



HAL
open science

The free bifibration on a functor

Bryce Clarke, Gabriel Scherer, Noam Zeilberger

► **To cite this version:**

| Bryce Clarke, Gabriel Scherer, Noam Zeilberger. The free bifibration on a functor. 2025. <hal-05357986v2>

HAL Id: hal-05357986

<https://inria.hal.science/hal-05357986v2>

Preprint submitted on 27 Dec 2025

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.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

THE FREE BIFIBRATION ON A FUNCTOR

BRYCE CLARKE ^a, GABRIEL SCHERER ^b, AND NOAM ZEILBERGER ^c

^a Tallinn University of Technology, Tallinn, Estonia
e-mail address: bryce.clarke@taltech.ee

^b Inria, IRIF, Université Paris Cité, Paris, France
e-mail address: gabriel.scherer@inria.fr

^c LIX, CNRS, Inria, École Polytechnique, Institut Polytechnique de Paris, Palaiseau, France
e-mail address: noam.zeilberger@polytechnique.edu

ABSTRACT. We consider the problem of constructing the free bifibration generated by a functor of categories $p : \mathcal{D} \rightarrow \mathcal{C}$. This problem was previously considered by Lamarche, and is closely related to the problem, considered by Dawson, Paré, and Pronk, of “freely adjoining adjoints” to a category. We develop a proof-theoretic approach to the problem, beginning with a construction of the free bifibration $\Lambda_p : \mathcal{Bif}(p) \rightarrow \mathcal{C}$ in which objects of $\mathcal{Bif}(p)$ are formulas of a primitive “bifibrational logic”, and arrows are derivations in a cut-free sequent calculus modulo a notion of permutation equivalence. We show that instantiating the construction to the identity functor generates a *zigzag double category* $\mathbb{Z}(\mathcal{C})$, which is also the free double category with companions and conjoints (or fibrant double category) on \mathcal{C} . The approach adapts smoothly to the more general task of building $(\mathcal{P}, \mathcal{N})$ -fibrations, where one only asks for pushforwards along arrows in \mathcal{P} and pullbacks along arrows in \mathcal{N} for some subsets of arrows, which encompasses Kock and Joyal’s notion of *ambifibration* when $(\mathcal{P}, \mathcal{N})$ form a factorization system. We establish a series of progressively stronger normal forms, guided by ideas of *focusing* from proof theory, and obtain a canonicity result under assumption that the base category is factorization preordered relative to \mathcal{P} and \mathcal{N} . This canonicity result allows us to decide the word problem and to enumerate relative homsets without duplicates. Finally, we describe several examples of a combinatorial nature, including a category of plane trees generated as a free bifibration over ω , and a category of increasing forests generated as a free ambifibration over Δ , which contains the lattices of noncrossing partitions as quotients of its fibers by the Beck-Chevalley condition.

INTRODUCTION

A functor $p : \mathcal{D} \rightarrow \mathcal{C}$ between two categories is a *bifibration* when, roughly speaking, objects of \mathcal{D} may be pushed and pulled along arrows of \mathcal{C} . If we visualize the category \mathcal{D} as living over the category \mathcal{C} and mapping downwards via the functor, then the structure of a bifibration on p allows the arrows of \mathcal{C} to be lifted upwards into the category \mathcal{D} , in two different kinds

Key words and phrases: bifibrations, double categories, free constructions, proof theory, combinatorics.

of ways. Formally, for any arrow $f : A \rightarrow B$ in \mathcal{C} and any object S in \mathcal{D} such that $p(S) = A$, there should be an object $f^+ S$ and an arrow $f_S : S \rightarrow f^+ S$ of \mathcal{D} such that $p(f_S) = f$,

$$\begin{array}{ccc} \mathcal{D} & & S \overset{f_S}{\dashrightarrow} f^+ S \\ p \downarrow & & \\ \mathcal{C} & & A \xrightarrow{f} B \end{array}$$

which are universal in the sense that for any arrow $g : B \rightarrow C$ in \mathcal{C} and arrow $\alpha : S \rightarrow T$ in \mathcal{D} such that $p(\alpha) = fg$,¹ there is a unique arrow $\beta : f^+ S \rightarrow T$ such that $\alpha = f_S \beta$.

$$\begin{array}{ccc} S \xrightarrow{\alpha} T & = & S \xrightarrow{f_S} f^+ S \overset{\beta}{\dashrightarrow} T \\ A \xrightarrow{f} B \xrightarrow{g} C & & A \xrightarrow{f} B \xrightarrow{g} C \end{array} \quad (0.1)$$

We refer to the object $f^+ S$ as the *pushforward* of S along f , and to the arrow $f_S : S \rightarrow f^+ S$ as a *+cartesian lifting* of f to S . Note that the universal property of +cartesian liftings ensures that the pushforward $f^+ S$ is determined up to unique isomorphism lying over the identity arrow id_B , which is what justifies speaking of “the” pushforward of S along f . Dually, for any arrow $g : B \rightarrow C$ in \mathcal{C} and object T in \mathcal{D} such that $p(T) = C$, there should be an object $g^- T$ and an arrow $\bar{g}_T : g^- T \rightarrow T$ of \mathcal{D} such that $p(\bar{g}_T) = g$,

$$\begin{array}{ccc} g^- T \overset{\bar{g}_T}{\dashrightarrow} T & & \\ & & B \xrightarrow{g} C \end{array}$$

again universal in the sense that for any arrow $f : A \rightarrow B$ in \mathcal{C} and arrow $\alpha : S \rightarrow T$ in \mathcal{D} such that $p(\alpha) = fg$, there is a unique arrow $\beta : S \rightarrow g^- T$ such that $\alpha = \beta \bar{g}_T$.

$$\begin{array}{ccc} S \xrightarrow{\alpha} T & = & S \overset{\beta}{\dashrightarrow} g^- T \xrightarrow{\bar{g}_T} T \\ A \xrightarrow{f} B \xrightarrow{g} C & & A \xrightarrow{f} B \xrightarrow{g} C \end{array} \quad (0.2)$$

We refer to the object $g^- T$ as the *pullback* of T along g , and to the arrow $\bar{g}_T : g^- T \rightarrow T$ as a *--cartesian lifting* of g to T . Again, the universal property of --cartesian liftings ensures that the pullback $g^- T$ is determined up to unique isomorphism lying over id_B .

An immediate consequence of the definitions is that if $p : \mathcal{D} \rightarrow \mathcal{C}$ is a bifibration then there is a one-to-one correspondence

$$\begin{array}{ccc} S \longrightarrow f^- T & \iff & S \longrightarrow T & \iff & f^+ S \longrightarrow T \\ & & \iff & & \\ A \xrightarrow{\text{id}_A} A & & A \xrightarrow{f} B & & B \xrightarrow{\text{id}_B} B \end{array} \quad (0.3)$$

¹Here and throughout the paper we notate composition of arrows by juxtaposition in diagrammatic order, except in a few places where we explicitly write the composition symbol \circ to mean $g \circ f = fg$.

of arrows in \mathcal{D} living over the corresponding arrows in \mathcal{C} . As a consequence, every arrow $f : A \rightarrow B$ of \mathcal{C} induces an adjunction

$$\begin{array}{ccc} & f^+ & \\ \mathcal{D}_A & \xrightarrow{\quad} & \mathcal{D}_B \\ & \perp & \\ & f^- & \end{array}$$

realized by the operations of pushing and pulling along f , where \mathcal{D}_A and \mathcal{D}_B are the *fiber categories* of A and B relative to the functor p , defined as the subcategories of \mathcal{D} spanned by the arrows living over the identities id_A and id_B in \mathcal{C} . More precisely, the functors f^+ and f^- can be defined given any specific choice of cartesian liftings in \mathcal{D} , corresponding to a choice of bifibrational structure on $p : \mathcal{D} \rightarrow \mathcal{C}$. Such a bifibrational structure is called a “cleavage” or “cleaving”, and p is said to be a *cloven* bifibration. We will always be working with bifibrational structures in this paper and so we usually omit the adjective “cloven”.

The categorical notion of bifibration was originally introduced by Grothendieck [20, VI], together with *fibration* and *cofibration*, the latter nowadays more commonly called *opfibration*. A functor $p : \mathcal{D} \rightarrow \mathcal{C}$ is a fibration (respectively opfibration) when one can pull (respectively, push) objects of \mathcal{D} along arrows of \mathcal{C} , so that p is a bifibration iff it is both a fibration and an opfibration. As established by Grothendieck and which motivated the original definition, the data of a cloven bifibration $p : \mathcal{D} \rightarrow \mathcal{C}$ is equivalent to the data of a pseudofunctor $F : \mathcal{C} \rightarrow \text{Adj}$ into the 2-category of small categories and adjunctions. A bifibration over \mathcal{C} is recovered from F by defining the *category of elements* $\int_{\mathcal{C}} F$ (also called the *Grothendieck construction*) to have objects given by pairs (A, S) of an object $A \in \mathcal{C}$ and an object $S \in F(A)$, and arrows $(A, S) \rightarrow (B, T)$ given by pairs (f, α) of an arrow $A \rightarrow B \in \mathcal{C}$ and an arrow $\alpha : F(f)^+(S) \rightarrow T \in F(B)$ or equivalently an arrow $\alpha : S \rightarrow F(f)^-(T) \in F(A)$, where we write $F(f)^+ : F(A) \rightarrow F(B)$ and $F(f)^- : F(B) \rightarrow F(A)$ respectively for the left and right adjoint associated to the adjunction $F(f)$. The category of elements is equipped with an evident projection functor $\pi_F : \int_{\mathcal{C}} F \rightarrow \mathcal{C}$, which is a bifibration.

We are interested in the problem of constructing the free bifibration on a functor, meaning the following. Let $p : \mathcal{D} \rightarrow \mathcal{C}$ be an arbitrary functor. The *free bifibration on p* is a category $\text{Bif}(p)$ equipped with a functor $\Lambda_p : \text{Bif}(p) \rightarrow \mathcal{C}$ that is a bifibration, together with a functor $\eta_p : \mathcal{D} \rightarrow \text{Bif}(p)$ commuting with the projections to \mathcal{C} :

$$\begin{array}{ccc} \mathcal{D} & \xrightarrow{\eta_p} & \text{Bif}(p) \\ & \searrow p & \swarrow \Lambda_p \\ & \mathcal{C} & \end{array}$$

Moreover, this data is universal in the sense that for any other bifibration $q : \mathcal{E} \rightarrow \mathcal{C}$ equipped with a functor $\theta : \mathcal{D} \rightarrow \mathcal{E}$ such that $q \circ \theta = p$, there is a unique morphism of bifibrations making the diagram below commute:

$$\begin{array}{ccc} \mathcal{D} & \xrightarrow{\eta_p} & \text{Bif}(p) \\ & \searrow p & \downarrow \Lambda_p \\ & \mathcal{C} & \end{array} \quad \begin{array}{ccc} & & \mathcal{E} \\ & \swarrow \theta & \uparrow \llbracket - \rrbracket_{\theta} \\ & & \mathcal{E} \\ & \swarrow \theta & \downarrow q \text{ bifibration} \\ & \mathcal{C} & \end{array} \tag{0.4}$$

By *morphism of bifibrations*, we mean a functor $\llbracket - \rrbracket_\theta : \mathcal{Bif}(p) \rightarrow \mathcal{E}$ sending chosen \pm -cartesian liftings to chosen \pm -cartesian liftings. A standard argument implies the universal property of $\Lambda_p : \mathcal{Bif}(p) \rightarrow \mathcal{C}$ extends to bifibrations over any base category into which \mathcal{C} can be mapped, in the sense that $\llbracket - \rrbracket_\theta$ below is uniquely determined from $\theta : \mathcal{D} \rightarrow \mathcal{E}$, $q : \mathcal{E} \rightarrow \mathcal{B}$, and $\delta : \mathcal{C} \rightarrow \mathcal{B}$:

$$\begin{array}{ccccc}
 \mathcal{D} & \xrightarrow{\eta_p} & \mathcal{Bif}(p) & & \\
 & \searrow & \downarrow \Lambda_p & \searrow \llbracket - \rrbracket_\theta & \\
 & & \mathcal{C} & & \mathcal{E} \\
 & \searrow p & & & \downarrow q \text{ bifibration} \\
 & & & & \mathcal{B} \\
 & & & \nearrow \delta & \\
 & & & & \mathcal{C}
 \end{array}
 \tag{0.5}$$

In particular, (0.5) can be reduced to (0.4) by first pulling back q along δ to a bifibration over \mathcal{C} , using the fact that bifibrations are closed under pullback along arbitrary functors.

Whereas the free fibration $\mathcal{C} \downarrow p \rightarrow \mathcal{C}$ and the free opfibration $p \downarrow \mathcal{C} \rightarrow \mathcal{C}$ on a functor have simple and well-known descriptions, the free bifibration has been relatively little studied, and is a more complex object. As far as we are aware there is only one direct construction in the literature, in an unpublished manuscript by François Lamarche [32]. This construction was motivated by his work on a model of type theory in which dependent types were Grothendieck bifibrations [33], and identity types were modelled as bifibrations $P(\mathcal{C}) \rightarrow \mathcal{C} \times \mathcal{C}$ from a certain “path category”, of which a detailed analysis was given in the manuscript [32]. Otherwise, the problem of building the free bifibration on a functor is also closely related to the problem studied by Dawson, Paré, and Pronk [14, 15] of freely adjoining right adjoints to a category, which may be seen as a weaker version of building the free groupoid over a category. Their construction, which generalizes Schanuel and Street’s definition of the *free adjunction* [48], takes a category \mathcal{C} and embeds it into a 2-category $\Pi_2(\mathcal{C})$ in which every arrow of the original category is equipped with a right adjoint.

