as *variables* themselves. Hence, the \exists moves carry terms having as free variables the \forall moves in their causal history. For instance the diagram of Figure 1 (right) 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 bélar's events.

Besides the fact that they are not assumed finite, Σ -strategies are more general than expansion trees: 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 $a_1 \rightarrow_\sigma a_2$ such that $a_1 \notin \text{fv}(\lambda_\sigma(a_2))$, then $a_1 \rightarrow_A a_2$ as well. In a minimal Σ -strategy $\sigma : A$, the ordering \leq_σ 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, we aim to capture that the substitution (by terms from the strategies) of the expansion of the original formula is a tautology.

► **Definition 6.** A **game** \mathcal{A} is an arena A , with $\mathcal{W}_\mathcal{A} : (x \in \mathcal{C}^\infty(A)) \rightarrow \text{QF}_\Sigma^\infty(x)$ expressing **winning conditions**, where $\text{QF}_\Sigma^\infty(x)$ denotes the **infinitary quantifier-free formulas** – obtained from $\text{QF}_\Sigma(x)$ by adding infinitary connectives $\bigvee_{i \in I} \varphi_i$ and $\bigwedge_{i \in I} \varphi_i$, with I countable.

For a game interpreting a formula φ , the winning conditions associate configurations of the arena $\llbracket \varphi \rrbracket$ with the propositional part of the corresponding *expansion* of φ . For instance:

$$\begin{aligned} \mathcal{W}_{\llbracket DF \rrbracket^\exists}(\{\exists_3, \forall_3, \exists_6, \forall_6\}) &= (\neg P(\exists_3) \vee P(\forall_3)) \vee (\neg P(\exists_6) \vee P(\forall_6)) \\ \mathcal{W}_{\llbracket DF \rrbracket^\exists}(\{\exists_3, \forall_3, \exists_6\}) &= (\neg P(\exists_3) \vee P(\forall_3)) \vee \top \end{aligned}$$

recalling that the arena for DF appears in Figure 4. In the second clause, \top (the true formula) comes from \forall bélar 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 7.** If $\sigma : A$ is a Σ -strategy and $x \in \mathcal{C}^\infty(\sigma)$, we say that x is **tautological** in σ if the formula $\mathcal{W}_\mathcal{A}(x)[\lambda_\sigma]$ corresponding to the substitution of $\mathcal{W}_\mathcal{A}(x) \in \text{QF}_\Sigma^\infty(x)$ by $\lambda_\sigma : x \rightarrow \text{Tm}_\Sigma(x)$, is a (possibly infinite) tautology.

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

Finally, a Σ -strategy $\sigma : A$ is **top-winning** if $|\sigma| \in \mathcal{C}^\infty(\sigma)$ is tautological.

2.3 Constructions on games and Herbrand's theorem

To complete our statement of Herbrand's theorem with Σ -strategies, it remains to set the interpretation of formulas as games. To that end we introduce a few constructions on games, first at the level of arenas and then enriched with winning conditions. We write \emptyset for the **empty arena**. If A is an arena, A^\perp is its **dual**, with same events and causality but polarity reversed. We review some other constructions.

► **Definition 8.** The **simple parallel composition** $A_1 \parallel A_2$ of A_1 and A_2 has as events the tagged disjoint union $\{1\} \times |A_1| \uplus \{2\} \times |A_2|$, as causal order that given by $(i, a) \leq_{A_1 \parallel A_2} (j, a')$ iff $i = j$ and $a \leq_{A_i} a'$, and, as polarity $\text{pol}_{A_1 \parallel A_2}((i, a)) = \text{pol}_{A_i}(a)$.

$$\begin{array}{llll} \llbracket \top \rrbracket_{\mathcal{V}}^{\exists} = 1 & \llbracket \mathsf{P}(t_1, \dots, t_n) \rrbracket_{\mathcal{V}}^{\exists} = \mathsf{P}(t_1, \dots, t_n) & \llbracket \exists x \varphi \rrbracket_{\mathcal{V}}^{\exists} = ?\exists x. \llbracket \varphi \rrbracket_{\mathcal{V} \uplus \{x\}} & \llbracket \varphi_1 \vee \varphi_2 \rrbracket_{\mathcal{V}}^{\exists} = \llbracket \varphi_1 \rrbracket_{\mathcal{V}}^{\exists} \wp \llbracket \varphi_2 \rrbracket_{\mathcal{V}}^{\exists} \\ \llbracket \perp \rrbracket_{\mathcal{V}}^{\exists} = \perp & \llbracket \neg \mathsf{P}(t_1, \dots, t_n) \rrbracket_{\mathcal{V}}^{\exists} = \neg \mathsf{P}(t_1, \dots, t_n) & \llbracket \forall x \varphi \rrbracket_{\mathcal{V}}^{\exists} = \forall x. \llbracket \varphi \rrbracket_{\mathcal{V} \uplus \{x\}} & \llbracket \varphi_1 \wedge \varphi_2 \rrbracket_{\mathcal{V}}^{\exists} = \llbracket \varphi_1 \rrbracket_{\mathcal{V}}^{\exists} \otimes \llbracket \varphi_2 \rrbracket_{\mathcal{V}}^{\exists} \end{array}$$

■ **Figure 5** \exists -biased interpretation of formulas.

Configurations $x \in \mathcal{C}^{\infty}(A \parallel B)$ have the form $\{1\} \times x_A \cup \{2\} \times x_B$ with $x_A \in \mathcal{C}^{\infty}(A)$ and $x_B \in \mathcal{C}^{\infty}(B)$, which we write $x = x_A \parallel x_B$. This construction has a general counterpart $\parallel_{i \in I} A_i$ with I at most countable, defined likewise. In particular we will later use the uniform countably infinite parallel composition $\parallel_{\omega} A$. Another important construction is *prefixing*.

► **Definition 9.** For $\alpha \in \{\forall, \exists\}$ and A an 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 $(i, a) = (1, \alpha)$; i.e., $(1, \alpha)$ is the unique minimal event. Its polarity is $\text{pol}_{\alpha.A}((1, \alpha)) = \alpha$ and $\text{pol}_{\alpha.A}((2, a)) = \text{pol}_A(a)$.

Configurations $x \in \mathcal{C}^{\infty}(\alpha.A)$ are \emptyset , or $\{(1, \alpha)\} \cup \{2\} \times x_A$ ($x_A \in \mathcal{C}^{\infty}(A)$), written $\alpha.x_A$.

Now, let us enrich these with winning, yielding the constructions on games used for interpreting formulas. Importantly, the inductive interpretation of formulas requires us to consider formulas with free variables. For \mathcal{V} a finite set, a **\mathcal{V} -game** is defined as a game \mathcal{A} (Def. 6), except that winning may also depend on \mathcal{V} : for $x \in \mathcal{C}^{\infty}(A)$, $\mathcal{W}_{\mathcal{A}}(x) \in \mathbf{QF}_{\Sigma \uplus \mathcal{V}}^{\infty}(x)$.

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

► **Definition 10.** For \mathcal{A} and \mathcal{B} \mathcal{V} -games, we define two \mathcal{V} -games with arena $A \parallel B$ and winning conditions $\mathcal{W}_{\mathcal{A} \otimes \mathcal{B}}(x_A \parallel x_B) = \mathcal{W}_{\mathcal{A}}(x_A) \wedge \mathcal{W}_{\mathcal{B}}(x_B)$ and $\mathcal{W}_{\mathcal{A} \wp \mathcal{B}}(x_A \parallel x_B) = \mathcal{W}_{\mathcal{A}}(x_A) \vee \mathcal{W}_{\mathcal{B}}(x_B)$.

Note the implicit renaming so that $\mathcal{W}_{\mathcal{A}}(x_A), \mathcal{W}_{\mathcal{B}}(x_B)$ are in $\mathbf{QF}_{\Sigma \uplus \mathcal{V}}^{\infty}(x_A \parallel x_B)$ rather than $\mathbf{QF}_{\Sigma \uplus \mathcal{V}}^{\infty}(x_A), \mathbf{QF}_{\Sigma \uplus \mathcal{V}}^{\infty}(x_B)$ respectively – we will often keep such renamings implicit.