The connection between these problems is revealed by taking the free bifibration on the identity functor $\text{id}_{\mathcal{C}}$: as we will see, this defines a category $\mathcal{Z}(\mathcal{C}) := \mathcal{Bif}(\text{id}_{\mathcal{C}})$ bifibered over \mathcal{C} that may be equipped with the structure of a double category $\mathbb{Z}(\mathcal{C}) := \mathcal{Z}(\mathcal{C}) \rightrightarrows \mathcal{C}$. We call $\mathbb{Z}(\mathcal{C})$ the *zigzag double category of \mathcal{C}* , since its vertical arrows (corresponding to the objects of $\mathcal{Z}(\mathcal{C})$) are zigzags of arrows in \mathcal{C} . In fact, $\mathcal{Z}(\mathcal{C})$ is equivalent to Lamarche’s path category $P(\mathcal{C})$, and as he observed [32, p.22], any free bifibration may be recovered from this double category by the formula $\Lambda_p \equiv \text{tgt}^+ \text{src}^- p$, or diagrammatically:

$$\begin{array}{ccc}
 \mathcal{D} & \longleftarrow & \mathcal{Bif}(p) \\
 p \downarrow & \lrcorner & \downarrow \Lambda_p \\
 \mathcal{C} & \xleftarrow{\text{src}} & \mathcal{Z}(\mathcal{C}) \xrightarrow{\text{tgt}} \mathcal{C}
 \end{array}
 \tag{0.6}$$

In turn, taking the underlying vertical 2-category of $\mathbb{Z}(\mathcal{C})$ recovers a 2-category equivalent to Dawson, Paré, and Pronk’s $\Pi_2(\mathcal{C})$, thus reducing the problem of freely adjoining right adjoints to the problem of understanding the free bifibration on the identity functor.

Our approach to the problem differs from prior work in that we study it from a proof-theoretic perspective. It is natural to interpret the pushforward f^+ and pullback g^- operations

as unary logical connectives, abstractions of existential and universal quantification as well as of diamond and box modalities. The rules of this primitive “bifibrational logic” may be presented by a simple sequent calculus, leading us to a description of the free bifibration wherein the objects of $\mathcal{Bif}(p)$ are formulas and its arrows are equivalence classes of sequent calculus derivations. See Figure 1 for an example derivation in this sequent calculus parameterized by a specific functor $p : \mathcal{D} \rightarrow \mathcal{C}$, representing an arrow $\beta : f^- f^+ X \rightarrow h^- g^+ Y$ in $\mathcal{Bif}(p)$ such that $\Lambda_p(\beta) = \text{id}_A$. The close connection with the zigzag double category also leads to an isomorphic representation of derivations as stacks of double cells in $\mathbb{Z}(\mathcal{C})$, as well as to a string diagram calculus.

It should be emphasized that we do *not* postulate the Beck-Chevalley condition. Although bifibrations satisfying the BC condition are commonly used in categorical logic after Lawvere [36] to model extensions of predicate logic with either substitution and existential quantification or substitution and universal quantification (see, e.g., Jacobs [24]), we are rather interested in free bifibrations where the BC condition fails. These model more general situations including ones where the adjoint pair of pushforward and pullback operations $f^+ \dashv f^-$ themselves represent generalized existential and universal quantifiers $\exists_f := f^+$, $\forall_f := f^-$, and hence by alternating these operations one can in some sense define objects of arbitrary quantifier complexity.




Figure 1: Upper left: an example functor, depicting objects and arrows of the category above over their images below. Lower left: a sequent calculus proof representing an arrow of $\mathcal{Bif}(p)$. Right: the same proof depicted as a stack of double cells in $\mathbb{Z}(\mathcal{C})$ acting on an arrow in \mathcal{D} , overlaid with the corresponding string diagram.

One of our original motivations for studying this problem from this perspective was the observation, made independently by Mellies and Zeilberger [40, 41, 42] and by Licata, Shulman, and Riley [37], that substructural and modal logics may be naturally modelled in certain bifibrations: specifically bifibrations of monoidal closed categories or multicategories in which the base moreover contains an object with some algebraic structure. As a typical example, any monoid (A, m, e) in the base of a bifibration of multicategories $p : \mathcal{E} \rightarrow \mathcal{B}$ generates a monoidal closed fiber category \mathcal{E}_A , with the tensor product and internal hom defined by pushing and pulling along the multiplication map $m : A, A \rightarrow A$. For more recent work in this spirit see Nicolas Blanco’s PhD thesis [8] developing the theory of bifibrations of polycategories as models of classical linear logic, as well as Shulman [50] on LNL doctrines, which axiomatize a wide variety of categorical structures and type theories. The problem of constructing free bifibrations may therefore be seen as a natural problem in proof theory (cf. [50, §8]), and we were motivated to study the “pure” version of the problem for functors of categories as an eventual stepping stone to the more sophisticated case of monoidal closed categories or generalized multicategories and polycategories. In the present paper we do already consider a simple but practically significant generalization of the problem, namely to free $(\mathcal{P}, \mathcal{N})$ -fibrations where one only asks for $+$ -cartesian liftings (resp. $-$ -cartesian liftings) of the arrows in \mathcal{P} (resp. \mathcal{N}). In particular when the pair $(\mathcal{P}, \mathcal{N})$ form a factorization system for the base category this defines what Joachim Kock and André Joyal refer to as an *ambifibration*.

As it turns out, even very simple functors can generate free bifibrations (and free ambifibrations) with surprisingly rich combinatorial structure. For instance, let $p_2 : \mathbf{1} \rightarrow \mathbf{2}$ be the functor from the terminal category to the interval category sending the unique object of $\mathbf{1}$ to the initial object of $\mathbf{2}$:

$$\begin{array}{ccc} \mathbf{1} & & * \\ p_2 \downarrow & & \\ \mathbf{2} & & 0 \xrightarrow{f} 1 \end{array}$$

Now consider the free bifibration generated by p_2 . It is easy to see that any object of $\mathcal{Bif}(p_2)$ must be isomorphic to an alternating sequence of pushes and pulls along the arrow f starting at the point $*$, with even-length sequences yielding an object in the fiber over 0, and odd-length sequences an object in the fiber over 1. Let us write $n := (f^- f^+)^n *$ and $n' := f^+ (f^- f^+)^n *$ for the objects in the respective fibers (e.g., $3 = f^- f^+ f^- f^+ f^- f^+ *$). From the fact that the morphisms in the fibers are generated by the unit and counit of the push-pull adjunction $f^+ \dashv f^-$, with a bit of mental exercise (or familiarity with Schanuel and Street’s paper [48], building on Auderset’s earlier work [4]) one can check that an arrow $m \rightarrow n$ in $\mathcal{Bif}(p_2)_0$ should correspond to an order-preserving map $\{1, \dots, m\} \rightarrow \{1, \dots, n\}$. Indeed, $\mathcal{Bif}(p_2)_0 \cong \Delta$ is equivalent to the (augmented) simplex category of finite ordinals and order-preserving maps, while $\mathcal{Bif}(p_2)_1 \cong \Delta_\perp$ is equivalent to the category of non-empty finite ordinals and order-and-least-element-preserving maps. The push-pull adjunction of the free bifibration instantiates to the free / forgetful adjunction

$$\begin{array}{ccc} & L & \\ \Delta & \xrightarrow{\quad} & \Delta_\perp \\ & \perp & \\ & R & \end{array}$$

where L freely adds a least element \perp to turn an order-preserving map $m \rightarrow n$ into an order-and-least-element-preserving map $m' \rightarrow n'$, while the right adjoint R interprets m'

as $1+m$ and forgets that a map is \perp -preserving. This much may be unshocking to readers familiar with the free adjunction \mathcal{A} , which can be obtained by freely adjoining adjoints to the interval category, $\mathcal{A} = \Pi_2(2)$, as Dawson, Paré, and Pronk already observed [14, p.139]—although we emphasize the free bifibration is another way of organizing the data.

More strikingly, this example generalizes in the following way. Replace the interval category 2 by the ordinal ω (that is, the total order on the natural numbers) and consider the free bifibration generated by the functor $p_\omega : 1 \rightarrow \omega$ mapping the point to 0 , as depicted below:

$$\begin{array}{ccc} 1 & & * \\ p_\omega \downarrow & & \\ \omega & & 0 \xrightarrow{f_0} 1 \xrightarrow{f_1} 2 \xrightarrow{f_2} \dots \end{array}$$

What is the fiber of $\mathcal{Bif}(p_\omega)$ over 0 ? Objects are isomorphic to sequences of pushes and pulls along f_i 's, such that, reading a formula from right to left, the running total of pulls never exceeds the running total of pushes, and the total numbers of pushes and pulls are equal. In other words, up to isomorphism the objects of $\mathcal{Bif}(p_\omega)_0$ may be read as *Dyck words*. Via a well-known bijection (see Stanley [52, V2:169–170]), Dyck words are in one-to-one correspondence with rooted plane trees, and thus $\mathcal{Bif}(p_\omega)_0$ may be interpreted as a category of trees. Under a natural encoding of plane trees as functors $T : \omega^{op} \rightarrow \Delta$, it turns out that this fiber category of the free bifibration is equivalent to a full subcategory of $[\omega^{op}, \Delta]$, and embeds as a wide subcategory into a category of finite plane trees defined by Joyal [26].

An example of a different nature is obtained by considering the free $(\Delta_{\text{epi}}, \Delta_{\text{mono}})$ -fibration generated by the inclusion $i : \mathbb{N} \rightarrow \Delta$ of the set of natural numbers (seen as a discrete category) into the simplex category. This is a free ambifibration, and its universal property ensures that it is equipped with a canonical morphism into the so-called “fat Delta” category introduced by Kock [29]. The category $\mathcal{Bif}(i, \Delta_{\text{epi}}, \Delta_{\text{mono}})$ appears to have a rich combinatorial structure that we only begin to analyze here. For example, the *lattices of noncrossing partitions* (see Stanley [52, V1:515, V2:226]) are recovered by quotienting this free ambifibration by the Beck-Chevalley condition.

Outline of the paper. After introducing the sequent calculus in Section 1 and after defining and analyzing the double category of zigzags in Section 2, the technical core of the paper is in Section 3, where we establish a series of progressively stronger normal form results, guided by ideas from proof theory. The strongest of these (Theorem 3.23), which holds under the assumption that the base category \mathcal{C} is *factorization preordered* (FP), leads to an inductive definition of the relative homsets of the free bifibration that does not use any kind of quotient by an equivalence relation, representing arrows by *maximally multifocused* derivations. The FP condition also showed up in Dawson, Paré, and Pronk’s work on the $\Pi_2(\mathcal{C})$ construction, where they observed that equality of 2-cells is undecidable in general but is decidable when \mathcal{C} is factorization preordered [15]. As a corollary of our canonicity theorem, we obtain both an analogous decidability result for equality of arrows in $\mathcal{Bif}(p)$ as well as a procedure for enumerating relative homsets without duplicates. Section 4 briefly describes how our constructions and results adapt to free $(\mathcal{P}, \mathcal{N})$ -fibrations. This is technically straightforward but has the useful consequence that one can weaken the assumption needed for canonicity, only requiring that the base category be factorization preordered relative to each class \mathcal{P} and \mathcal{N} independently. Finally, in Section 5 we use the tools developed in the paper to analyze

the three examples described above. Section 6 concludes and suggests some directions for further exploration.

Remark. This work mixes domains and techniques that will be familiar to different audiences, and we expect most readers to be unfamiliar with some of them. We tried to find a presentation that would be accessible to people only familiar but not expert in category theory, and people unfamiliar with sequent calculus and proof theory. We ask the reader for forgiveness if they find the exposition too basic in places, for the small doses of redundancy this introduces, and for the length of the resulting presentation.

1. SEQUENT CALCULUS

In this section we will explain how the free bifibration $\Lambda_p : \mathcal{Bif}(p) \rightarrow \mathcal{C}$ on an arbitrary functor $p : \mathcal{D} \rightarrow \mathcal{C}$ may be constructed proof-theoretically, as an inductively-defined sequent calculus for a primitive “bifibrational logic”, starting from the data of p . The idea will be to take the objects of $\mathcal{Bif}(p)$ to be formulas of the sequent calculus and its arrows to be proofs modulo a notion of *permutation equivalence* to be defined below. This sequent calculus is cut-free, but the cut rule is admissible, giving a definition of composition of arrows in the category $\mathcal{Bif}(p)$. We will establish a series of properties leading to the main result of this section (Theorem 1.17), that the sequent calculus indeed gives a presentation of the free bifibration on p .

1.1. Formulas and unary sequents. To define the formulas of the logic, we find it helpful to introduce the judgment $S \sqsubset A$ to mean that S is a formula (= object of $\mathcal{Bif}(p)$) lying over (or “refining”) the object A of \mathcal{C} . We usually refer to formulas simply as “formulas”, but sometimes for clarity we refer to them as *bifibrational formulas*, being formulas of bifibrational logic. Bifibrational formulas are defined inductively by the following rules:

$$\frac{X \in \mathcal{D} \quad p(X) = A}{X \sqsubset A} \qquad \frac{S \sqsubset A \quad f : A \rightarrow B}{f^+ S \sqsubset B} \qquad \frac{g : B \rightarrow C \quad T \sqsubset C}{g^- T \sqsubset B}$$

Here we are using standard notation for inference rules: the judgment below the horizontal bar is the *conclusion* of the inference rule, and the judgments above the bar are its *premises*. Each rule asserts that if its premises are valid, then its conclusion is valid. We refer to formulas of the form X as *atomic formulas*, and formulas of the form $f^+ S$ or $g^- T$ as *push formulas* and *pull formulas* respectively. All of the formulas S' appearing in the inductive construction of a bifibrational formula S are said to be *subformulas* of S . We denote the subformula relation $S' \preceq S$, and the strict subformula relation $S' < S$.

In set-theoretic terms, each object A of \mathcal{C} determines a set $\text{Frm}_A = \{S \mid S \sqsubset A\}$ of formulas lying over A , and the collection of sets of formulas $(\text{Frm}_A)_{A \in \mathcal{C}}$ is the least family of sets closed under the following three conditions:

- if $X \in \mathcal{D}$ and $p(X) = A$ then $X \in \text{Frm}_A$;
- if $S \in \text{Frm}_A$ and $f : A \rightarrow B$ is an arrow of \mathcal{C} then $f^+ S \in \text{Frm}_B$;
- if $T \in \text{Frm}_C$ and $g : B \rightarrow C$ is an arrow of \mathcal{C} then $g^- T \in \text{Frm}_B$.