Observe that \otimes and \wp are De Morgan duals, i.e., $(\mathcal{A} \otimes \mathcal{B})^{\perp} = \mathcal{A}^{\perp} \wp \mathcal{B}^{\perp}$. We write these operations \otimes and \wp rather than \wedge and \vee , because they behave more like the connectives of linear logic [12] than those of classical logic; for each \mathcal{V} the \otimes and \wp will form the basis of a $*$ -autonomous structure and hence a model of multiplicative linear logic (see Section 3).

To interpret classical logic however, we will need *replication*.

► **Definition 11.** For \mathcal{V} -game \mathcal{A} , we define the \mathcal{V} -games $! \mathcal{A}, ? \mathcal{A}$ with arena $\parallel_{\omega} A$ and winning:

$$\mathcal{W}_{! \mathcal{A}}(\parallel_{i \in \omega} x_i) = \bigwedge_{i \in \omega} \mathcal{W}_{\mathcal{A}}(x_i) \quad \mathcal{W}_{? \mathcal{A}}(\parallel_{i \in \omega} x_i) = \bigvee_{i \in \omega} \mathcal{W}_{\mathcal{A}}(x_i)$$

Though $\mathcal{W}_{! \mathcal{A}}(x)$ (resp. $\mathcal{W}_{? \mathcal{A}}(x)$) is an infinite conjunction (resp. disjunction), it simplifies to a finite one when x visits finitely many copies (with cofinitely many copies of $\mathcal{W}_{\mathcal{A}}(\emptyset)$).

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

► **Definition 12.** Let \mathcal{A} a $(\mathcal{V} \uplus \{x\})$ -game, we define the \mathcal{V} -game $\forall x. \mathcal{A}$ and its dual $\exists x. \mathcal{A}$ with arenas $\forall. A$ and $\exists. A$ respectively, with $\mathcal{W}_{\forall x. \mathcal{A}}(\emptyset) = \top$, $\mathcal{W}_{\exists x. \mathcal{A}}(\emptyset) = \perp$, and:

$$\mathcal{W}_{\forall x. \mathcal{A}}(\forall. x_A) = \mathcal{W}_{\mathcal{A}}(x_A)[\forall/x] \quad \mathcal{W}_{\exists x. \mathcal{A}}(\exists. x_A) = \mathcal{W}_{\mathcal{A}}(x_A)[\exists/x]$$

Finally, we regard a literal φ as a \mathcal{V} -game on arena \emptyset , with $\mathcal{W}_{\varphi}(\emptyset) = \varphi$. We write 1 and \perp for the unit \mathcal{V} -games on arena \emptyset with winning conditions respectively \top and \perp .

Putting these together, we give in Figure 5 the \exists -biased interpretation of a formula $\varphi \in \text{Form}_{\Sigma}(\mathcal{V})$ as a \mathcal{V} -game. Note the difference between the case of existential and universal

\mathcal{V}-MLL			
$\text{Ax} \frac{}{\vdash^{\mathcal{V}} \varphi^{\perp}, \varphi} \text{fv}(\varphi) \subseteq \mathcal{V}$	$\text{Cut} \frac{\vdash^{\mathcal{V}} \Gamma, \varphi \quad \vdash^{\mathcal{V}} \varphi^{\perp}, \Delta}{\vdash^{\mathcal{V}} \Gamma, \Delta}$	$\text{Ex} \frac{\vdash^{\mathcal{V}} \Gamma, \varphi, \psi, \Delta}{\vdash^{\mathcal{V}} \Gamma, \psi, \varphi, \Delta}$	
$\text{TI} \frac{}{\vdash^{\mathcal{V}} \top}$	$\text{II} \frac{\vdash^{\mathcal{V}} \Gamma}{\vdash^{\mathcal{V}} \Gamma, \perp}$	$\text{LI} \frac{\vdash^{\mathcal{V}} \Gamma, \varphi \quad \vdash^{\mathcal{V}} \psi, \Delta}{\vdash^{\mathcal{V}} \Gamma, \varphi \wedge \psi, \Delta}$	$\text{VI} \frac{\vdash^{\mathcal{V}} \Gamma, \varphi, \psi, \Delta}{\vdash^{\mathcal{V}} \Gamma, \varphi \vee \psi, \Delta}$
First-order MLL (MLL₁)		LK	
$\forall\text{I} \frac{\vdash^{\mathcal{V}\{x\}} \Gamma, \varphi}{\vdash^{\mathcal{V}} \Gamma, \forall x. \varphi} x \notin \text{fv}(\Gamma)$	$\exists\text{I} \frac{\vdash^{\mathcal{V}} \Gamma, \varphi[t/x]}{\vdash^{\mathcal{V}} \Gamma, \exists x. \varphi} t \in \text{TM}_{\Sigma}(\mathcal{V})$	$\text{C} \frac{\vdash^{\mathcal{V}} \Gamma, \varphi, \varphi}{\vdash^{\mathcal{V}} \Gamma, \varphi}$	$\text{W} \frac{\vdash^{\mathcal{V}} \Gamma}{\vdash^{\mathcal{V}} \Gamma, \varphi}$

■ **Figure 6** Rules for the sequent calculus LK.

formulas, reflecting the bias towards \exists loïse. This is indeed compatible with the examples given previously. We can now state our concurrent version of Herbrand's theorem.

► **Theorem 13.** *For any $\varphi \in \text{Form}_{\Sigma}$, $\models \varphi$ iff there exists a finite, top-winning $\sigma : \llbracket \varphi \rrbracket^{\exists}$.*

Besides the game-theoretic language, the difference with expansion trees is superficial: on φ , expansion trees essentially coincide with the *minimal* top-winning Σ -strategies $\sigma : \llbracket \varphi \rrbracket^{\exists}$. The effort to change view point, from a syntactic construction to a (game) semantic one, will however pay off now, when we show how to *compose* Σ -strategies.

2.4 Compositional Herbrand's theorem

Unlike expansion trees, strategies can be *composed*. Whereas Theorem 13 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's theorem where the witnesses are obtained truly *compositionally* from any sequent proof, without first eliminating cuts. In other words, strategies will come naturally from the interpretation of the classical sequent calculus in a semantic model.

To compose Σ -strategies, we must restore the symmetry between \exists loïse and \forall bélar in the interpretation of formulas. The *non-biased* interpretation $\llbracket \varphi \rrbracket_{\mathcal{V}}$ of $\varphi \in \text{Form}_{\Sigma}(\mathcal{V})$ is defined as for $\llbracket \varphi \rrbracket_{\mathcal{V}}^{\exists}$, except for $\llbracket \forall x \varphi \rrbracket_{\mathcal{V}} = !\forall x. \llbracket \varphi \rrbracket_{\mathcal{V}\{x\}}$. Thus we lose finiteness: \exists loïse must be reactive to the infinite number of copies potentially opened by \forall bélar. But we can now state:

► **Theorem 14.** *For φ closed, the following are equivalent: (1) $\models \varphi$, (2) there exists a finite, top-winning Σ -strategy $\sigma : \llbracket \varphi \rrbracket^{\exists}$, (3) there exists a winning Σ -strategy $\sigma : \llbracket \varphi \rrbracket$.*

Proof. That (2) implies (1) is easy, as a finite top-winning $\sigma : \llbracket \varphi \rrbracket^{\exists}$ directly informs a proof.

That (3) implies (2) is more subtle: first, one may restrict a winning $\sigma : \llbracket \varphi \rrbracket$ to $\llbracket \varphi \rrbracket^{\exists}$ to obtain a finite top-winning strategy. However, this top-winning strategy *may not be finite*. Yet, it follows by compactness that there is always a finite top-winning sub-strategy that may be effectively computed from σ . See the Appendix C for details.