Besides the judgment $S \sqsubset A$ describing well-formedness of formulas, the main judgment of interest of the calculus is the *unary sequent arrow* of the form

$$S \xRightarrow[h]{} T \tag{1.1}$$

where $S \sqsubset A$ and $T \sqsubset C$ are formulas lying over objects A and C respectively, and where $h : A \rightarrow C$ is an arrow of \mathcal{C} . For comparison, it is worth keeping in mind traditional sequents of intuitionistic sequent calculus, which take the form

$$P_1, \dots, P_n \Longrightarrow Q$$

(often written with \vdash in place of an arrow) and which may be read as asserting an entailment from the conjunction of the formulas on the left-hand side to the formula on the right-hand side. In the bifibrational calculus, unary sequents of the form (1.1) may be read as asserting an entailment from S to T “fibered over h ”. We refer to S as the *left side*, to T as the *right side*, and to h as the *base* of the sequent. The definition of the sequent calculus will ensure that a proof of the judgment $S \Longrightarrow_h T$ corresponds to an arrow $\alpha : S \rightarrow T$ in $\mathcal{Bif}(p)$ such that $\Lambda_p(\alpha) = h$.

1.2. Inference rules and derivations. The calculus includes four logical inference rules:

$$\frac{S \Longrightarrow T}{f^+ S \Longrightarrow_g T} \text{L}f^+ \quad \frac{S' \Longrightarrow S}{S' \Longrightarrow_{f'f} f^+ S} \text{R}f^+ \quad \frac{T \Longrightarrow T'}{g^- T \Longrightarrow_{gg'} T'} \text{L}g^- \quad \frac{S \Longrightarrow T}{S \Longrightarrow_f g^- T} \text{R}g^-$$

together with an *initial axiom* for every arrow of \mathcal{D} :

$$\frac{\delta : X \rightarrow Y \in \mathcal{D} \quad p(\delta) = f}{X \Longrightarrow_f Y} \delta$$

In general in sequent calculus, one can define a *derivation* to be a finite rooted tree whose edges are labelled by sequents, and whose nodes are labelled by valid instances of the inference rules. Moreover, one can distinguish between *open* and *closed* derivations, where open derivations have premises corresponding to open edges of the tree. Another name for a closed derivation is a *proof*. A special property of the bifibrational sequent calculus is that every inference rule has at most one premise, so that derivations are completely linear (i.e., they are just *lists* rather than proper trees), and a derivation is closed just in case it is terminated by a unique initial axiom.

The overall format of the rules follows the tradition of sequent calculus with left rules and right rules for every connective, which in this case are the pushforward f^+ and pullback g^- connectives parameterized by arrows f and g of \mathcal{C} .

In set-theoretic terms, the rules may be seen as specifying closure conditions that inductively define a family of sets of proofs of judgments. For example, the $\text{L}f^+$ rule says that for any formulas S and T and arrows f and g , given a proof of $S \Longrightarrow_{fg} T$ there should be a proof of $f^+ S \Longrightarrow_g T$. Note that the rules have implicit side-conditions that the judgments are well-formed. For example, in the $\text{L}f^+$ rule, the arrows f and g should be composable, with $f : A \rightarrow B$ and $g : B \rightarrow C$ for some objects A, B, C such that $S \sqsubset A$ and $T \sqsubset C$.

We refer the reader back to Figure 1 from the Introduction, which gives an example derivation in the sequent calculus relative to a simple functor.

Finally, we emphasize that the underlying base arrows are more than a mere annotation on the sequent judgments: they really do restrict the applicability of the inference rules. For example, applying the $\text{L}f^+$ rule to a derivation of a judgment $S \Longrightarrow_h T$ requires factoring the arrow as a composite $h = fg$.

We do not include explicit identity and cut rules in the definition of the sequent calculus,

$$\frac{S \sqsubset A}{S \Longrightarrow_S S} \text{id}_S \qquad \frac{S \Longrightarrow_f U \quad U \Longrightarrow_g T}{S \Longrightarrow_{fg} T} \text{cut}$$

but these rules are *admissible* in the proof-theoretic sense as we will explain below. Since derivations are cut-free, they satisfy a variation of the standard subformula property:

Proposition 1.1 (Subformula property). *A derivation of $S \Longrightarrow_f T$ involves only sequents of the form $S' \Longrightarrow_g T'$ where $S' \leq S$ and $T' \leq T$ are subformulas of S and T .*

We now define closed derivations corresponding to the identity rule below, deferring the treatment of cut to Section 1.6.

Definition 1.2 (identity). For any formula $S \sqsubset A$, there is a proof id_S of $S \Longrightarrow_{\text{id}_A} S$ constructed by induction on S as follows:

$$\begin{aligned} \frac{X \in \mathcal{D} \quad p(X) = A}{X \sqsubset A} &\Longrightarrow \text{id}_X := \frac{}{X \Longrightarrow_{\text{id}_A} X} \text{id}_X \\ \\ \frac{S \sqsubset A \quad f : A \rightarrow B}{f^+ S \sqsubset B} &\Longrightarrow \text{id}_{f^+ S} := \frac{\frac{\frac{}{S \Longrightarrow_S S} \text{id}_S}{S \Longrightarrow_{\text{id}_A} S} \text{R}f^+}{f^+ S \Longrightarrow_{\text{id}_B} f^+ S} \text{L}f^+ \\ \\ \frac{g : B \rightarrow C \quad T \sqsubset C}{g^- T \sqsubset B} &\Longrightarrow \text{id}_{g^- T} := \frac{\frac{\frac{\frac{}{T \Longrightarrow_T T} \text{id}_T}{g^- T \Longrightarrow_{\text{id}_C} T} \text{L}g^-}{g^- T \Longrightarrow_{\text{id}_B} g^- T} \text{R}g^- \end{aligned}$$

□

Observe that proofs corresponding to the unit and counit of the canonical family of adjunctions $f^+ \dashv f^-$ of a bifibration may be constructed as

$$\frac{\frac{\frac{}{S \Longrightarrow_S S} \text{id}_S}{S \Longrightarrow_{\text{id}_A} S} \text{R}f^+}{S \Longrightarrow_{\text{id}_A} f^- f^+ S} \text{R}f^- \qquad \frac{\frac{\frac{}{U \Longrightarrow_U U} \text{id}_U}{f^- U \Longrightarrow_{\text{id}_B} U} \text{L}f^-}{f^+ f^- U \Longrightarrow_{\text{id}_B} U} \text{L}f^+$$

where we make use of the identity derivations from Definition 1.2 to terminate the proofs.

$$\begin{array}{c}
\frac{\delta : X \rightarrow Y \in \mathcal{D} \quad p(\delta) = f}{\delta : X \xRightarrow{f} Y} \\
\\
\frac{\alpha : S \xRightarrow{fg} T}{f \setminus_g \alpha : f^+ S \xRightarrow{g} T} \qquad \frac{\alpha : S' \xRightarrow{f'} S}{\alpha.f : S' \xRightarrow{f'f} f^+ S} \\
\\
\frac{\alpha : T \xRightarrow{g'} T'}{\bar{g}.\alpha : g^- T \xRightarrow{gg'} T'} \qquad \frac{\alpha : S \xRightarrow{fg} T}{\alpha.f/\bar{g} : S \xRightarrow{f} g^- T}
\end{array}$$

Figure 2: Term-annotated inference rules of the bifibrational calculus.

1.3. A term syntax for proofs. A term syntax is a compact notation to describe derivations that can be convenient for defining some operations and equations. We introduce here a syntax of terms generated by the following grammar:

$$\begin{array}{l}
\alpha ::= \delta \quad \text{for } \delta : X \rightarrow Y \in \mathcal{D} \\
\quad | f \setminus_g \alpha \\
\quad | \alpha.f \\
\quad | \bar{g}.\alpha \\
\quad | \alpha.f/\bar{g}
\end{array}$$

which are assigned to derivations as shown in Figure 2.

Formally, this defines a new judgment $\alpha : S \xRightarrow{h} T$ whose derivations are in one-to-one correspondence with derivations of $S \xRightarrow{h} T$. In other words, for a judgment $S \xRightarrow{h} T$, a derivation of the judgment uniquely determines a term α and a derivation of $\alpha : S \xRightarrow{h} T$. We will only consider *valid* terms, that do correspond to a derivation, and will use the term α and the underlying derivation interchangeably.

We sometimes call *multiplications* the term-formers $\alpha.f$ and $\bar{g}.\alpha$ and *divisions* the term-formers $f \setminus_g \alpha$ and $\alpha.f/\bar{g}$. The reason why our division notation takes an explicit subscript f or g is that its meaning depends on a choice of decomposition of the arrow h underlying α as $h = fg$; there is no such ambiguity for multiplications.

Examples. The terms for the identity derivations can be defined recursively as follows:

$$\text{id}_X := \text{id}_X \in \mathcal{D} \qquad \text{id}_{f^+ S} := f \setminus_{\text{id}_B} (\text{id}_S.f) \qquad \text{id}_{g^- T} := (\bar{g}.\text{id}_T)_{\text{id}_B/\bar{g}}$$

and the unit and co-unit derivations as follows:

$$\eta_S := (\text{id}_S.f)_{\text{id}_A/\bar{f}} \qquad \epsilon_T := f \setminus_{\text{id}_B} (\bar{f}.\text{id}_U)$$

These examples use subterms of the form $\text{id}_S.f$ and $\bar{f}.\text{id}_S$, which we will later show correspond to cartesian liftings in the free bifibration:

$$f_S : S \xRightarrow{f} f^+ S \quad := \quad \text{id}_S.f \qquad \bar{g}_T : g^- T \xRightarrow{g} T \quad := \quad \bar{g}.\text{id}_T$$

1.4. Interpreting formulas and proofs in any bifibration on p . In this section we suppose given an arbitrary bifibration $q : \mathcal{E} \rightarrow \mathcal{C}$ equipped with a functor $\theta : \mathcal{D} \rightarrow \mathcal{E}$ such that $q \circ \theta = p$. As explained in the Introduction, we expect θ to factor uniquely via the free bifibration in the diagram below:

$$\begin{array}{ccc}
 \mathcal{D} & \xrightarrow{\eta_p} & \mathcal{Bif}(p) \\
 & \searrow \theta & \downarrow \Lambda_p \\
 & & \mathcal{E} \\
 & \searrow p & \downarrow q \text{ bifibration} \\
 & & \mathcal{C}
 \end{array}$$

We describe now how to interpret formulas and proofs of the sequent calculus as objects and arrows of \mathcal{E} , which will determine the functor $\llbracket - \rrbracket_\theta$ once we have completed the definition of the category $\mathcal{Bif}(p)$.

We interpret any formula S as an object $\llbracket S \rrbracket_\theta$ of \mathcal{E} , and any proof $\alpha : S \Longrightarrow_h T$ as an arrow $\llbracket \alpha \rrbracket_\theta : \llbracket S \rrbracket_\theta \rightarrow \llbracket T \rrbracket_\theta$ such that $q(\llbracket \alpha \rrbracket_\theta) = h$.

Interpretation of formulas. Defined by induction on the formula:

$$\llbracket X \rrbracket_\theta = \theta(X) \qquad \llbracket f^+ S \rrbracket_\theta = f^+ \llbracket S \rrbracket_\theta \qquad \llbracket g^- T \rrbracket_\theta = g^- \llbracket T \rrbracket_\theta$$

That is, we interpret the pushforward and pullback connectives using pushforward and pullback in \mathcal{E} , and atomic formulas using the action of the functor $\theta : \mathcal{D} \rightarrow \mathcal{E}$ on objects.

Interpretation of terms. By gentle abuse of notation, for any arrow $\alpha : X \rightarrow Y \in \mathcal{E}$ such that $q(\alpha) = fg$, let us write $f \setminus_g \alpha$ for the unique $\beta : f^+ X \rightarrow Y$ such that $\alpha = f_X \beta$, and $\alpha \setminus_f \bar{g}$ for the unique $\beta : X \rightarrow g^- Y$ such that $\alpha = \beta \bar{g}_Y$, both of these derived using the respective universal properties (0.1) and (0.2) of cartesian liftings. We thus reuse the notation for division to denote arrows in an arbitrary bifibration. Let us also allow ourselves to write f for f_X and \bar{g} for \bar{g}_Y , when the object can be deduced from the context.

The interpretation of terms is then defined by induction as follows:

$$\begin{aligned}
 \llbracket \delta \rrbracket_\theta &= \theta(\delta) \\
 \llbracket \alpha \cdot f \rrbracket_\theta &= \llbracket \alpha \rrbracket_\theta \cdot f & \llbracket f \setminus_g \alpha \rrbracket_\theta &= f \setminus_g \llbracket \alpha \rrbracket_\theta \\
 \llbracket \bar{g} \cdot \alpha \rrbracket_\theta &= \bar{g} \cdot \llbracket \alpha \rrbracket_\theta & \llbracket \alpha \setminus_f \bar{g} \rrbracket_\theta &= \llbracket \alpha \rrbracket_\theta \setminus_f \bar{g}
 \end{aligned}$$

Initial axioms are interpreted using the action of θ on arrows of \mathcal{D} , multiplications (corresponding to the inference rules $\mathsf{L}g^-$ and $\mathsf{R}f^+$) are interpreted by pre- and post-composition with cartesian liftings in \mathcal{E} (we remind the reader that we write composition by juxtaposition in diagrammatic order), and divisions ($\mathsf{L}f^+$ and $\mathsf{R}g^-$) are interpreted using division in \mathcal{E} , i.e., using the universal properties of the cartesian liftings.

1.5. Permutation equivalence. Let us begin by observing that some algebraic laws relating multiplication and division hold in any bifibration $q : \mathcal{E} \rightarrow \mathcal{B}$.