The proof that (1) implies (3) is our main contribution: a winning strategy will be computed from a proof using our denotational model of classical proofs. ◀

Our source sequent calculus (Figure 6) is fairly standard, one-sided, with rules presented in the multiplicative style. A notable variation is that sequents carry a set \mathcal{V} of free variables, that may appear freely in formulas. 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? Ignoring the \mathcal{V} annotations, the first group is nothing but Multiplicative Linear Logic (MLL). Propositional (\mathcal{V} -)MLL can be interpreted in a $*$ -autonomous category [3]. Accordingly, in Section 3, we first construct a $*$ -autonomous category Ga of games and winning Σ -strategies. Then, in Section 4, we build the structure required for the interpretation of quantifiers, still ignoring contraction and weakening. For each set of variables \mathcal{V} we construct a $*$ -autonomous category $\mathcal{V}\text{-Ga}$, with a fibred structure to link the $\mathcal{V}\text{-Ga}$ together for distinct \mathcal{V} s and suitable structure to deal with quantifiers, obtaining a model of first-order MLL. Finally in Section 5 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.

3 A $*$ -autonomous category

The following theorem, on cut reduction for MLL, is folklore.

► **Theorem 15.** *There is a set of reduction rules on MLL sequent proofs, written $\rightsquigarrow_{\text{MLL}}$, such that for any proof π of a sequent $\vdash \Gamma$, there is a cut-free π' of Γ such that $\pi \rightsquigarrow_{\text{MLL}}^* \pi'$.*

The reduction $\rightsquigarrow_{\text{MLL}}$ comprises *logical* reductions, reducing a cut on a formula φ/φ^\perp , between two proofs starting with the introduction rule for the main connective of φ/φ^\perp ; and *structural* reductions, consisting in commutations between rules so as to reach the logical steps. We assume 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 this is essentially the same as producing a *$*$ -autonomous category*. We opt here for the equivalent formulation by Cockett and Seely as a *symmetric linearly distributive category with negation* [6].

► **Definition 16.** A **symmetric linearly distributive category** is a category \mathcal{C} with two symmetric monoidal structures $(\otimes, 1)$ and (\wp, \perp) which *distribute*: there is a natural $\delta_{A,B,C} : A \otimes (B \wp C) \xrightarrow{\mathcal{C}} (A \otimes B) \wp C$, 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 $\eta_A : 1 \xrightarrow{\mathcal{C}} A^{\perp} \wp A$ and $\epsilon_A : A \otimes A^{\perp} \xrightarrow{\mathcal{C}} \perp$ such that the canonical composition $A \rightarrow A \otimes (A^{\perp} \wp A) \rightarrow (A \otimes A^{\perp}) \wp A \rightarrow A$, and its dual $A^{\perp} \rightarrow A^{\perp}$, 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 (\wp, \perp) coincide.

Abusing terminology, we will refer to *symmetric linearly distributive categories with negation* by the shorter *$*$ -autonomous categories*. This should not create any confusion in the light of their equivalence [6]. If \mathcal{C} a $*$ -autonomous category comes with a choice of $\llbracket P(t_1, \dots, t_n) \rrbracket$ (an object of \mathcal{C}) for all closed literal, then this interpretation can be extended to all closed quantifier-free formulas following Figure 5. For all such φ , we have $\llbracket \varphi^{\perp} \rrbracket = \llbracket \varphi \rrbracket^{\perp}$.

The interpretation of MLL proofs in a $*$ -autonomous category \mathcal{C} is standard [29]: a proof π of a *MLL sequent* $\vdash \varphi_1, \dots, \varphi_n$ is interpreted as a morphism $\llbracket \pi \rrbracket : 1 \xrightarrow{\mathcal{C}} \llbracket \varphi_1 \rrbracket \wp \dots \wp \llbracket \varphi_n \rrbracket$. This interpretation is sound *w.r.t.* provability: if φ is provable, then $1 \rightarrow_{\mathcal{C}} \llbracket \varphi \rrbracket$ is inhabited. Furthermore, the categorical laws make this interpretation invariant under cut reduction.

► **Theorem 17.** *If $\pi \rightsquigarrow_{\text{MLL}} \pi'$ are proofs of $\vdash \Gamma$, $\llbracket \pi \rrbracket = \llbracket \pi' \rrbracket$.*



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

So a proof has the same denotation as its cut-free form obtained by Theorem 15. In the rest of this section we construct a concrete $*$ -autonomous category of games and winning Σ -strategies; supporting the interpretation of MLL. This is done in three stages: first we focus on composition of Σ -strategies (without winning), then we extend this to a compact closed category. Finally, adding back winning, we split \parallel into two \otimes and \wp , and prove $*$ -autonomy.

3.1 Composition of Σ -strategies

We construct a category Ar_Σ having arenas as objects, and as morphisms from A to B the Σ -strategies $\sigma : A^\perp \parallel B$, also written $\sigma : A \xrightarrow{\text{Ar}_\Sigma} B$. The composition of $\sigma : A \xrightarrow{\text{Ar}_\Sigma} B$ and $\tau : B \xrightarrow{\text{Ar}_\Sigma} C$ will be computed in two stages: first, the *interaction* $\tau \otimes \sigma$ 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 $\exists_5^{f(g(c), h(c))}$.

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

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

► **Definition 18.** A **neutral arena** is an arena, without polarities. **Neutral strategies** $\sigma : A$, are defined as in Definition 4 without (2), (3). **Neutral Σ -strategies** additionally have $\lambda_\sigma : (s \in |\sigma|) \rightarrow \text{Tm}_\Sigma([s]_\sigma)$, and are **idempotent**: for all $a \in |a|$, $\lambda_\sigma(a)[\lambda_\sigma] = \lambda_\sigma(a)$.

Forgetting polarities, every Σ -strategy is a neutral one. Given σ and τ , $\tau \otimes \sigma$ is a *minimal strengthening* of σ and τ , regarding both the causal structure and term annotations, i.e., a *meet* for the partial order (*idempotence* above is required for it to be antisymmetric):

► **Definition 19.** For $\sigma, \tau : A$ neutral Σ -strategies, we write $\sigma \preceq \tau$ iff $|\sigma| \subseteq |\tau|$, $\mathcal{C}^\infty(\sigma) \subseteq \mathcal{C}^\infty(\tau)$, and for all $x \in \mathcal{C}(|\sigma|)$, $\lambda_\tau \upharpoonright x$ subsumes $\lambda_\sigma \upharpoonright x$ (regarded as substitutions $x \xrightarrow{\Sigma} x$).

Ignoring terms, any two σ and τ have a meet $\sigma \wedge \tau$; this is a simplification of the *pullback* in the category of event structures, exploiting the absence of conflict [31]. The partial order $(|\sigma \wedge \tau|, \leq_{\sigma \wedge \tau})$ has events all common moves of σ and τ with a causal history compatible with both \leq_σ and \leq_τ , and for $\leq_{\sigma \wedge \tau}$ the minimal causal order compatible with both.

However, two neutral Σ -strategies do not necessarily have a meet for \preceq (see Example 45 in Appendix A). Hence, we focus on the meets occurring from compositions of Σ -strategies and show that for $\sigma : A$ and $\tau : A^\perp$ dual Σ -strategies the meet *does* exist:

► **Lemma 20.** *Any two Σ -strategies $\sigma : A$ and $\tau : A^\perp$ have a meet $\sigma \wedge \tau$.*

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

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

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 Σ -courtesy, if $a^\exists \in |\sigma|$ then $\lambda_\sigma(a) \in \text{Tm}_\Sigma([a]_\sigma)$. ◀

However this is not sufficient: 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 $(\sigma \parallel C^\perp) \wedge (A \parallel \tau)$, where the parallel composition of Definition 8 is extended with terms in the obvious way, and where $\lambda_A(a) = a$ 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, but Σ -courtesy from Σ -strategies is relaxed to idempotence. Yet, Lemma 20 still holds since, from idempotence, if $a^\exists \in |\sigma|$ then either $\lambda_\sigma(a) \in \text{Tm}_\Sigma([a]_\sigma)$ or $\lambda_\sigma(a) = a$. Hence, we can define $\tau \otimes \sigma = (\sigma \parallel C^\perp) \wedge (A \parallel \tau) : A \parallel B \parallel C$.

Variables appearing in $\lambda_{\tau \otimes \sigma}$ cannot be events in B – they must be negative in $A^\perp \parallel C$. So we can define $\tau \odot \sigma = (\tau \otimes \sigma) \cap (A \parallel C)$ the restriction of $\tau \otimes \sigma$ to $A \parallel C$, with same causal order and term annotation. The pair $(|\tau \odot \sigma|, \leq_{\tau \odot \sigma})$ is a strategy, as an instance of the constructions in [4], and this extends to terms so that $\tau \odot \sigma : A^\perp \parallel C$ is a Σ -strategy, the **composition** of σ and τ . Because interaction is defined as a meet for \preceq , it follows that it is compatible with it, i.e., if $\sigma \preceq \sigma'$, then $\tau \otimes \sigma \preceq \tau \otimes \sigma'$. This is preserved by projection, and hence $\tau \odot \sigma \preceq \tau \odot \sigma'$ as well. This compatibility of composition with \preceq will be used later on, together with the easy fact that \preceq is more constrained on Σ -strategies:

► **Lemma 21.** *For $\sigma, \sigma' : A$ Σ -strategies, if $\sigma \preceq \sigma'$, then $\lambda_\sigma(s) = \lambda_{\sigma'}(s)$ for all $s \in |\sigma|$.*

To complete our category, we also define the *copycat strategy*.

► **Definition 22.** For an arena A , the **copycat Σ -strategy** $\alpha_A : A^\perp \parallel A$ has events $|\alpha_A| = A^\perp \parallel A$. Writing $(i, a) = (3 - i, a)$, its partial order \leq_{α_A} is the transitive closure of $\leq_{A^\perp \parallel A} \cup \{(c, \bar{c}) \mid c^\forall \in |A^\perp \parallel A|\}$ and its labelling function is $\lambda_{\alpha_A}(c^\forall) = c$, $\lambda_{\alpha_A}(c^\exists) = \bar{c}$.

The proof of categorical laws are variations on construction of the bicategory in [4].

► **Proposition 23.** *There is a poset-enriched category Ar_Σ with arenas as objects, and Σ -strategies as morphisms.*

3.2 Compact closed structure

We show that Ar_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 notations, we also write $\sigma_1 \parallel \sigma_2 : (A_1 \parallel A_2)^\perp \parallel (B_1 \parallel B_2)$ for the obvious renaming. It is not difficult to prove:

► **Proposition 24.** *Simple parallel composition yields an enriched functor $\parallel : \text{Ar}_\Sigma \times \text{Ar}_\Sigma \rightarrow \text{Ar}_\Sigma$.*

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

► **Definition 25.** 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 the corresponding **copycat strategy** is $\alpha_f = (A^\perp \parallel f) * \alpha_A : A^\perp \parallel B$. We use this to define the structural morphisms for the symmetric monoidal structure of Ar_Σ . For instance, the iso $\alpha_{A,B,C} : (A \parallel B) \parallel C \cong A \parallel (B \parallel C)$ yields $\alpha_{\alpha_{A,B,C}} : (A \parallel B) \parallel C \xrightarrow{\text{Ar}_\Sigma} A \parallel (B \parallel C)$. The other structural morphisms arise similarly. Coherence and naturality then follows from the key *copycat lemma*:

► **Lemma 26.** For $\sigma : A^\perp \parallel B$ a Σ -strategy and $f : B \cong C$, $\alpha_f \odot \sigma = (A^\perp \parallel f) * \sigma : A^\perp \parallel C$.

As a corollary we get coherence for the structural morphisms (following from those on isomorphisms), and naturality. For all A we get $\eta_A : \emptyset \xrightarrow{\text{Ar}_\Sigma} A^\perp \parallel A$ and $\epsilon_A : A \parallel A^\perp \xrightarrow{\text{Ar}_\Sigma} \emptyset$ as the obvious renamings of copycat. Checking the law for compact closed categories is a variation of the idempotence of copycat. Overall:

► **Proposition 27.** Ar_Σ is a poset-enriched compact closed category.

3.3 A linearly distributive category with negation

Finally, we reinstate winning conditions. We first note:

► **Proposition 28.** There is a (poset-enriched) category Ga_Σ with objects the games (Definition 6) on Σ , and morphisms Σ -strategies $\sigma : \mathcal{A}^\perp \wp \mathcal{B}$, also written $\sigma : \mathcal{A} \xrightarrow{\text{Ga}_\Sigma} \mathcal{B}$.

That copycat is winning boils down to the excluded middle. That $\tau \odot \sigma : \mathcal{A}^\perp \wp \mathcal{C}$ is winning if $\sigma : \mathcal{A}^\perp \wp \mathcal{B}$ and $\tau : \mathcal{B}^\perp \wp \mathcal{C}$ are, is as in [5]: for $x \in \mathcal{C}(\tau \odot \sigma)$ \exists -maximal we find a witness $y \in \mathcal{C}(\tau \otimes \sigma)$ (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; and apply transitivity of implication. The equations follow from Ar_Σ . Likewise:

► **Proposition 29.** The functor $\parallel : \text{Ar}_\Sigma \times \text{Ar}_\Sigma \rightarrow \text{Ar}_\Sigma$ splits into $\otimes, \wp : \text{Ga}_\Sigma \times \text{Ga}_\Sigma \rightarrow \text{Ga}_\Sigma$.

It suffices to check winning, which is straightforward. It remains to prove that all structural morphisms from Ar_Σ (copycat strategies) are winning, which boils down to the following sufficient conditions to hold: For \mathcal{A}, \mathcal{B} games, a **win-iso** $f : \mathcal{A} \rightarrow \mathcal{B}$ is an iso $f : A \cong B$ such that $(\mathcal{W}_A(x))^\perp \vee \mathcal{W}_B(f x)$ is a tautology, for all $x \in \mathcal{C}^\infty(A)$.

► **Lemma 30.** If $f : \mathcal{A} \rightarrow \mathcal{B}$ is a win-iso, then $\alpha_f : \mathcal{A}^\perp \wp \mathcal{B}$ is a winning Σ -strategy.

This easily entails that all structural morphisms (including linear distributivity) are winning. Finally $\eta_A : \mathbf{1} \xrightarrow{\text{Ga}_\Sigma} \mathcal{A}^\perp \wp \mathcal{A}$ and $\epsilon_A : \mathcal{A} \otimes \mathcal{A}^\perp \xrightarrow{\text{Ga}_\Sigma} \mathbf{1}$ are winning, which concludes:

► **Proposition 31.** Ga_Σ is a poset-enriched *-autonomous category.

4 A model of first-order MLL

We move on to MLL_1 , i.e., all rules except for contraction and weakening. Before developing the interpretation, we discuss cut elimination. There are three new cut reduction rules, displayed in Figure 8: the new *logical* reduction (\forall/\exists), and two for the propagation of cuts past introduction rules for \forall and \exists . Writing $\pi \rightsquigarrow_{\text{MLL}_1} \pi'$ for the reduction obtained with these new rules together with $\rightsquigarrow_{\text{MLL}}$:

► **Proposition 32.** Let π be any MLL_1 proof of $\vdash^\vee \Gamma$. Then, there is a cut-free proof π' of $\vdash^\vee \Gamma$ s.t. $\pi \rightsquigarrow_{\text{MLL}_1}^* \pi'$.

are games on the signature $\Sigma \uplus \mathcal{V}$, i.e., the \mathcal{V} -games of Section 2.3. *Morphisms* between \mathcal{V} -games \mathcal{A} and \mathcal{B} are winning $(\Sigma \uplus \mathcal{V})$ -strategies on $\mathcal{A}^\perp \wp \mathcal{B}$ regarded as a game on signature $\Sigma \uplus \mathcal{V}$ – also called **winning Σ -strategies on the \mathcal{V} -game $\mathcal{A}^\perp \wp \mathcal{B}$** .