Proposition 1.3. For $\alpha : X \rightarrow Y \in \mathcal{E}$ such that $q(\alpha) = fg$, and $\beta : X \rightarrow Y'$ such that $q(\beta) = fgg'$, the following equalities hold:

$$\begin{aligned} \overline{f'} (\alpha_{f/\overline{g}}) &= (\overline{f'} \alpha)_{f'f/\overline{g}} \\ (f \setminus_g \alpha) g' &= f \setminus_{gg'} (\alpha g') \\ (f \setminus_{gg'} \beta)_{g/\overline{g'}} &= f \setminus_g (\beta_{fg/\overline{g'}}) \end{aligned}$$

Proof. The uniqueness condition on the arrows obtained by division can be expressed as an equivalence: $\beta = (\alpha_{f/\overline{g}}) \iff \beta \overline{g} = \alpha$ and $\beta = (f \setminus_g \alpha) \iff f \beta = \alpha$. The proofs for the desired equalities are then as follows:

$$\begin{aligned} \iff \frac{f'}{\overline{f'}} (\alpha_{f/\overline{g}}) \overline{g} &= f' \alpha & \iff (f \setminus_g \alpha) g' &= f \setminus_{gg'} (\alpha g') \\ \iff (\alpha_{f/\overline{g}}) \overline{g} &= \alpha & \iff f (f \setminus_g \alpha) &= \alpha \\ \iff \alpha_{f/\overline{g}} &= \alpha_{f/\overline{g}} & \iff f \setminus_g \alpha &= f \setminus_g \alpha \end{aligned}$$

$$\begin{aligned} \iff (f \setminus_{gg'} \beta)_{g/\overline{g'}} &= f \setminus_g (\beta_{fg/\overline{g'}}) \\ \iff f (f \setminus_{gg'} \beta)_{g/\overline{g'}} &= \beta_{fg/\overline{g'}} \\ \iff f ((f \setminus_{gg'} \beta)_{g/\overline{g'}}) \overline{g'} &= \beta \\ \iff ((f \setminus_{gg'} \beta)_{g/\overline{g'}}) \overline{g'} &= f \setminus_{gg'} \beta \\ \iff ((f \setminus_{gg'} \beta)_{g/\overline{g'}}) &= (f \setminus_{gg'} \beta)_{g/\overline{g'}} \end{aligned}$$

□

These laws imply that different derivations of the sequent calculus $\alpha_1, \alpha_2 : S \implies_h T$ will be interpreted by the same arrow $\llbracket \alpha_1 \rrbracket_\theta = \llbracket \alpha_2 \rrbracket_\theta : \llbracket S \rrbracket_\theta \rightarrow \llbracket T \rrbracket_\theta$ in \mathcal{E} under the bifibrational interpretation described in the previous section. In order for derivations to form a bifibration themselves, we must quotient them by a *permutation equivalence* relation (\sim) that relates any such pair of derivations equated in all interpretations.

Definition 1.4. We define **permutation equivalence** (\sim) as the smallest congruence that contains the four generating equations below. A congruence is an equivalence relation that is also closed under the derivation term-formers: if $\alpha_1 \sim \alpha_2$, then we also have $\alpha_1.g' \sim \alpha_2.g'$, and $f \setminus_g \alpha_1 \sim f \setminus_g \alpha_2$, etc., whenever both sides of the equation are valid terms. The generating equations of (\sim) are (whenever both sides of the equation are valid):

$$(\overline{f}. \alpha).h \sim \overline{f}.(\alpha.h) \quad \text{for } \alpha \text{ over } g \quad (1.2)$$

$$(f \setminus_g \alpha).h \sim f \setminus_{gh} (\alpha.h) \quad \text{for } \alpha \text{ over } fg \quad (1.3)$$

$$(\overline{f}. \alpha)_{fg/\overline{h}} \sim \overline{f}.(\alpha_{g/\overline{h}}) \quad \text{for } \alpha \text{ over } gh \quad (1.4)$$

$$(f \setminus_{gh} \alpha)_{g/\overline{h}} \sim f \setminus_g (\alpha_{fg/\overline{h}}) \quad \text{for } \alpha \text{ over } fgh \quad (1.5)$$

The generators of the permutation equivalence relation may be defined equivalently (albeit less compactly) as the following relations on derivations:

$$\begin{array}{cccc}
\frac{S \Longrightarrow T}{f^- S \xrightarrow{fg} h^+ T} \text{L}f^- & \frac{S \Longrightarrow T}{f^- S \xrightarrow{gh} h^+ T} \text{R}h^+ & \frac{S \Longrightarrow T}{f^+ S \xrightarrow{fg} h^+ T} \text{L}f^+ & \frac{S \Longrightarrow T}{f^+ S \xrightarrow{gh} h^+ T} \text{R}h^+ \\
\frac{S \xrightarrow{g} T}{f^- S \xrightarrow{fg} T} \text{L}f^- & \frac{S \xrightarrow{g} T}{S \xrightarrow{gh} h^+ T} \text{R}h^+ & \frac{S \xrightarrow{fg} T}{f^+ S \xrightarrow{g} T} \text{L}f^+ & \frac{S \xrightarrow{fg} T}{S \xrightarrow{gh} h^+ T} \text{R}h^+ \\
\frac{S \xrightarrow{gh} T}{f^- S \xrightarrow{fg} h^- T} \text{R}h^- & \frac{S \xrightarrow{gh} T}{f^- S \xrightarrow{g} h^- T} \text{L}f^- & \frac{S \xrightarrow{fgh} T}{f^+ S \xrightarrow{gh} T} \text{L}f^+ & \frac{S \xrightarrow{fgh} T}{f^+ S \xrightarrow{g} h^- T} \text{R}h^-
\end{array}$$

Observe that each of the relations permutes a left rule with a right rule.

Lemma 1.5. *Interpretation in any bifibration respects permutation equivalence: if $\alpha_1 \sim \alpha_2$, then $\llbracket \alpha_1 \rrbracket_\theta = \llbracket \alpha_2 \rrbracket_\theta$.*

Proof. The equivalence and congruence rules are obviously respected by the interpretation. It suffices to check that the four generating equations are also respected. The equation with two multiplications holds by associativity of composition:

$$\llbracket (\bar{f}'.\alpha).g' \rrbracket_\theta = \bar{f}'_{S'} \alpha g'_T = \llbracket \bar{f}'.(\alpha.g') \rrbracket_\theta$$

The three other equations exactly correspond to the equalities in Proposition 1.3, which hold in any bifibration. \square

Remark 1.6. The first permutation rule (1.2) is really an associativity property for left and right multiplication: a term outline $\bar{f}.\alpha.h$ can be read in two ways depending on the parenthesis placement, the two terms are valid (if α is) and are equivalent.

The second permutation rule (1.3) can be oriented from left to right: if the term $(f \setminus_g \alpha).h$ is valid, then one can “push the multiplication inside”, in the sense that the term $f \setminus_{gh} (\alpha.h)$ is also valid and is equivalent. But going in the other direction, the validity of a term of the form $f \setminus_k (\alpha.h)$ does not uniquely determine a corresponding term of the form $(f \setminus_l \alpha).h$. This depends on the solutions l of the equation $k = lh$: it may be that no such l exists, and there is no valid term of this shape, or that several l exist giving different possible rewrites. For example, we saw that for any $S \sqsubset A$ and $f : A \rightarrow B$ the identity derivation $\text{id}_{f^+ S} : f^+ S \xrightarrow{\text{id}_B} f^+ S$ is of the form $\text{id}_{f^+ S} = f \setminus_{\text{id}_B} (\text{id}_S.f)$, but there are not necessarily any derivations of the form $(f \setminus_g \text{id}_S).f$. Indeed, $(f \setminus_g \text{id}_S).f \sim \text{id}_{f^+ S}$ is an instance of permutation rule (1.3) just in case $g : B \rightarrow A$ exists and is inverse to f .

Similarly, the third permutation rule (1.4) can be oriented from right to left: one can always “push the multiplication inside” by moving from a term of the shape $\bar{f}.(\alpha / \bar{h})$ to an equivalent term of the shape $(\bar{f}.\alpha) / \bar{h}$, but the other direction is not always possible or is non-deterministic.

Finally, the fourth permutation rule (1.5) is also an associativity-like property that can be read in both directions: the two terms of the form $f \setminus \alpha / \bar{h}$ are equally valid (not always valid, but in the same situations) and equivalent. For example, let us assume that a term of the form $(f \setminus \alpha) / \bar{h}$ is valid and is over some morphism g ; then the term must be of the form $(f \setminus \alpha)_g / \bar{h}$, so $f \setminus \alpha$ must be over the morphism gh : our original term is uniquely determined

to be $(f \setminus_{gh} \alpha)_{g/\bar{h}}$, and α must be over fgh . By the same reasoning, any valid term of the form $f \setminus (\alpha / \bar{h})$ over g must be $f \setminus_g (\alpha_{fg/\bar{h}})$, this term is always valid (assuming that the other term is) and equivalent.

Remark 1.7. This situation, where multiplications can be pushed inside divisions but not necessarily outside, is reminiscent of Lambek calculus [34], which inspires our notation. Indeed, the generators of the permutation equivalence relation are analogous to the natural transformations

$$(A \bullet B) \bullet C \xrightarrow{\sim} A \bullet (B \bullet C) \quad A \setminus (B / C) \xrightarrow{\sim} (A \setminus B) / C \quad (1.6)$$

$$(A \setminus B) \bullet C \longrightarrow A \setminus (B \bullet C) \quad A \bullet (B / C) \longrightarrow (A \bullet B) / C \quad (1.7)$$

that exist in any monoidal biclosed category, where the transformations (1.6) are always invertible, but the transformations (1.7) in general are not. For example, considering the special case of cartesian closed categories, the former correspond to the natural isomorphisms $(A \times B) \times C \cong A \times (B \times C)$ and $(B^C)^A \cong (B^A)^C$, while the latter correspond to the natural transformations $B^A \times C \longrightarrow (B \times C)^A$ and $A \times B^C \longrightarrow (A \times B)^C$ that may be derived in any ccc, but are very rarely invertible.

1.6. Cut. From a logical perspective, cut is an expected rule for any logic, that ensures that sequences of reasoning can be chained as expected. The surprising insight behind Gentzen's invention of sequent calculus was that the cut rule is in a sense logically redundant. The original approach for demonstrating this was to include a cut rule directly in the inference system and then show *cut elimination*: any closed proof can be rewritten without using the cut rule. We use the alternative but nearly equivalent approach of first defining a cut-free system and then showing that the cut rule is admissible. Recall that a reasoning rule is *admissible* if, given closed derivations for its premises, there is a closed derivation for its conclusion.

Definition 1.8. The following **cut rule** is admissible

$$\frac{\alpha : S \xRightarrow{g} T \quad \beta : T \xRightarrow{h} U}{\alpha \cdot \beta : S \xRightarrow{gh} U}$$

via a composition function $\alpha \cdot \beta$ defined up to permutation equivalence (\sim).

Proof. The cut, or horizontal composition, written $\alpha \cdot \beta$, is defined by induction in the usual style of cut-elimination arguments: we reason by case analysis on the derivations α and β , and we may use instances of the cut rule on strictly smaller arguments (more precisely: in each case both arguments are at least as small and one of them is strictly smaller).

Principal cuts:

$$\frac{\frac{\delta : X \rightarrow Y \in \mathcal{D}}{\delta : X \xRightarrow{g} Y} \quad \frac{\epsilon : X \rightarrow Y \in \mathcal{D}}{\epsilon : Y \xRightarrow{h} Z}}{\delta \cdot \epsilon : X \xRightarrow{gh} Z} \quad := \quad \frac{\delta \epsilon : X \rightarrow Z \in \mathcal{D}}{\delta \epsilon : X \xRightarrow{gh} Z}$$

$$\frac{\frac{\alpha : S \Longrightarrow_g T}{\alpha.f : S \Longrightarrow_{gf} f^+ T} \quad \frac{\beta : T \Longrightarrow_{fh} U}{f \setminus_h \beta : f^+ T \Longrightarrow_h U}}{\alpha.f \cdot f \setminus_h \beta : S \Longrightarrow_{gfh} U} \quad := \quad \frac{\alpha : S \Longrightarrow_g T \quad \beta : T \Longrightarrow_{fh} U}{\alpha \cdot \beta : S \Longrightarrow_{gfh} U}$$

$$\frac{\frac{\alpha : S \Longrightarrow_{gf} T}{\alpha_g / \bar{f} : S \Longrightarrow_g f^- T} \quad \frac{\beta : T \Longrightarrow_h U}{\bar{f}.\beta : f^- T \Longrightarrow_{fh} U}}{\alpha_g / \bar{f} \cdot \bar{f}.\beta : S \Longrightarrow_{gfh} U} \quad := \quad \frac{\alpha : S \Longrightarrow_{gf} T \quad \beta : T \Longrightarrow_h U}{\alpha \cdot \beta : S \Longrightarrow_{gfh} U}$$

Commutative cuts:

$$\frac{\frac{\alpha : S \Longrightarrow_g T}{\bar{f}.\alpha : f^- S \Longrightarrow_{fg} T} \quad \beta : T \Longrightarrow_h U}{(\bar{f}.\alpha) \cdot \beta : f^- S \Longrightarrow_{fgh} U} \quad := \quad \frac{\alpha : S \Longrightarrow_g T \quad \beta : T \Longrightarrow_h U}{\alpha \cdot \beta : S \Longrightarrow_{gh} U} \quad \frac{\alpha \cdot \beta : S \Longrightarrow_{gh} U}{\bar{f}.\alpha \cdot \beta : f^- S \Longrightarrow_{fgh} U}$$

$$\frac{\frac{\alpha : S \Longrightarrow_{fg} T}{f \setminus_g \alpha : f^+ S \Longrightarrow_g T} \quad \beta : T \Longrightarrow_h U}{(f \setminus_g \alpha) \cdot \beta : f^+ S \Longrightarrow_{gh} U} \quad := \quad \frac{\alpha : S \Longrightarrow_{fg} T \quad \beta : T \Longrightarrow_h U}{\alpha \cdot \beta : S \Longrightarrow_{fgh} U} \quad \frac{\alpha \cdot \beta : S \Longrightarrow_{fgh} U}{f \setminus_{gh} (\alpha \cdot \beta) : f^+ S \Longrightarrow_{gh} U}$$

$$\frac{\alpha : S \Longrightarrow_g T \quad \frac{\beta : T \Longrightarrow_h U}{\beta.f : T \Longrightarrow_{hf} f^+ U}}{\alpha \cdot \beta.f : S \Longrightarrow_{ghf} f^+ U} \quad := \quad \frac{\alpha : S \Longrightarrow_g T \quad \beta : T \Longrightarrow_h U}{\alpha \cdot \beta : S \Longrightarrow_{gh} U} \quad \frac{\alpha \cdot \beta : S \Longrightarrow_{gh} U}{(\alpha \cdot \beta).f : S \Longrightarrow_{ghf} f^+ U}$$

$$\frac{\alpha : S \Longrightarrow_g T \quad \frac{\beta : T \Longrightarrow_{hf} U}{\beta_h / \bar{f} : T \Longrightarrow_h U}}{\alpha \cdot \beta_h / \bar{f} : S \Longrightarrow_{gh} f^- U} \quad := \quad \frac{\alpha : S \Longrightarrow_g T \quad \beta : T \Longrightarrow_{hf} U}{\alpha \cdot \beta : S \Longrightarrow_{ghf} U} \quad \frac{\alpha \cdot \beta : S \Longrightarrow_{ghf} U}{(\alpha \cdot \beta)_{gh} / \bar{f} : S \Longrightarrow_{gh} f^- U}$$

These rules cover all possible cases for a pair of derivations α, β whose judgments are composable judgments—the codomain of α is the domain of β . In particular, notice that when α starts with a right rule, for example a multiplication $\alpha = \alpha'.f$, β must start with the opposite rule on the *same* arrow, here a division by f : $\beta = f \setminus_h \beta'$. This is the only possibility because we know that the two judgments must compose at a type of the form $f^+ T'$, and this formula admits only one left (respectively right) rule.