Finally, for \mathcal{A} a \mathcal{V}_2 -game and $\gamma : \mathcal{V}_1 \rightarrow \mathcal{V}_2$ a substitution, the game $\mathcal{T}(\gamma)(\mathcal{A}) = \mathcal{A}[\gamma]$ is defined as having arena A , and, for $x \in \mathcal{C}^\infty(A)$, $\mathcal{W}_{\mathcal{A}[\gamma]}(x) = \mathcal{W}_{\mathcal{A}}(x)[\gamma] \in \mathbf{QF}_{\Sigma \uplus \mathcal{V}_1}^\infty(x)$. Likewise, given \mathcal{A} and \mathcal{B} two \mathcal{V} -games and $\sigma : \mathcal{A}^\perp \wp \mathcal{B}$, $\sigma[\gamma]$ has the same components as σ , but term annotations $\lambda_{\sigma[\gamma]}(s) = \lambda(s)[\gamma] \in \mathbf{Tm}_{\Sigma \uplus \mathcal{V}_1}(x)$. It is a simple verification to prove:

► **Proposition 34.** *For any $\gamma : \mathcal{V}_1 \xrightarrow{\mathcal{S}} \mathcal{V}_2$, $\mathcal{T}(\gamma) : \mathcal{T}(\mathcal{V}_2) \rightarrow \mathcal{T}(\mathcal{V}_1)$ is a strict $*$ -autonomous functor preserving the order.*

4.2 Quantifiers

Finally, we give the interpretation of $\forall I$ and $\exists I$. For now, we consider a *linear* interpretation $\llbracket - \rrbracket_{\mathcal{V}}^\ell$ of formulas defined like $\llbracket - \rrbracket_{\mathcal{V}}^\exists$ except for $\llbracket \exists x \varphi \rrbracket_{\mathcal{V}}^\ell = \exists x. \llbracket \varphi \rrbracket_{\mathcal{V}}^\ell$.

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

► **Lemma 35.** *Let $\varphi \in \mathbf{Form}_\Sigma(\mathcal{V}_2)$ and $\gamma : \mathcal{V}_1 \rightarrow \mathcal{V}_2$ a substitution, then $\llbracket \varphi[\gamma] \rrbracket_{\mathcal{V}_1}^\ell = \llbracket \varphi \rrbracket_{\mathcal{V}_2}^\ell[\gamma]$.*

This will be used implicitly from now on. The definition of quantifiers on games of Definition 12 extends to functors $\forall_{\mathcal{V},x}, \exists_{\mathcal{V},x} : \mathcal{T}(\mathcal{V} \uplus \{x\}) \rightarrow \mathcal{T}(\mathcal{V})$. From $\sigma : \mathcal{A}^\perp \wp \mathcal{B}$, $\forall_{\mathcal{V},x}(\sigma) : (\forall x. \mathcal{A})^\perp \wp \forall x. \mathcal{B}$ plays copycat on the initial \forall , then plays as σ (similarly for $\exists_{\mathcal{V},x}(\sigma)$). Following Lawvere [20], one expects adjunctions $\exists_{\mathcal{V},x} \dashv \mathcal{T}(w_{\mathcal{V},x}) \dashv \forall_{\mathcal{V},x}$. Unfortunately, this fails – we present this failure later as the non-preservation of $\rightsquigarrow_{\text{CUT}/\forall}$.

We now interpret $\forall I$ and $\exists I$. First, we give a strategy introducing a witness t .

► **Definition 36.** The $(\Sigma \uplus \mathcal{V})$ -strategy $\exists_A^t : A^\perp \parallel \exists. A$ is $(|A^\perp \parallel \exists. A|, \leq_{\exists_A^t}, \lambda_{\exists_A^t})$ where $\leq_{\exists_A^t}$ includes \leq_{α_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 α_A plus $\lambda_{\exists_A^t}((2, \exists)) = t$.

In other words, \exists_A^t plays \exists annotated with t , then proceeds as copycat on A . We have:

► **Proposition 37.** *Let \mathcal{A} be a \mathcal{V} -game, and $t \in \mathbf{Tm}_\Sigma(\mathcal{V})$. Then, $\exists_A^t : \mathcal{A}[t/x] \xrightarrow{\mathcal{V}\text{-Gas}} \exists x. \mathcal{A}$.*


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

To a strategy σ , the operation interpreting $\forall I$ adds \forall as new minimal event, and sets it as a dependency for all events whose annotation comprise the distinguished variable x .

► **Definition 38.** For σ a $(\Sigma \uplus \mathcal{V} \uplus \{x\})$ -strategy on $A^\perp \parallel B$, the $(\Sigma \uplus \mathcal{V})$ -strategy $\forall_{A,B}^x(\sigma) : A^\perp \parallel \forall. B$ has events $|\sigma| \uplus \{(2, \forall)\}$, term assignment $\lambda((2, \forall)) = (2, \forall)$ and causality $\lambda(s) = \lambda_\sigma(s)[(2, \forall)/x]$ ($s \in |\sigma|$), and $\leq = \leq_\sigma \cup \{((2, \forall), s) \mid s \in \forall. B \quad \vee \quad \exists s' \leq_\sigma s, x \in \text{fv}(\lambda_\sigma(s'))\}$.

► **Proposition 39.** *If σ is winning on a $(\mathcal{V} \uplus \{x\})$ -game $\mathcal{A}[w_{\mathcal{V},x}] \wp \mathcal{B}$, then $\forall_{A,B}^x(\sigma)$ is winning on the \mathcal{V} -game $\mathcal{A} \wp \forall x. \mathcal{B}$.*

Indeed, if \forall bélarde does not play $(2, \forall)$ we get a tautology, otherwise the remaining configuration is in σ and so is tautological. This completes the interpretation of MLL_1 . This interpretation leaves $\rightsquigarrow_{\forall/\exists}$ invariant, but fails $\rightsquigarrow_{\text{CUT}/\forall}$. This stems from the fact that the *minimal* Σ -strategies are not stable under composition (see Example 46 in Appendix A). The interpretation of cut-free proofs yield minimal Σ -strategies. In contrast, in compositions



■ **Figure 9** Two examples of contraction.

interpreting cuts, causality may flow through the syntax tree of the cut formula, and create causal dependencies not reflected in the variables. Hence, cut reduction may weaken the causal structure.

► **Lemma 40.** *For $\sigma : A \xrightarrow{\text{Ar}\Sigma} B$ and $\tau : B \xrightarrow{\text{Ar}\Sigma_{\omega\{x\}}} C$, we have $\forall I_{A,C}^x(\tau \odot \sigma) \preceq \forall I_{B,C}^x(\tau) \odot \sigma$.*

By Lemma 21 these two have the same terms on common events. In fact, $\forall I_{A,C}^x(\tau \odot \sigma)$ and $\forall I_{B,C}^x(\tau) \odot \sigma$ also have the same *events* – they correspond to the same *expansion tree*, only the acyclicity witness differs. But the variant of \preceq with $|\sigma_1| = |\sigma_2|$ is not a congruence: relaxing causality of σ in $\tau \odot \sigma$ may unlock new events, previously part of causal loops.

As \preceq is preserved by all operations on Σ -strategies, we deduce:

► **Theorem 41.** *If $\pi \rightsquigarrow_{\text{MLL}_1} \pi'$, then $\llbracket \pi' \rrbracket \preceq \llbracket \pi \rrbracket$.*

For MLL_1 , we conjecture that “having the same expansion tree” (i.e., same events and term annotations) is actually a congruence, yielding a **-autonomous hyperdoctrine*. As this would not hold in the presence of contraction and weakening, we leave this for future work.

5 Contraction and weakening

In this section we reinstate $!$ and $?$ in the interpretation of quantifiers, i.e., $\llbracket \forall x. \varphi \rrbracket_{\mathcal{V}} = !\forall x. \llbracket \varphi \rrbracket_{\mathcal{V}_{\omega\{x\}}}$ and $\llbracket \exists x \varphi \rrbracket_{\mathcal{V}} = ?\exists x \llbracket \varphi \rrbracket_{\mathcal{V}_{\omega x}}$ – this is reminiscent of Mellies’ discussion on the interaction between quantifiers and exponential modalities in a polarized setting [22].