The definition exhibits some non-determinism on the raw syntax of derivations, as it may be the case that two commutative cut rules are applicable. We consider it well-defined as a function *up to* permutation equivalence (\sim): whenever several of the cases above would apply, all derivations that can be defined from the definition are permutation equivalent.

Note that this happens only in the cases where the derivation on the left starts with the left rule, and the derivation on the right starts with a right rule. There are four such critical pairs (shapes of derivations that can be composed in two different ways), corresponding to all four permutations of generators. Two pairs are symmetric to each other, so we only need to check three.

$$\begin{array}{ccc}
 & (\bar{f}.\alpha) \cdot (\beta.i) & \\
 \swarrow \text{def} & & \searrow \text{def} \\
 \bar{f}.\alpha \cdot (\beta.i) & & ((\bar{f}.\alpha) \cdot \beta).i \\
 \downarrow \text{def} & & \downarrow \text{def} \\
 \bar{f}.\alpha \cdot (\beta.i) & \sim & (\bar{f}.\alpha \cdot \beta).i
 \end{array}$$

$$\begin{array}{ccc}
 & (\bar{f}.\alpha) \cdot (\beta_{h/\bar{i}}) & \\
 \swarrow \text{def} & & \searrow \text{def} \\
 \bar{f}.\alpha \cdot (\beta_{h/\bar{i}}) & & ((\bar{f}.\alpha) \cdot \beta)_{fgh/\bar{i}} \\
 \downarrow \text{def} & & \downarrow \text{def} \\
 \bar{f}.\alpha \cdot (\beta_{gh/\bar{i}}) & \sim & (\bar{f}.\alpha \cdot \beta)_{fgh/\bar{i}}
 \end{array}$$

$$\begin{array}{ccc}
 & (f \setminus_g \alpha) \cdot (\beta_{h/\bar{i}}) & \\
 \swarrow \text{def} & & \searrow \text{def} \\
 f \setminus_{gh} (\alpha \cdot (\beta_{h/\bar{i}})) & & ((f \setminus_g \alpha) \cdot \beta)_{gh/\bar{i}} \\
 \downarrow \text{def} & & \downarrow \text{def} \\
 f \setminus_{gh} (\alpha \cdot \beta)_{fgh/\bar{i}} & \sim & (f \setminus_{ghi} (\alpha \cdot \beta))_{gh/\bar{i}}
 \end{array}$$

□

We will refer to Definition 1.8 as a “cut-elimination procedure”, even if technically speaking it is a proof of cut-admissibility.

Lemma 1.9. *Cut is compatible with permutation equivalence. If $\alpha_1 \sim \alpha_2$ then $\alpha_1 \cdot \beta \sim \alpha_2 \cdot \beta$, and $\beta \cdot \alpha_1 \sim \beta \cdot \alpha_2$,*

Proof sketch. The proof (detailed in Appendix A.1) proceeds by equational reasoning on proof terms. The following piece of equational reasoning is representative of the rest of the proof, and many other proofs in this section:

$$\begin{aligned}
& (\bar{f}.(\alpha.h)) \cdot \beta \\
\sim & \bar{f}.((\alpha.h) \cdot \beta) \\
= & \bar{f}.((\alpha.h) \cdot (h \setminus_i \beta')) \\
= & \bar{f}.(\alpha \cdot \beta') \\
\sim & (\bar{f}.\alpha) \cdot \beta' \\
= & ((\bar{f}.\alpha).h) \cdot (h \setminus_i \beta') \\
= & ((\bar{f}.\alpha).h) \cdot \beta
\end{aligned}$$

Each reasoning step is by definition of cut-elimination. Principal cuts are defined unambiguously, for example we have $(\alpha.h) \cdot \beta = (\alpha.h) \cdot (h \setminus_i \beta')$ by definition. Commutative cuts (left rule in the left derivation, or right rule in the right derivation) are only defined up to permutation equivalence, for example we only have $\bar{f}.(\alpha \cdot \beta') \sim (\bar{f}.\alpha) \cdot \beta'$. \square

Lemma 1.10 (identity is neutral). *The identity derivations are neutral elements for cut:*

$$\text{id}_S \cdot \alpha \quad \sim \quad \alpha \quad \sim \quad \alpha \cdot \text{id}_S$$

Proof. See Appendix A.1. \square

Lemma 1.11 (η -expansion).

For any $\alpha : f^+ S \Longrightarrow_g T$, for any $\beta : S \Longrightarrow_f g^- T$, we have:

$$\alpha \sim f \setminus_g (f_S \cdot \alpha) \qquad \beta \sim (\beta \cdot \bar{g}_T) f / \bar{g}$$

Proof. See Appendix A.1. \square

Lemma 1.12. *Cut is associative: $(\alpha \cdot \beta) \cdot \gamma \sim \alpha \cdot (\beta \cdot \gamma)$.*

Proof. See Appendix A.1. \square

1.7. The free bifibration. We can now formally define $\mathcal{Bif}(p)$ as a category where:

- The objects are the valid formulas $S \sqsubset A$, (Section 1.1).
- The arrows between $S \sqsubset A$ and $T \sqsubset B$ are the valid (closed) derivations of the judgment $S \Longrightarrow_f T$ for any $f : A \rightarrow B \in \mathcal{C}$ (Section 1.2), quotiented by the permutation equivalence relation (\sim) (Section 1.5).
- The identity arrow for $S \sqsubset A$ is (the equivalence class of) the derivation id_S (Defn. 1.2).
- The composition of $\alpha : S \Longrightarrow_g T$ and $\beta : T \Longrightarrow_h U$ is the cut $\alpha \cdot \beta : S \Longrightarrow_{gh} U$ (Defn. 1.8).

Lemma 1.13. *$\mathcal{Bif}(p)$ is a category.*

Proof. We proved that composition is well-defined on derivations quotiented by permutation equivalence (Lemma 1.9), and that identity and cut satisfy neutrality (Lemma 1.10) and associativity (Lemma 1.12). \square

Next we show that $\mathcal{Bif}(p)$ is the domain of a bifibration $\Lambda_p : \mathcal{Bif}(p) \rightarrow \mathcal{C}$ equipped with a functor $\eta_p : \mathcal{D} \rightarrow \mathcal{Bif}(p)$ such that $\Lambda_p \circ \eta_p = p$. We define Λ_p by

- (1) $\Lambda_p(S \sqsubset A) := A$
- (2) $\Lambda_p(\alpha : S \Longrightarrow_f T) := f$

Lemma 1.14. *$\Lambda_p : \mathcal{Bif}(p) \rightarrow \mathcal{C}$ is a bifibration.*

Proof. For any arrow $f : A \rightarrow B$ in \mathcal{C} and formula $S \sqsubset A$, the pushforward of S along f is precisely the formula $f^+ S \sqsubset B$, and we have previously defined the arrow $f_S : S \rightarrow f^+ S$ as (the equivalence class of) $\text{id}_S \cdot f$. Let us show universality, i.e., that f_S is $+$ -cartesian.

For any $g : B \rightarrow C$ and $\alpha : S \Longrightarrow_{fg} T$, we can take $\beta := f \setminus_g \alpha : f^+ S \Longrightarrow_g T$. We have to show $\alpha \sim f_S \cdot \beta$:

$$\begin{aligned} & f_S \cdot \beta \\ &= (\text{id}_S \cdot f) \cdot (f \setminus_g \alpha) \\ &= \text{id}_S \cdot \alpha \\ &\sim \alpha \end{aligned}$$

as well as unicity: if $\alpha \sim f_S \cdot \beta'$ then

$$\begin{aligned} & \beta' \\ &\sim f \setminus_g (f_S \cdot \beta') \quad \text{by } \eta\text{-expansion (Lemma 1.11)} \\ &\sim f \setminus_g \alpha \\ &= \beta \end{aligned}$$

The case for $\bar{g}_T : g^- T \rightarrow T$ is symmetrical. \square

We define the functor $\eta_p : \mathcal{D} \rightarrow \mathcal{Bif}(p)$ by interpreting every object X of \mathcal{D} as an atomic formula and every arrow $\delta : X \rightarrow Y$ as an initial axiom, and it is immediate that $\Lambda_p \circ \eta_p = p$.

Lastly, we need to show that $\Lambda_p : \mathcal{Bif}(p) \rightarrow \mathcal{C}$ is universal in the sense explained before, that for any bifibration $q : \mathcal{E} \rightarrow \mathcal{C}$ equipped with a functor $\theta : \mathcal{D} \rightarrow \mathcal{E}$ such that $q \circ \theta = p$, there exists a unique morphism of bifibrations $\llbracket - \rrbracket_\theta : \Lambda_p \rightarrow q$ making the diagram below commute:

$$\begin{array}{ccc} \mathcal{D} & \xrightarrow{\eta_p} & \mathcal{Bif}(p) \\ & \searrow \theta & \downarrow \Lambda_p \\ & & \mathcal{C} \\ & \swarrow p & \uparrow q \text{ bifibration} \\ & & \mathcal{E} \end{array} \quad \begin{array}{c} \llbracket - \rrbracket_\theta \\ \swarrow \\ \mathcal{E} \end{array}$$

We already defined the interpretation $\llbracket - \rrbracket_\theta$ of formulas and proof terms in Section 1.4, and showed that it respects permutation equivalence (Lemma 1.5). We now establish that it is a functor $\llbracket - \rrbracket_\theta : \mathcal{Bif}(p) \rightarrow \mathcal{E}$, and that it is the unique morphism of bifibrations making the diagram above commute.

Lemma 1.15. $\llbracket - \rrbracket_\theta : \mathcal{Bif}(p) \rightarrow \mathcal{E}$ is a functor.

Proof. We prove the interpretation respects identities $\llbracket \text{id}_S \rrbracket_\theta = \text{id}_{\llbracket S \rrbracket_\theta}$ by induction on S :

$$\begin{aligned} & \llbracket \text{id}_X \rrbracket_\theta = \theta(\text{id}_X) = \text{id}_{\theta(X)} = \text{id}_{\llbracket X \rrbracket_\theta} \\ &= \llbracket \text{id}_{f^+ S} \rrbracket_\theta \\ &= \llbracket f \setminus_{\text{id}_B} (\text{id}_S \cdot f) \rrbracket_\theta \\ &= f \setminus_{\text{id}_B} \left(\llbracket \text{id}_S \rrbracket_\theta \quad f \llbracket S \rrbracket_\theta \right) \\ &= f \setminus_{\text{id}_B} \left(\text{id}_{\llbracket S \rrbracket_\theta} \quad f \llbracket S \rrbracket_\theta \right) \quad \text{by induction on } S \\ &= f \setminus_{\text{id}_B} f \llbracket S \rrbracket_\theta \\ &= \text{id}_{\llbracket S \rrbracket_\theta} \quad \text{property of cartesian liftings} \end{aligned}$$

(The case of the identity on $g^- T$ is symmetrical.)

We prove that the interpretation respects composition $\llbracket \alpha \cdot \beta \rrbracket_\theta = \llbracket \alpha \rrbracket_\theta \llbracket \beta \rrbracket_\theta$ by simultaneous induction on the derivations α and β . Representative cases follow:

$$\begin{aligned}
\llbracket \delta \cdot \epsilon \rrbracket_\theta &= \llbracket \delta \epsilon \rrbracket_\theta = \theta(\delta \epsilon) = \llbracket \delta \rrbracket_\theta \llbracket \epsilon \rrbracket_\theta & & \begin{aligned} &= \llbracket \alpha.f \rrbracket_\theta \llbracket f \setminus_h \beta \rrbracket_\theta \\ &= \llbracket \alpha \rrbracket_\theta f_{\llbracket S \rrbracket_\theta} (f \setminus_h \llbracket \beta \rrbracket_\theta) \\ &= \llbracket \alpha \rrbracket_\theta \llbracket \beta \rrbracket_\theta \\ &= \llbracket \alpha \cdot \beta \rrbracket_\theta \\ &= \llbracket (\alpha.f) \cdot (f \setminus_h \beta) \rrbracket_\theta \end{aligned} & & \text{by (0.1)} \\
\llbracket \alpha \cdot (\beta.f) \rrbracket_\theta &= \llbracket (\alpha \cdot \beta).f \rrbracket_\theta & & \begin{aligned} &= \llbracket \alpha \cdot (\beta \setminus_h \bar{f}) \rrbracket_\theta \\ &= \llbracket (\alpha \cdot \beta) \setminus_{gh} \bar{f} \rrbracket_\theta \\ &= \llbracket \alpha \cdot \beta \rrbracket_\theta \setminus_{gh} \bar{f} \\ &= (\llbracket \alpha \rrbracket_\theta \llbracket \beta \rrbracket_\theta) \setminus_{gh} \bar{f} \\ &= \llbracket \alpha \rrbracket_\theta (\llbracket \beta \rrbracket_\theta \setminus_h \bar{f}) \\ &= \llbracket \alpha \rrbracket_\theta \llbracket \beta \setminus_h \bar{f} \rrbracket_\theta \end{aligned} & & \text{by Proposition 1.3} \\
\llbracket \alpha \cdot \beta \rrbracket_\theta f_{\llbracket U \rrbracket_\theta} &= \llbracket \alpha \rrbracket_\theta \llbracket \beta \rrbracket_\theta f_{\llbracket U \rrbracket_\theta} & & \\
\llbracket \alpha \rrbracket_\theta \llbracket \beta.f \rrbracket_\theta &= \llbracket \alpha \rrbracket_\theta \llbracket \beta \rrbracket_\theta \llbracket f \rrbracket_\theta & &
\end{aligned}$$

□

Lemma 1.16. $\llbracket - \rrbracket_\theta : \Lambda_p \rightarrow q$ is the unique morphism of bifibrations such that $\llbracket - \rrbracket_\theta \circ \eta_p = \theta$.

Proof. Recall that the interpretation was defined as follows on terms:

$$\begin{aligned}
\llbracket \delta \rrbracket_\theta &= \theta(\delta) \\
\llbracket \alpha.f \rrbracket_\theta &= \llbracket \alpha \rrbracket_\theta f & \llbracket f \setminus_g \alpha \rrbracket_\theta &= f \setminus_g \llbracket \alpha \rrbracket_\theta \\
\llbracket \bar{g}.\alpha \rrbracket_\theta &= \bar{g} \llbracket \alpha \rrbracket_\theta & \llbracket \alpha \setminus_f \bar{g} \rrbracket_\theta &= \llbracket \alpha \rrbracket_\theta \setminus_f \bar{g}
\end{aligned}$$

The equation $\llbracket - \rrbracket_\theta \circ \eta_p = \theta$ holds immediately by construction.

To be a morphism of bifibrations, the interpretation must send cartesian arrows to cartesian arrows of the same type. This follows from the definition and the previous lemma. For $+$ -cartesian arrows for example, we have:

$$\llbracket f_S \rrbracket_\theta = \llbracket \text{id}_S.f \rrbracket_\theta = \llbracket \text{id}_S \rrbracket_\theta f_{\llbracket S \rrbracket_\theta} = \text{id}_{\llbracket S \rrbracket_\theta} f_{\llbracket S \rrbracket_\theta} = f_{\llbracket S \rrbracket_\theta}$$

Finally, for unicity we remark that each of the five equations defining the interpretation $\llbracket - \rrbracket_\theta$ must hold for any morphism of bifibrations $\tau : \Lambda_p \rightarrow q$ such that $\tau \circ \eta_p = \theta$, so necessarily $\tau = \llbracket - \rrbracket_\theta$. □

Putting it all together, we conclude with the main theorem of this section.

Theorem 1.17. $\Lambda_p : \text{Bif}(p) \rightarrow \mathcal{C}$ is the free bifibration on $p : \mathcal{D} \rightarrow \mathcal{C}$.

An immediate consequence—applying one of the equivalent characterizations of adjunctions [38, Theorem IV.1.2(ii), p.83]—is that we can construct a functor

$$\Lambda : \text{Cat}/\mathcal{C} \longrightarrow \text{BifCat}(\mathcal{C})$$

from the slice of the category of categories over \mathcal{C} into the category of cloven bifibrations over \mathcal{C} , which is left adjoint to the evident forgetful functor:

$$\begin{array}{ccc} & \Lambda & \\ & \curvearrowright & \\ \text{Cat}/\mathcal{C} & \perp & \text{BifCat}(\mathcal{C}) \\ & \curvearrowleft & \\ & U & \end{array} \quad (1.8)$$

This induces a monad on Cat/\mathcal{C} whose algebras are the cloven bifibrations over \mathcal{C} .

1.8. Some proof-theoretic consequences. Initiality of the $\text{Bif}(p)$ construction has as another direct consequence that the sequent calculus can be used to reason about both existence and equality of morphisms in arbitrary bifibrations. We state this as a corollary of Theorem 1.17.

Corollary 1.18. *Suppose given a functor $p : \mathcal{D} \rightarrow \mathcal{C}$, an arrow $f : A \rightarrow B$ of \mathcal{C} , and a pair of bifibrational formulas $S \sqsubset A$ and $T \sqsubset B$. Then:*

- (1) *the sequent $S \Longrightarrow_f T$ is derivable iff for every bifibration $q : \mathcal{E} \rightarrow \mathcal{C}$ and functor $\theta : \mathcal{D} \rightarrow \mathcal{E}$ such that $q \circ \theta = p$, there is a morphism $\alpha : \llbracket S \rrbracket_\theta \rightarrow \llbracket T \rrbracket_\theta$ in \mathcal{E} such that $q(\alpha) = f$;*
- (2) *two derivations $\alpha_1, \alpha_2 : S \Longrightarrow_f T$ are permutation equivalent $\alpha_1 \sim \alpha_2$ iff for every bifibration $q : \mathcal{E} \rightarrow \mathcal{C}$ and functor $\theta : \mathcal{D} \rightarrow \mathcal{E}$ such that $q \circ \theta = p$, they have the same interpretation $\llbracket \alpha_1 \rrbracket_\theta = \llbracket \alpha_2 \rrbracket_\theta$ as morphisms in \mathcal{E} .*

As recalled in the Introduction, any bifibration $\mathcal{E} \rightarrow \mathcal{C}$ determines and is determined by the data of a pseudofunctor $\mathcal{C} \rightarrow \text{Adj}$ into the category of categories and adjunctions. This pseudofunctor sends every object A of \mathcal{C} to its fiber category \mathcal{E}_A consisting of the arrows lying over id_A , and every arrow $f : A \rightarrow B$ of \mathcal{C} to the push/pull adjunction $f^+ : \mathcal{E}_A \rightarrow \mathcal{E}_B \dashv f^- : \mathcal{E}_B \rightarrow \mathcal{E}_A$ between fiber categories. In sequent calculus terms, the action of the push/pull functors on arrows is described by the following open derivations:

$$\begin{array}{ccc} \frac{S_1 \Longrightarrow_{\text{id}_A} S_2}{S_1 \Longrightarrow_f f^+ S_2} \text{R}f^+ & \frac{T_1 \Longrightarrow_{\text{id}_B} T_2}{g^- T_1 \Longrightarrow_g T_2} \text{L}g^- \\ \frac{f^+ S_1 \Longrightarrow_{\text{id}_B} f^+ S_2}{f^+ S_1 \Longrightarrow_{\text{id}_A} f^+ S_2} \text{L}f^+ & \frac{g^- T_1 \Longrightarrow_{\text{id}_A} g^- T_2}{g^- T_1 \Longrightarrow_{\text{id}_B} g^- T_2} \text{R}g^- \end{array} \quad (1.9)$$

The fact that the fiber functor $\mathcal{C} \rightarrow \text{Adj}$ associated to a bifibration $\mathcal{E} \rightarrow \mathcal{C}$ is only a *pseudo*-functor in general means that the identities $(g \circ f)^+ \equiv g^+ \circ f^+$ and $(g \circ f)^- \equiv f^- \circ g^-$ hold only up to natural isomorphism, as do the identities $\text{id}_A^+ \equiv \text{id}_{\mathcal{E}_A} \equiv \text{id}_A^-$. To describe this in the sequent calculus, let us first define *logical equivalence* of formulas (corresponding to what is often called “vertical isomorphism” in the categorical literature on fibrations).

Definition 1.19. Let $S_1, S_2 \sqsubset A$ be two bifibrational formulas over the same object A . We say that S_1 and S_2 are **logically equivalent**, written $S_1 \equiv_A S_2$ or simply $S_1 \equiv S_2$, if there are a pair of derivations $\alpha : S_1 \Longrightarrow_{\text{id}_A} S_2$ and $\beta : S_2 \Longrightarrow_{\text{id}_A} S_1$ such that $\alpha \cdot \beta \sim \text{id}_{S_1}$ and $\beta \cdot \alpha \sim \text{id}_{S_2}$.

Proposition 1.20. *Two formulas $S_1, S_2 \sqsubset A$ are logically equivalent $S_1 \equiv_A S_2$ iff they are isomorphic in the fiber category $\text{Bif}(p)_A$.*

Logical equivalence should be distinguished from the weaker notion of isomorphism $S \cong T$ given by isomorphism in the total category $\mathcal{Bif}(p)$, i.e., a pair of derivations $\alpha : S \Longrightarrow_f T$ and $\beta : T \Longrightarrow_g S$ such that $\alpha \cdot \beta \sim \text{id}_S$ and $\beta \cdot \alpha \sim \text{id}_T$ (where one does not require f and g to be identity arrows or that $S \sqsubset A$ and $T \sqsubset B$ lie over the same object).

Proposition 1.21 (Pseudofunctoriality). *The following logical equivalences hold*

$$(g \circ f)^+ S \equiv g^+ f^+ S \quad (1.10)$$

$$(f g)^- T \equiv f^- g^- T \quad (1.11)$$

$$\text{id}^+ U \equiv U \equiv \text{id}^- U \quad (1.12)$$

for all f, g, S, T, U of the appropriate type.

Proof. By Proposition 1.20, this is a corollary of Lemma 1.14, since these isomorphisms hold in the fiber categories of any bifibration. Working out the derivation explicitly, the equivalence (1.10), for instance, is witnessed by the following derivations (parameterized by $S \sqsubset A$, $f : A \rightarrow B$, and $g : B \rightarrow C$), whose compositions are permutation equivalent to the identity derivations on $(g \circ f)^+ S$ and $g^+ f^+ S$:

$$\frac{\frac{\frac{S \Longrightarrow S}{\text{id}_A}}{S \Longrightarrow f^+ S}}{S \Longrightarrow g^+ f^+ S}}{(g \circ f)^+ S \Longrightarrow g^+ f^+ S} \quad \frac{\frac{\frac{S \Longrightarrow S}{\text{id}_A}}{S \Longrightarrow (g \circ f)^+ S}}{f^+ S \Longrightarrow (g \circ f)^+ S}}{g^+ f^+ S \Longrightarrow (g \circ f)^+ S}$$

□

More generally, by Theorem 1.17 we have

Corollary 1.22. *Given a functor $p : \mathcal{D} \rightarrow \mathcal{C}$, two bifibrational formulas $S_1, S_2 \sqsubset A$ are logically equivalent $S_1 \equiv S_2$ iff their interpretations are isomorphic $\llbracket S_1 \rrbracket_\theta \cong \llbracket S_2 \rrbracket_\theta$ in every fiber category \mathcal{E}_A for every bifibration $q : \mathcal{E} \rightarrow \mathcal{C}$ and functor $\theta : \mathcal{D} \rightarrow \mathcal{E}$ such that $q \circ \theta = p$.*

Finally, let us observe that the free bifibration construction is *conservative*.

Proposition 1.23 (Conservativity). *The functor $\eta_p : \mathcal{D} \rightarrow \mathcal{Bif}(p)$ is full and faithful.*

Conservativity is obvious by inspection of the sequent calculus, since the only way to prove an atomic sequent $X \Longrightarrow_f Y$ is by providing an arrow $\delta : X \rightarrow Y$ in \mathcal{D} such that $p(\delta) = f$. Faithfulness of η_p is also a corollary of the following stronger statement, whose proof only uses the universal property of the free bifibration, and which will serve us in the next section.

Proposition 1.24. *The functor $\eta_p : \mathcal{D} \rightarrow \mathcal{Bif}(p)$ is equipped with a left inverse $\chi_p : \mathcal{Bif}(p) \rightarrow \mathcal{D}$, such that the pairing $(\chi_p, \Lambda_p) : \mathcal{Bif}(p) \rightarrow \mathcal{D} \times \mathcal{C}$ is a morphism of bifibrations from Λ_p to the trivial bifibration $\pi_2 : \mathcal{D} \times \mathcal{C} \rightarrow \mathcal{C}$.*

Proof. Projection functors $\pi_2 : \mathcal{D} \times \mathcal{C} \rightarrow \mathcal{C}$ are always bifibrations, with cartesian liftings defined by

$$(X, A) \xrightarrow{(\text{id}_X, f)} (X, B) =: f^+(X, A) \quad g^-(X, C) := (X, B) \xrightarrow{(\text{id}_X, g)} (X, C)$$

$$A \xrightarrow{f} B \quad B \xrightarrow{g} C$$

By the universal property of the free bifibration, we therefore obtain a unique morphism of bifibrations $\llbracket - \rrbracket_\theta : \Lambda_p \rightarrow \pi_2$

$$\begin{array}{ccc} \mathcal{D} & \xrightarrow{\eta_p} & \text{Bif}(p) \\ & \searrow \theta & \downarrow \Lambda_p \\ & & \mathcal{D} \times \mathcal{C} \\ & \swarrow p & \nwarrow \pi_2 \\ & & \mathcal{C} \end{array}$$

where $\theta := (\text{id}_{\mathcal{D}}, p)$. This proves the proposition, taking $\chi_p := \pi_1 \circ \llbracket - \rrbracket_\theta$ and using the universal property of the product of categories. \square

In terms of our concrete description of the free bifibration, $\chi_p : \text{Bif}(p) \rightarrow \mathcal{D}$ is the functor sending a bifibrational formula to its atomic source and sending an equivalence class of derivations to the necessarily unique initial axiom with which all of the derivations begin.

1.9. Related work. There is a well-known construction, originally due to Gray [19], of the free fibration on a functor (or dually the free opfibration) as a comma category. The standard construction actually defines a *split* (op)fibration, in the sense that the pseudofunctoriality laws (1.10)–(1.12) collapse to strict equalities, although one can represent non-split fibrations and opfibrations as pseudo-algebras for the corresponding pseudomonads (see the account by Anders Kock [28], as well as recent work of Emmenegger, Mesiti, Rosolini, and Streicher on an alternative construction of fibrations [17]). Conversely, our construction produces a non-split bifibration, but one can recover a split bifibration by restricting to *nonidentity strictly alternating formulas* (see Section 3.3 below, and in particular Theorem 3.5). It is worth mentioning that since a functor is a bifibration just in case it is both a fibration and an opfibration, the free bifibration monad is formally the coproduct of the free fibration and the free opfibration monads—but this characterization does not give much help in constructing it explicitly.