Unlike for MLL_1 , we only aim to map proofs to Σ -strategies on the appropriate game, with no preservation of reduction. We must interpret contraction and weakening, but also revisit the interpretation of rules for quantifiers as the interpretation of formulas has changed.

Weakening is easy: for any game \mathcal{A} , any Σ -strategy $\sigma : \mathcal{A} \rightarrow 1$ is winning; for definiteness, we use the minimal $e_{\mathcal{A}} : \mathcal{A} \rightarrow 1$, 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 (without term annotations). The first looks like the usual contraction of 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 (mimicking the effect of cut reduction).

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

Recall $!$ and $?$ from Definition 11, both based on arena $\llbracket_{\omega} A$. First, we examine their functorial action. Let $\sigma : A \xrightarrow{\text{Ar}\Sigma} B$. Then, $\llbracket_{\omega} \sigma : \llbracket_{\omega} (A^{\perp} \parallel B)$ which is isomorphic to $(\llbracket_{\omega} A)^{\perp} \parallel (\llbracket_{\omega} B)$; overloading notion we still write $\llbracket_{\omega} \sigma : \llbracket_{\omega} A \xrightarrow{\text{Ar}\Sigma} \llbracket_{\omega} B$.

► **Lemma 42.** *Let $\sigma : \mathcal{A} \xrightarrow{\text{Ga}\Sigma} \mathcal{B}$. Then, we have $!\sigma = \llbracket_{\omega} \sigma : !\mathcal{A} \xrightarrow{\text{Ga}\Sigma} !\mathcal{B}$ and $?\sigma = \llbracket_{\omega} \sigma : ?\mathcal{A} \xrightarrow{\text{Ga}\Sigma} ?\mathcal{B}$.*

Rather than defining directly the contraction, we build $co_{\varphi} : \llbracket \varphi \rrbracket_{\mathcal{V}} \xrightarrow{\text{Ga}\Sigma_{\omega\mathcal{V}}} !\llbracket \varphi \rrbracket_{\mathcal{V}}$ by induction on $\varphi \in \text{Form}_{\Sigma}(\mathcal{V})$. For φ quantifier-free, the empty $co_{\varphi} : \llbracket \varphi \rrbracket_{\mathcal{V}} \rightarrow !\llbracket \varphi \rrbracket_{\mathcal{V}}$ is winning. We

$$\begin{aligned}
 \left[\left[\frac{\pi}{\frac{\vdash^{\mathcal{V}} \Gamma, \varphi, \varphi}{\vdash^{\mathcal{V}} \Gamma, \varphi}} \right] \right]_{\mathcal{C}} &= \Gamma^{\perp} \xrightarrow{\mathcal{T}(\mathcal{V})} \llbracket \pi \rrbracket \varphi \wp \varphi \xrightarrow{\mathcal{T}(\mathcal{V})} \delta_{\varphi^{\perp}}^{\perp} \varphi \\
 \left[\left[\frac{\pi}{\frac{\vdash^{\mathcal{V} \uplus \{x\}} \Gamma, \varphi}{\vdash^{\mathcal{V}} \Gamma, \forall x. \varphi}} \right] \right]_{\forall \mathbf{I}} &= \Gamma^{\perp} \xrightarrow{\mathcal{T}(\mathcal{V})} \text{co}_{\Gamma^{\perp}} \text{!} \Gamma^{\perp} \xrightarrow{\mathcal{T}(\mathcal{V})} \text{!(}\forall \mathbf{I}(\llbracket \pi \rrbracket)\text{)} \text{!} \forall x. \varphi \\
 \left[\left[\frac{\pi}{\frac{\vdash^{\mathcal{V}} \Gamma, \varphi[t/x]}{\vdash^{\mathcal{V}} \Gamma, \exists x. \varphi}} \right] \right]_{\forall \mathbf{I}} &= \Gamma^{\perp} \xrightarrow{\mathcal{T}(\mathcal{V})} \llbracket \pi \rrbracket \varphi[t/x] \xrightarrow{\mathcal{T}(\mathcal{V})} \exists x. \varphi \xrightarrow{\mathcal{T}(\mathcal{V})} \exists x. \varphi \text{?} \exists x. \varphi
 \end{aligned}$$

■ **Figure 10** Interpretation of the remaining rules of LK.

$$\begin{array}{ccccc}
 \text{!}\mathcal{A} \rightarrow \text{!!}\mathcal{A} & \text{!}\mathcal{A} \rightarrow \text{!}\mathcal{A} \otimes \text{!}\mathcal{A} & \text{?}\mathcal{A} \rightarrow \text{!}\mathcal{A} & \text{!}\mathcal{A} \otimes \text{!}\mathcal{B} \rightarrow \text{!(}\mathcal{A} \otimes \mathcal{B}\text{)} & \text{!}\mathcal{A} \wp \text{!}\mathcal{B} \rightarrow \text{!(}\mathcal{A} \wp \mathcal{B}\text{)} \\
 ((i, j), a) \mapsto (i, (j, a)) & (2i, a) \mapsto (1, (i, a)) & (i, (j, a)) \mapsto (j, (i, a)) & (j, (i, a)) \mapsto (i, (j, a)) & (j, (i, a)) \mapsto (i, (j, a)) \\
 & (2i + 1, a) \mapsto (2, (i, a)) & & &
 \end{array}$$

■ **Figure 11** Some win-isos with exponentials whose lifting are used in the interpretation.

get $\text{co}_{\forall x. \varphi} : \text{!}\forall x. \llbracket \varphi \rrbracket_{\mathcal{V}} \rightarrow \text{!!}\forall x. \llbracket \varphi \rrbracket_{\mathcal{V}}$ as a particular case of $\text{!}\mathcal{A} \rightarrow \text{!!}\mathcal{A}$ from Figure 11. We get $\text{co}_{\varphi \wedge \psi}$ and $\text{co}_{\varphi \vee \psi}$ by induction and composition with $\text{!}\mathcal{A} \otimes \text{!}\mathcal{B} \rightarrow \text{!(}\mathcal{A} \otimes \mathcal{B}\text{)}$, $\text{!}\mathcal{A} \wp \text{!}\mathcal{B} \rightarrow \text{!(}\mathcal{A} \wp \mathcal{B}\text{)}$.

Finally, $\text{co}_{\exists x. \llbracket \varphi \rrbracket_x}$ is obtained analogously to the contraction on the right of Figure 9.

► **Lemma 43.** *For any $(\mathcal{V} \uplus \{x\})$ -game \mathcal{A} , there is a winning $\mu_{\mathcal{A}, x} : \exists x. \text{!}\mathcal{A} \xrightarrow{\mathcal{V}\text{-Ga}} \text{!}\exists x. \mathcal{A}$.*

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 $\omega_{\mathcal{A}}$. ◀

We get $\text{co}_{\exists x. \llbracket \varphi \rrbracket_x}$ by induction, post-composition with $\text{?}\mu_{\llbracket \varphi \rrbracket, x}$ and distribution of ? over ! .

► **Proposition 44.** *For any $\varphi \in \text{Form}_{\Sigma}(\mathcal{V})$, there is a winning $\text{co}_{\llbracket \varphi \rrbracket_{\mathcal{V}}} : \llbracket \varphi \rrbracket_{\mathcal{V}} \xrightarrow{\mathcal{V}\text{-Ga}} \text{!}\llbracket \varphi \rrbracket_{\mathcal{V}}$.*

Combining Proposition 44 with other primitives (including $\text{!}\mathcal{A} \rightarrow \mathcal{A}$, playing copy-cat between \mathcal{A} and the 0th copy on the left, closed under receptivity), we get $\delta_{\llbracket \varphi \rrbracket_{\mathcal{V}}} : \llbracket \varphi \rrbracket_{\mathcal{V}} \rightarrow \llbracket \varphi \rrbracket_{\mathcal{V}} \otimes \llbracket \varphi \rrbracket_{\mathcal{V}}$ for $\varphi \in \text{Form}_{\Sigma}(\mathcal{V})$. We complete the interpretation in Figure 10, omitting W, which is by post-composition with $e_{\mathcal{A}}$ and silently using the isomorphism between winning Σ -strategies from 1 to $\Gamma \wp \mathcal{A}$ and from Γ^{\perp} to \mathcal{A} . This concludes the proof of Theorem 14.

6 Conclusion

For LK there is no hope of preserving unrestricted cut reduction without collapsing to a boolean algebra [13]. There are non-degenerate models for classical logic with an involutive negation, *e.g.* Führman and Pym's *classical categories* [9] with reduction only preserved in a lax sense; but our model does not preserve reduction even in this weaker sense. Besides, our semantics is infinitary: from the *structural dilemma* in [8] we obtained a proof of some $\exists x. \varphi$ with φ quantifier-free (no \forall bélard moves) yielding an infinite Σ -strategy (see Appendix B).

Both phenomena could be avoided by adopting a polarized model, abandoning however our faithfulness to the raw Herbrand content of proofs. It is a fascinating open question whether one can find a non-polarized model of classical first-order logic that remains finitary – this is strongly related to the actively investigated question of finding a strongly normalizing reduction strategy on syntaxes for expansion trees [15, 21, 16].

References

- 1 Samson Abramsky, Radha Jagadeesan, and Pasquale Malacaria. Full abstraction for PCF. *Inf. Comput.*, 163(2):409–470, 2000.
- 2 Franco Barbanera and Stefano Berardi. A symmetric lambda calculus for "classical" program extraction. In Masami Hagiya and John C. Mitchell, editors, *Theoretical Aspects of Computer Software, International Conference TACS '94, Sendai, Japan, April 19-22, 1994, Proceedings*, volume 789 of *Lecture Notes in Computer Science*, pages 495–515. Springer, 1994. doi:10.1007/3-540-57887-0_112.
- 3 Michael Barr. *-autonomous categories and linear logic. *Mathematical Structures in Computer Science*, 1(2):159–178, 1991.
- 4 Simon Castellan, Pierre Clairambault, Silvain Rideau, and Glynn Winskel. Games and strategies as event structures. *Logical Methods in Computer Science*, 13(3), 2017.
- 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*, 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 Gerhard Gentzen. Untersuchungen über das logische schließen. i. *Mathematische zeitschrift*, 39(1):176–210, 1935.
- 11 Philipp Gerhardy and Ulrich Kohlenbach. Extracting herbrand disjunctions by functional interpretation. *Arch. Math. Log.*, 44(5):633–644, 2005.
- 12 Jean-Yves Girard. Linear logic. *Theor. Comput. Sci.*, 50:1–102, 1987.
- 13 Jean-Yves Girard. A new constructive logic: Classical logic. *Mathematical Structures in Computer Science*, 1(3):255–296, 1991.
- 14 Kurt Gödel. Über eine bisher noch nicht benützte erweiterung des finiten standpunktes. *dialectica*, 12(3-4):280–287, 1958.
- 15 Willem Heijltjes. Classical proof forestry. *Ann. Pure Appl. Logic*, 161(11):1346–1366, 2010.
- 16 Stefan Hetzl and Daniel Weller. Expansion trees with cut. *CoRR*, abs/1308.0428, 2013.
- 17 Ulrich Kohlenbach. On the no-counterexample interpretation. *J. Symb. Log.*, 64(4):1491–1511, 1999.
- 18 Jean-Louis Krivine. Realizability in classical logic. *Panoramas et synthèses*, 27:197–229, 2009.
- 19 Olivier Laurent. Game semantics for first-order logic. *Logical Methods in Computer Science*, 6(4), 2010.
- 20 F William Lawvere. Adjointness in foundations. *Dialectica*, 23(3-4):281–296, 1969.

- 21 Richard McKinley. Proof nets for Herbrand's theorem. *ACM Trans. Comput. Log.*, 14(1):5, 2013.
- 22 Paul-André Mellies. Categorical semantics of linear logic. *Panoramas et synthèses*, 27:15–215, 2009.
- 23 Paul-André Mellies 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.
- 24 Dale A Miller. A compact representation of proofs. *Studia Logica*, 46(4):347–370, 1987.
- 25 Samuel Mimram. The structure of first-order causality. *Mathematical Structures in Computer Science*, 21(1):65–110, 2011.
- 26 Michel Parigot. Lambda-my-calculus: An algorithmic interpretation of classical natural deduction. In Andrei Voronkov, editor, *Logic Programming and Automated Reasoning, International Conference LPAR'92, St. Petersburg, Russia, July 15-20, 1992, Proceedings*, volume 624 of *Lecture Notes in Computer Science*, pages 190–201. Springer, 1992. doi:10.1007/BFb0013061.
- 27 Silvain Rideau and Glynn Winskel. Concurrent strategies. In *Proceedings of the 26th Annual IEEE Symposium on Logic in Computer Science, LICS 2011*, pages 409–418, 2011.
- 28 Robert A. G. Seely. Hyperdoctrines, natural deduction and the Beck condition. *Math. Log.*, 29(10):505–542, 1983.
- 29 Robert AG Seely. *Linear logic, *-autonomous categories and cofree coalgebras*. Ste. Anne de Bellevue, Quebec: CEGEP John Abbott College, 1987.
- 30 J.R. Shoenfield. *Mathematical Logic*. Addison-Wesley, 1967.
- 31 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.

A Counter-examples

In this section, we detail a few counter-examples referred to in the main text.

► **Example 45.** The neutral Σ -strategies $\sigma_1 = \begin{matrix} e_1^{e_1} & e_2^{e_2} \\ \forall & \exists \\ e_3^{f(e_1)} \end{matrix}$ and $\sigma_2 = \begin{matrix} e_1^{e_1} & e_2^{e_2} \\ \forall & \exists \\ e_3^{f(e_2)} \end{matrix}$, have no meet.

Assume they have a meet σ . Necessarily, since $e_1^{e_1} e_2^{e_2} \preceq \sigma_1, \sigma_2$, then σ must comprise the events-with-annotations $e_1^{e_1}$ and $e_2^{e_2}$. But we also have

$$\begin{matrix} e_1^c & & e_2^c \\ \forall & & \exists \\ & e_3^{f(c)} & \\ & \preceq \sigma_1, \sigma_2 & \end{matrix}$$

for any constant symbol c . Therefore, σ must also include event-with-annotation e_3^c . But t must be an instance of $f(e_1), f(e_2)$; and must instantiate to $f(c)$ for all constant symbol c . So t must have the form $f(e)$ for some $e \in [e_3]$, i.e., $e \in \{e_1, e_2, e_3\}$. It is direct to check that none of those options gives a neutral Σ -strategy that is below both σ_1 and σ_2 for \preceq .

► **Example 46.** Consider $\sigma : \forall_1 1 \multimap \forall_2 \forall_3 1$ and $\tau : \forall_2 \forall_3 1 \multimap \forall_4 1$ two Σ -strategies:

$$\begin{matrix} \forall_1 1 \multimap \forall_2 \forall_3 1 & \forall_2 \forall_3 1 \multimap \forall_4 1 \\ \forall_2 & \forall_4 \\ \forall_3 & \exists_2^{\forall_4} \\ \exists_1^{\forall_3} & \forall_5 \\ & \exists_3 \end{matrix}$$

where we omit the annotation of negative events, forced by Σ -receptivity.

Their composition has $\forall_4 \rightarrow \exists_1^c$, which is not a minimal strategy since c does not have \forall_4 as a free variable.