We will discuss Lamarche’s construction of the free bifibration on a functor along with Dawson, Paré, and Pronk’s closely related $\Pi_2(\mathcal{C})$ construction, both already referenced in the Introduction, further in Section 2.

The notation $S \sqsubset A$ to denote an object $S \in \text{Bif}(p)$ such that $\Lambda_p(S) = A$ is taken from Mellès and Zeilberger’s work on functors as type refinement systems [40], and more broadly our proof-theoretic construction is inspired in part from their suggestive type-theoretic reading of bifibrations [41, 42], as well as from the bifibrational calculus of Licata, Shulman, and Riley [37]. This approach of course owes much to the tradition of categorical proof theory as pioneered by Lambek [35].

2. THE DOUBLE CATEGORY OF ZIGZAGS

As mentioned in the Introduction, Lamarche [32, 33] defined a *path category* $\mathcal{P}(\mathcal{C})$ for each category \mathcal{C} , showed that it admits the structure of a double category, and explained how to use this structure to construct the free bifibration on a functor. In this section we will start the other way around, defining a category $\mathcal{Z}(\mathcal{C})$ equivalent to Lamarche's path category as the free bifibration on the identity functor, and showing that it admits a strict double category structure. The objects and tight morphisms of $\mathcal{Z}(\mathcal{C})$ are the objects and arrows in \mathcal{C} , while the loose morphisms are zigzags (or paths) of arrows in \mathcal{C} , and thus $\mathcal{Z}(\mathcal{C})$ is called the *zigzag double category*. Its underlying 2-category of loose morphisms is equivalent to $\Pi_2(\mathcal{C})$, the 2-category studied by Dawson, Paré, and Pronk [14] that is obtained by freely adjoining right adjoints to \mathcal{C} . To prove this equivalence, we establish that $\mathcal{Z}(\mathcal{C})$ has another characterization as the free *fibrant double category* on \mathcal{C} .

Whereas Lamarche and Dawson, Paré, and Pronk gave rather intricate combinatorial descriptions of what it means to be a morphism of zigzags (i.e., to be a double cell in $\mathcal{Z}(\mathcal{C})$, or a 2-cell in $\Pi_2(\mathcal{C})$), we will extract a simple presentation of $\mathcal{Z}(\mathcal{C})$ by translating the sequent calculus definition to double categorical syntax. In turn we will see how that presentation can be translated backwards to obtain a topologically meaningful string diagrammatic representation of proofs in the bifibrational calculus.

2.1. Deriving a double category structure from a right action. For the rest of the paper we define $\mathcal{Z}(\mathcal{C}) := \mathcal{Bif}(\text{id}_{\mathcal{C}})$ to be the total category of the free bifibration on the identity functor for \mathcal{C} . Observe that by Proposition 1.24, $\mathcal{Z}(\mathcal{C})$ is automatically equipped with the structure of a reflexive graph object in Cat

$$\begin{array}{ccc} & \mathcal{Z}(\mathcal{C}) & \\ \text{src} \downarrow & \uparrow U & \downarrow \text{tgt} \\ & \mathcal{C} & \end{array}$$

taking $\text{src} := \chi_p$, $U := \eta_p$, and $\text{tgt} := \Lambda_p$, for $p = \text{id}_{\mathcal{C}}$. We use the letter Z to range over the objects of $\mathcal{Z}(\mathcal{C})$, which we refer to as *zigzag formulas* or simply *zigzags*, and write $Z : A \rightsquigarrow B$ to indicate that $\text{src}(Z) = A$ and $\text{tgt}(Z) = B$. We use the letter ζ to range over arrows $Z \rightarrow Z'$ of $\mathcal{Z}(\mathcal{C})$. Such an arrow may be depicted as a *cell*

$$\begin{array}{ccc} A & \xrightarrow{f} & B \\ \zeta \downarrow & & \downarrow \zeta' \\ C & \xrightarrow{g} & D \end{array}$$

where $f = \text{src}(\zeta)$ and $g = \text{tgt}(\zeta)$. We always depict the tight morphisms (corresponding to the underlying arrows of \mathcal{C}) horizontally and the loose morphisms (corresponding to objects of $\mathcal{Z}(\mathcal{C})$, i.e., zigzags of arrows in \mathcal{C}) vertically. We will develop this diagrammatic picture further in Section 2.2, but for now we work with the purely formal definition building on the results of Section 1.

To extend this reflexive graph structure on $\mathcal{Z}(\mathcal{C})$ to a double category (i.e., to an internal category in Cat), we need to exhibit a vertical composition functor

$$\odot : \mathcal{Z}(\mathcal{C})_{\text{tgt} \times \text{src}} \mathcal{Z}(\mathcal{C}) \rightarrow \mathcal{Z}(\mathcal{C})$$

and prove that it is associative and neutral relative to U . We will recover vertical composition \odot as the special case $p = \text{id}_{\mathcal{C}}$ of a more general action

$$\otimes_p : \mathcal{Bif}(p)_{\Lambda_p \times_{\text{src}}} \mathcal{Z}(\mathcal{C}) \rightarrow \mathcal{Bif}(p)$$

of $\mathcal{Z}(\mathcal{C})$ on any free bifibration $\mathcal{Bif}(p)$.

To define this action, let us begin by observing that the sequent calculus for $\mathcal{Z}(\mathcal{C})$ shares the same logical inference rules as any free bifibration $\mathcal{Bif}(p)$ over the same base category: their sequent calculi only differ in the choice of atomic formulas and initial axioms. Moreover, both atomic formulas and initial axioms of $\mathcal{Z}(\mathcal{C})$ carry no information, since they are uniquely determined by the underlying object/arrow of the base category. These observations make clear that we can define the action of $\mathcal{Z}(\mathcal{C})$ on $\mathcal{Bif}(p)$ as follows:

- at the level of formulas, given a formula $S \sqsubset A$ and a zigzag $Z : A \rightsquigarrow B$, the formula $S \otimes_p Z \sqsubset B$ is obtained by substituting S for the unique atom A in Z ;
- at the level of proofs (i.e., representatives of arrows), given $\alpha \in \mathcal{Bif}(p)$ such that $\Lambda_p(\alpha) = f$ and $\zeta \in \mathcal{Z}(\mathcal{C})$ such that $\text{src}(\zeta) = f$ and $\text{tgt}(\zeta) = g$, the proof $\alpha \otimes_p \zeta \in \mathcal{Bif}(p)$ such that $\Lambda_p(\alpha \otimes_p \zeta) = g$ is obtained by substituting α for the unique initial axiom f in ζ .

As a visual aid, let us write “ \square ” for the uniquely determined atoms and axioms in a zigzag formula/proof. Formally, the action can then be expressed inductively as follows:

$$\begin{aligned} S \otimes_p \square &= S & S \otimes_p (f^+ Z) &= f^+(S \otimes_p Z) & S \otimes_p (g^- Z) &= g^-(S \otimes_p Z) \\ \alpha \otimes_p \square &= \alpha & \alpha \otimes_p (\zeta.f) &= (\alpha \otimes_p \zeta).f & \alpha \otimes_p (f \setminus_g \zeta) &= f \setminus_g (\alpha \otimes_p \zeta) \\ \alpha \otimes_p (\bar{g}.\zeta) &= \bar{g}.\alpha \otimes_p \zeta & \alpha \otimes_p (\zeta_f / \bar{g}) &= (\alpha \otimes_p \zeta)_f / \bar{g} \end{aligned}$$

Proposition 2.1. *The action extends to a functor $\otimes_p : \mathcal{Bif}(p)_{\Lambda_p \times_{\text{src}}} \mathcal{Z}(\mathcal{C}) \rightarrow \mathcal{Bif}(p)$.*

Proof. Since \otimes_p is defined as a substitution operator, it is immediate that it respects permutation equivalence of proofs in the sense that we can define $[\alpha] \otimes_p [\zeta] := [\alpha \otimes_p \zeta]$ without ambiguity. To show that it is functorial, we need to verify the permutation equivalence

$$(\alpha \cdot \alpha') \otimes_p (\zeta \cdot \zeta') \sim (\alpha \otimes_p \zeta) \cdot (\alpha' \otimes_p \zeta') \quad (2.1)$$

for all $\alpha, \alpha' \in \mathcal{Bif}(p)$ and $\zeta, \zeta' \in \mathcal{Z}(\mathcal{C})$ for which these operations are defined. This follows from the definition of cut (Definition 1.8) by a straightforward induction on ζ and ζ' . \square

We can now define the vertical composition functor $\odot : \mathcal{Z}(\mathcal{C})_{\text{tgt} \times_{\text{src}}} \mathcal{Z}(\mathcal{C}) \rightarrow \mathcal{Z}(\mathcal{C})$ as the action of $\mathcal{Z}(\mathcal{C})$ on itself, $\odot := \otimes_{\text{id}_{\mathcal{C}}}$, and use that to define the double category of zigzags.

Proposition 2.2. *Vertical composition \odot is strictly associative and neutral relative to U .*

Proof. Immediate from the definition of \otimes_p as a substitution operator. \square

Definition 2.3. The **zigzag double category** of \mathcal{C} , notated $\mathbb{Z}(\mathcal{C})$, is defined as a strict double category by the data $\mathbb{Z}(\mathcal{C}) := (\mathcal{C}, \mathcal{Z}(\mathcal{C}), \text{src}, \text{tgt}, U, \odot)$, where

$$\mathcal{Z}(\mathcal{C}) := \mathcal{Bif}(p) \quad \text{src} := \chi_p \quad \text{tgt} := \Lambda_p \quad U := \eta_p \quad \odot = \otimes_p$$

as explained above, for $p = \text{id}_{\mathcal{C}}$.

Notice that with this definition, equation (2.1) specializes to the interchange axiom for the double category $\mathbb{Z}(\mathcal{C})$. We already saw in Section 1.5 that the permutation equivalence relations were necessary for obtaining a bifibration, but it is worth mentioning that this axiom also forces the definition of permutation equivalence. For example, the permutation relation (1.3) may be rederived as follows:

$$\begin{aligned}
(f \setminus_g \alpha).h &= (\alpha \circledast (f \setminus_g fg)) \cdot (\text{id}_T .h \circledast h) && \text{defn. of } \circledast \\
&\sim (\alpha \cdot (\text{id}_T .h)) \circledast ((f \setminus_g fg) \cdot h) && \text{by (2.1)} \\
&\sim (\alpha.h) \circledast (f \setminus_{gh} fgh) && \text{properties of cut and id} \\
&= f \setminus_{gh} (\alpha.h) && \text{defn. of } \circledast
\end{aligned}$$

Similar observations were made by Dawson, Paré, and Pronk [14, §1.5] and by Lamarche [32, p.12].

2.2. From inference rules to double cells. Having established that the free bifibration on the identity functor $\text{id}_{\mathcal{C}} : \mathcal{C} \rightarrow \mathcal{C}$ generates a double category of zigzags in \mathcal{C} , at this point let us develop a more diagrammatic presentation of $\mathbb{Z}(\mathcal{C})$ that is isomorphic to Definition 2.3 but will help us to better visualize it as a double category.

An easy but important observation is that what we have called “zigzag formulas” (that is, objects of the free bifibration $\mathbb{Z}(\mathcal{C}) := \text{Bif}(\text{id}_{\mathcal{C}})$) really are in one-to-one correspondence with zigzags of arrows in the category \mathcal{C} . Of course we have not defined what a “zigzag of arrows” is formally, but one possible definition is, say, as a graph homomorphism $Z : \tilde{L} \rightarrow \mathcal{C}$ from some orientation \tilde{L} of the line graph L_n with n edges and $n + 1$ vertices into the underlying graph of \mathcal{C} , for some $n \geq 0$. It is clear that such a graph homomorphism can be represented as a zigzag formula with n total pushes and pulls, and vice versa. For example, if \mathcal{C} is the category on the left and Z is the (preformally defined!) zigzag on the right,

$$\begin{array}{ccc}
A & \xrightarrow{f} & B & \begin{array}{l} \nearrow^g \\ \searrow_h \end{array} & C \\
& & & & D
\end{array}
\qquad
A \xrightarrow{f} B \xrightarrow{h} D \xleftarrow{h} B \xrightarrow{g} C$$

then Z can be encoded isomorphically either as a graph homomorphism $Z : \tilde{L} \rightarrow \mathcal{C}$ where

$$\begin{aligned}
\tilde{L} &= v_0 \xrightarrow{e_1} v_1 \xrightarrow{e_2} v_2 \xleftarrow{e_3} v_3 \xrightarrow{e_4} v_4 \\
Z(v_0) &= A \quad Z(v_1) = Z(v_3) = B \quad Z(v_2) = D \quad Z(v_4) = C \\
Z(e_1) &= f \quad Z(e_2) = h \quad Z(e_3) = g \quad Z(e_4) = g
\end{aligned}$$

or by the zigzag formula $Z = g^+ h^- h^+ f^+ A$. One important point is that both of these representations allow us to represent the *empty zigzag*, which corresponds to a choice of object in \mathcal{C} , and should not be confused with a one-step zigzag along an identity arrow.