This counter-example also means that we do not have the adjunction expected from categorical logic $\exists_{\mathcal{V},x} \dashv \mathcal{T}(w_{\mathcal{V},x}) \dashv \forall_{\mathcal{V},x}$. More precisely, Lemma 40 cannot be strengthened into an equality. Indeed, note that $\tau = \mathbb{M}_{(\forall_2 \forall_3 1),1}^x (\exists_2^x \rightarrow \exists_3^c)$. On the other hand, $\tau \odot \sigma = \forall_4 \rightarrow \exists_1^c$, which cannot be of the form $\mathbb{M}_{\forall_1 1,1}^x$ – this construction would put no causal link from \forall_4 to \exists_1^c , since c does not involve the variable x .

The intuition behind this failure is that $\mathbb{M}_{A,B}^x$ only introduces causal links that follow occurrences of a variable x . However, after composition, we may end up with Σ -strategies that are not *minimal*, i.e., they have immediate causal links not reflecting directly a syntactic dependency. In other words, in order to get an adjunction as one would expect, only the term information would have to be retained – but our interpretation remembers more.

B Non-finiteness of the interpretation

From the infinitary primitives in the interpretation, it is natural to expect the interpretation to be infinitary. It was surprisingly difficult to find such an example, however one can do so by revisiting standard pathological examples in the proof theory of classical logic, having arbitrarily large normal forms.

More precisely, we construct an LK proof of the formula $\exists x. \top$ whose interpretation is infinite, despite the fact that there is no move by \forall bélard in the game.

Our starting point is the following proof:

$$\varpi = \frac{\frac{\frac{\text{Ax}}{\text{AI}} \frac{}{\vdash \varphi, \varphi^\perp} \quad \frac{\text{Ax}}{\text{AI}} \frac{}{\vdash \varphi, \varphi^\perp}}{\text{C}} \frac{}{\vdash \varphi \wedge \varphi, \varphi^\perp, \varphi^\perp} \quad \frac{\frac{\text{Ax}}{\text{AI}} \frac{}{\vdash \varphi, \varphi^\perp} \quad \frac{\text{Ax}}{\text{AI}} \frac{}{\vdash \varphi, \varphi^\perp}}{\text{C}} \frac{}{\vdash \varphi, \varphi, \varphi^\perp \wedge \varphi^\perp}}{\text{CUT}} \frac{}{\vdash \varphi \wedge \varphi, \varphi^\perp} \quad \frac{}{\vdash \varphi, \varphi^\perp \wedge \varphi^\perp}}{\vdash \varphi \wedge \varphi, \varphi^\perp \wedge \varphi^\perp}$$

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 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 $\varphi = \forall x. \perp \vee \exists y. \top$, which to shorten notations we will just write as $\forall \vee \exists$.

Omitting details, here is the interpretation of the left branch of ϖ (we omit term annotations, which always coincide with the unique predecessor for Eloïse's moves).

$$\left[\left[\frac{\frac{\text{Ax}}{\text{AI}} \frac{}{\vdash \varphi, \varphi^\perp} \quad \frac{\text{Ax}}{\text{AI}} \frac{}{\vdash \varphi, \varphi^\perp}}{\text{C}} \frac{}{\vdash \varphi \wedge \varphi, \varphi^\perp, \varphi^\perp} \right] \right] = \begin{array}{c} (\forall \vee \exists) \wedge (\forall \vee \exists) , (\exists \wedge \forall) \\ \forall_i \xrightarrow{\quad} \exists_{\langle 0,i \rangle} \\ \hline \forall_j \xrightarrow{\quad} \exists_{\langle 1,j \rangle} \\ \hline \forall_k \xrightarrow{\quad} \exists_k \leftarrow \exists_k \end{array}$$


5:20 The True Concurrency of Herbrand's Theorem

The second branch of ϖ is symmetric, so we do not make it explicit. Now, we interpret the CUT rule and the composition yields $\llbracket \varpi \rrbracket$ below.




It is interesting to note that although ϖ has arbitrarily large cut-free forms, the corresponding strategy only plays finitely many \exists moves for every \forall move. However, we are on the right path to finding an infinitary Σ -strategy.


The next step is to set (with s some unary function symbol) the proof ϖ_2 below with interpretation



We now use these to compute the interpretation of ϖ_3 , a cut between ϖ and ϖ_2 :




We are almost there. It suffices now to note that ϖ_3 provides a proof of $(\exists x. \top \implies \exists x. \top) \wedge (\exists x. \top \implies \exists x. \top)$. These two implications can be *composed* by cutting ϖ_3 against the following proof ϖ_4 :



Write ϖ_5 for the proof of $\exists x. \top \vee \forall y. \perp$ obtained by cutting ϖ_3 and ϖ_4 . The interpretation of ϖ_5 is the composition of $\llbracket \varpi_3 \rrbracket$ and $\llbracket \varpi_4 \rrbracket$, 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 \parallel 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, Eloïse responds to an initial \forall bélar move ∇ by playing simultaneously all $\exists^{s^n(\nabla)}$, for $n \geq 1$. Finally, cutting ϖ_5 against a proof of $\exists x. \top$ playing a constant symbol 0, we get a proof ϖ_6 of $\vdash \exists x. \top$ whose interpretation plays simultaneously all $\exists^{s^n(0)}$ for $n \geq 1$.

C Compactness

Restricting any winning Σ -strategy $\sigma : \llbracket \varphi \rrbracket$ to $\llbracket \varphi \rrbracket^\exists$ (ignoring \forall bélar's replications) yields $\sigma^\exists : \llbracket \varphi \rrbracket^\exists$, not necessarily finite. Yet, we will show that it has a finite *top-winning* sub-strategy.

A game \mathcal{A} is a **prefix** of \mathcal{B} if $|A| \subseteq |B|$, and all the structure coincides on $|A|$. Notice that $\llbracket \varphi \rrbracket^\exists$ embeds (subject to renaming) as a prefix of $\llbracket \varphi \rrbracket$. Keeping the renaming silent, we have:

► **Lemma 47.** *For any winning $\sigma : \llbracket \varphi \rrbracket$, setting*

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

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

Proof. Most conditions are direct. For $\sigma^\exists : \llbracket \varphi \rrbracket^\exists$ winning we use that for any \exists -maximal $x \in \mathcal{C}^\infty(\sigma^\exists)$, $x \in \mathcal{C}^\infty(\sigma)$ \exists -maximal as well: this follows from $\llbracket \varphi \rrbracket^\exists$ being itself \exists -maximal in $\llbracket \varphi \rrbracket$. ◀

As mentioned above, the extracted σ^\exists may not be finite! Indeed there are classical proofs for which our interpretation yields infinite strategies, even after removing \forall bélar's replications (see Appendix B). 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 $\sigma : \llbracket \varphi \rrbracket^\exists$ any Σ -strategy, we denote $\mathcal{C}^\forall(\sigma)$ the set of \forall -**maximal** configurations of σ , i.e., they can only be extended in σ by Eloïse moves – inheriting all structure from σ they correspond to its *sub-strategies*, as they are automatically receptive. The proof relies on:

5:22 The True Concurrency of Herbrand's Theorem

► **Lemma 48.** *Let X be a directed set of \forall -maximal configurations. Then, $\mathcal{W}_{\llbracket \varphi \rrbracket^\exists}(\bigcup X)$ is logically equivalent to $\bigvee_{x \in X} \mathcal{W}_{\llbracket \varphi \rrbracket^\exists}(x)$.*

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

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

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

is a tautology. By the compactness theorem there is a finite $X = \{x_1, \dots, x_n\} \subseteq \mathcal{C}^\forall(\sigma)$ such that $\bigvee_{x \in X} \mathcal{W}_{\llbracket \varphi \rrbracket^\exists}(x)[\lambda_\sigma]$ is a tautology – *w.l.o.g.* X is directed as $\mathcal{C}^\forall(\sigma)$ is closed under union. By Lemma 48 again, $\mathcal{W}_{\llbracket \varphi \rrbracket^\exists}(\bigcup X)[\lambda_\sigma]$ is a tautology. So, restricting σ to events $\bigcup X$ gives a top-winning finite sub-strategy of σ .

Although this argument is non-constructive, the extraction of a finite sub-strategy can still be performed effectively: Σ -strategies and their operations can be effectively presented, and the finite top-winning sub-strategy can be computed by Markov's principle.