Now turning to double cells, observe that each logical inference rule of the sequent calculus, recalled below for reference, may be interpreted as a generating double cell of $\mathbb{Z}(\mathcal{C})$

(to make the notation lighter, we omit arrows from cell names):

$$\begin{array}{cccc}
\frac{S \rightrightarrows T}{\frac{f^+ S \rightrightarrows T}{g} \mathbb{L}f^+} & \frac{S' \rightrightarrows S}{\frac{S' \rightrightarrows f^+ S}{f'f} \mathbb{R}f^+} & \frac{T \rightrightarrows T'}{\frac{g^- T \rightrightarrows T'}{gg'} \mathbb{L}g^-} & \frac{S \rightrightarrows T}{\frac{S \rightrightarrows g^- T}{f} \mathbb{R}g^-} \\
\Downarrow & \Downarrow & \Downarrow & \Downarrow \\
\begin{array}{ccc} A \xrightarrow{fg} C \\ f \downarrow \quad \mathbb{L}^+ \quad \parallel \\ B \xrightarrow{g} C \end{array} & \begin{array}{ccc} A' \xrightarrow{f'} A \\ \parallel \quad \mathbb{R}^+ \quad \downarrow f \\ A' \xrightarrow{f'f} B \end{array} & \begin{array}{ccc} C \xrightarrow{g'} C' \\ g \uparrow \quad \mathbb{L}^- \quad \parallel \\ B \xrightarrow{gg'} C' \end{array} & \begin{array}{ccc} A \xrightarrow{fg} C \\ \parallel \quad \mathbb{R}^- \quad \uparrow g \\ A \xrightarrow{f} B \end{array}
\end{array} \tag{2.2}$$

Formally, since we defined $\mathbb{Z}(\mathcal{C})$ as an internal category, each of these cells corresponds to an arrow in $\mathbb{Z}(\mathcal{C})$ with given source and target arrows in \mathcal{C} . For example, the cell we have named \mathbb{L}^+ corresponds to an arrow $\gamma : f^+ A \rightarrow C$ with $\text{src}(\gamma) = fg$ and $\text{tgt}(\gamma) = g$. These four cells (really, a family of cells parameterized by the arrows f, g, f', g') generate the double category $\mathbb{Z}(\mathcal{C})$ in the sense that every representative ζ of an arrow $\mathbb{Z}(\mathcal{C})$ may be uniquely decomposed as a vertical composition of generators $\zeta = \gamma_1 \odot \cdots \odot \gamma_n$. Indeed, it is immediate from the inductive definition of the sequent calculus that any proof in any free bifibration over \mathcal{C} may be uniquely decomposed as the action of a list of generators on an initial axiom.

Proposition 2.4. *For any derivation $\alpha : S \rightrightarrows_f T$ there exists a unique series of generators $\gamma_1, \dots, \gamma_n \in \mathbb{Z}(\mathcal{C})$ of the form (2.2) such that $\alpha = \delta \circledast_p \gamma_1 \circledast_p \cdots \circledast_p \gamma_n$, where $\delta = \chi_p(\alpha)$.*

It is possible to go in the other direction and extract a direct presentation of $\mathbb{Z}(\mathcal{C})$, considering its double cells as formal vertical compositions of generators—we will refer to such lists of generators as “stacks”—modulo an equivalence relation. Our definition of permutation equivalence for the sequent calculus (Definition 1.4) translates directly to the following relations on stacks:

$$\begin{array}{cccc}
\begin{array}{ccc} \cdot \xrightarrow{g} \cdot \\ f \uparrow \quad \mathbb{L}^- \quad \parallel \\ \cdot \xrightarrow{-fg} \cdot \\ \parallel \quad \mathbb{R}^+ \quad \downarrow h \\ \cdot \xrightarrow{fgh} \cdot \end{array} & \sim & \begin{array}{ccc} \cdot \xrightarrow{g} \cdot \\ \parallel \quad \mathbb{R}^+ \quad \downarrow h \\ \cdot \xrightarrow{-gh} \cdot \\ f \uparrow \quad \mathbb{L}^- \quad \parallel \\ \cdot \xrightarrow{fgh} \cdot \end{array} & & \begin{array}{ccc} \cdot \xrightarrow{fg} \cdot \\ f \downarrow \quad \mathbb{L}^+ \quad \parallel \\ \cdot \xrightarrow{-g} \cdot \\ \parallel \quad \mathbb{R}^+ \quad \downarrow h \\ \cdot \xrightarrow{gh} \cdot \end{array} & \sim & \begin{array}{ccc} \cdot \xrightarrow{fg} \cdot \\ \parallel \quad \mathbb{R}^+ \quad \downarrow h \\ \cdot \xrightarrow{-fgh} \cdot \\ f \downarrow \quad \mathbb{L}^+ \quad \parallel \\ \cdot \xrightarrow{gh} \cdot \end{array} \\
\begin{array}{ccc} \cdot \xrightarrow{gh} \cdot \\ f \uparrow \quad \mathbb{L}^- \quad \parallel \\ \cdot \xrightarrow{-fgh} \cdot \\ \parallel \quad \mathbb{R}^- \quad \uparrow h \\ \cdot \xrightarrow{fg} \cdot \end{array} & \sim & \begin{array}{ccc} \cdot \xrightarrow{gh} \cdot \\ \parallel \quad \mathbb{R}^- \quad \uparrow h \\ \cdot \xrightarrow{-g} \cdot \\ f \uparrow \quad \mathbb{L}^- \quad \parallel \\ \cdot \xrightarrow{fg} \cdot \end{array} & & \begin{array}{ccc} \cdot \xrightarrow{fgh} \cdot \\ f \downarrow \quad \mathbb{L}^+ \quad \parallel \\ \cdot \xrightarrow{-gh} \cdot \\ \parallel \quad \mathbb{R}^- \quad \uparrow h \\ \cdot \xrightarrow{g} \cdot \end{array} & \sim & \begin{array}{ccc} \cdot \xrightarrow{fgh} \cdot \\ \parallel \quad \mathbb{R}^- \quad \uparrow h \\ \cdot \xrightarrow{-fg} \cdot \\ f \downarrow \quad \mathbb{L}^+ \quad \parallel \\ \cdot \xrightarrow{g} \cdot \end{array}
\end{array} \tag{2.3}$$

Our inductive definition of the identity derivation $\text{id}_S : S \Longrightarrow_{\text{id}_A} S$ for every bifibrational formula $S \sqsubset A$ (Definition 1.2) likewise translates to the following definition of identity double cells for every one-step zigzag, which can be vertically stacked to define the identity cell for an arbitrary length zigzag $Z : A \rightsquigarrow B$:

$$\begin{array}{ccc}
 A \text{ --- } A & & B \text{ --- } B \\
 \parallel & \mathbf{R}^+ & \downarrow f \\
 A \text{ --- } f \rightarrow B & & \\
 f \downarrow & \mathbf{L}^+ & \parallel \\
 B \text{ --- } B & &
 \end{array}
 \qquad
 \begin{array}{ccc}
 B \text{ --- } B & & \\
 \uparrow g & \mathbf{L}^- & \parallel \\
 A \text{ --- } g \rightarrow B & & \\
 \parallel & \mathbf{R}^- & \uparrow g \\
 A \text{ --- } A & &
 \end{array}
 \tag{2.4}$$

We elide the explicit definition of horizontal composition for $\mathbb{Z}(\mathcal{C})$ in terms of stacks, but suffice it to say that our cut-elimination procedure (Definition 1.8) can be repeated. We should mention that horizontal composition $\alpha \cdot \beta$ is an operation taking two stacks with matching zigzags along their boundaries $\text{tgt}(\alpha) = \text{src}(\beta)$, where the left and right boundaries of a stack can be computed by concatenating all of the one-step zigzags (i.e., up or down arrows) that appear on that side of the generators, ignoring empty zigzags.

Now, a corollary of Proposition 2.4 is that any free bifibration may be easily reconstructed from the double category of zigzags, as observed by Lamarche [32, p.22].

Corollary 2.5. $\text{Bif}(p) \cong \mathcal{D}_p \times_{\text{src}} \mathbb{Z}(\mathcal{C})$ and $\Lambda_p \equiv \text{tgt}^+ \text{src}^- p$, or diagrammatically:

$$\begin{array}{ccccc}
 \mathcal{D} & \longleftarrow & \text{Bif}(p) & & \\
 p \downarrow & & \downarrow & \searrow \Lambda_p & \\
 \mathcal{C} & \xleftarrow{\text{src}} & \mathbb{Z}(\mathcal{C}) & \xrightarrow{\text{tgt}} & \mathcal{C}
 \end{array}$$

This suggests an isomorphic representation of sequent calculus proofs, as a stack of generators in $\mathbb{Z}(\mathcal{C})$ paired with a single compatible arrow of \mathcal{D} . We refer the reader again back to the Introduction for an example of this representation on the right side of Figure 1 (we will explain the overlaid string diagram further below).

2.3. Aside: displayed bifibrations. To better understand the relationship between the zigzag double category and the bifibrational sequent calculus, we find it revealing to pursue a more conceptual analysis of the action defined in Section 2.1. This analysis will not be needed in the rest of the paper and so may be safely skipped, but could interest readers familiar with Bénabou’s work [6, §7] and so-called *displayed categories* [1].

Recall that we wrote Frm_A for the set of bifibrational formulas (that is, objects of $\text{Bif}(p)$) lying over an object $A \in \mathcal{C}$. Let us similarly write

$$\text{Pf}_h := \{ \alpha : S \Longrightarrow_h T \mid S \in \text{Frm}_A, T \in \text{Frm}_C \} / \sim$$

for the set of equivalence classes of proofs (that is, arrows of $\text{Bif}(p)$) lying over an arrow $h : A \rightarrow C$. This set is equipped with evident projection functions $\text{Frm}_A \leftarrow \text{Pf}_h \rightarrow \text{Frm}_C$ returning the domain and codomain of a proof. The action $\text{Bif}(p) \times_{\Lambda_p \times \text{src}} \mathbb{Z}(\mathcal{C}) \rightarrow \text{Bif}(p)$ now permits the definition of a *double functor*

$$\partial \Lambda_p : \mathbb{Z}(\mathcal{C})^\top \longrightarrow \text{Span}(\text{Set})$$

from the transpose of the zigzag double category into the double category of spans of sets, by the following mappings:

$$\begin{array}{l}
A \mapsto \text{Frm}_A \\
f : A \rightarrow B \mapsto \text{Frm}_A \leftarrow \text{Pf}_f \rightarrow \text{Frm}_B \\
Z : A \rightsquigarrow B \mapsto - \otimes Z : \text{Frm}_A \rightarrow \text{Frm}_B
\end{array}$$

$$\begin{array}{ccc}
A \xrightarrow{f} B & & \text{Frm}_A \longleftarrow \text{Pf}_f \longrightarrow \text{Frm}_B \\
\downarrow Z \quad \zeta \quad \downarrow Z' & \mapsto & \begin{array}{ccc} - \otimes Z \downarrow & \downarrow - \otimes \zeta & \downarrow - \otimes Z' \\ \text{Frm}_C \longleftarrow \text{Pf}_g \longrightarrow \text{Frm}_D & & \end{array}
\end{array}$$

Note the transpose is necessary because the tight morphisms of $\mathbb{Z}(\mathcal{C})$ are sent to loose morphisms of $\text{Span}(\text{Set})$.

In particular, $\partial\Lambda_p$ maps each generator of $\mathbb{Z}(\mathcal{C})$ to a double cell of spans:

$$\begin{array}{ccc}
\begin{array}{ccc} \text{Frm}_A \longleftarrow \text{Pf}_{fg} \longrightarrow \text{Frm}_C & & \text{Frm}_{A'} \longleftarrow \text{Pf}_{f'} \longrightarrow \text{Frm}_A \\ f^+ \downarrow & \downarrow L^+ & \parallel \\ \text{Frm}_B \longleftarrow \text{Pf}_g \longrightarrow \text{Frm}_C & & \text{Frm}_{A'} \longleftarrow \text{Pf}_{f'f} \longrightarrow \text{Frm}_B \end{array} & & \begin{array}{ccc} \text{Frm}_A \longleftarrow \text{Pf}_{fg} \longrightarrow \text{Frm}_C & & \text{Frm}_A \longleftarrow \text{Pf}_f \longrightarrow \text{Frm}_B \\ g^- \downarrow & \downarrow L^- & \parallel \\ \text{Frm}_B \longleftarrow \text{Pf}_{gg'} \longrightarrow \text{Frm}_{C'} & & \text{Frm}_A \longleftarrow \text{Pf}_f \longrightarrow \text{Frm}_B \end{array} \\
& & \text{(2.5)}
\end{array}$$

One might think of this interpretation as internalizing the idea of an inference rule as a transformation from a proof of the premise(s) into a proof of the conclusion. Notice that this double functor $\partial\Lambda_p : \mathbb{Z}(\mathcal{C})^\top \rightarrow \text{Span}(\text{Set})$

- *strictly preserves composition of zigzags*: since \otimes_p is a strict action, with the equation $x \otimes (a \odot b) = (x \otimes a) \otimes b$ corresponding to simple chaining of logical connectives and inference rules;
- *laxly preserves composition of arrows in \mathcal{C}* : since the span morphisms $\text{Pf}_f \cdot \text{Pf}_g \rightarrow \text{Pf}_{fg}$, which correspond to cut-elimination, are typically not invertible.

In fact, this construction is very general and can be applied to *any* bifibration $q : \mathcal{E} \rightarrow \mathcal{C}$ to derive a lax double functor $\partial q : \mathbb{Z}(\mathcal{C})^\top \rightarrow \text{Span}(\text{Set})$. On the horizontal category $\mathcal{H}(\mathbb{Z}(\mathcal{C})) = \mathcal{C}$, this functor sends every object $A \in \mathcal{C}$ to the set $\text{Ob}_A = \{ S \in \mathcal{E} \mid q(S) = A \}$ of objects lying over it, and similarly every arrow $f : A \rightarrow B$ to the set $\text{Arr}_f = \{ \alpha \in \mathcal{E} \mid q(\alpha) = f \}$ of arrows lying over it, which is equipped with evident projections $\text{Ob}_A \leftarrow \text{Arr}_f \rightarrow \text{Ob}_B$. This

extends to the generators of $\mathbb{Z}(\mathcal{C})$ by interpreting the span morphisms

$$\begin{array}{ccc}
\text{Ob}_A \longleftarrow \text{Arr}_{fg} \longrightarrow \text{Ob}_C & & \text{Ob}_{A'} \longleftarrow \text{Arr}_{f'} \longrightarrow \text{Ob}_A \\
f^+ \downarrow & \downarrow \mathbf{L}^+ & \parallel \\
\text{Ob}_B \longleftarrow \text{Arr}_g \longrightarrow \text{Ob}_C & & \text{Ob}_{A'} \longleftarrow \text{Arr}_{f'f} \longrightarrow \text{Ob}_B \\
& & \downarrow \mathbf{R}^+ \quad \downarrow f^+ \\
& & \text{Ob}_B \\
\text{Ob}_C \longleftarrow \text{Arr}_{g'} \longrightarrow \text{Ob}_{C'} & & \text{Ob}_A \longleftarrow \text{Arr}_{fg} \longrightarrow \text{Ob}_C \\
g^- \downarrow & \downarrow \mathbf{L}^- & \parallel \\
\text{Ob}_B \longleftarrow \text{Arr}_{gg'} \longrightarrow \text{Ob}_{C'} & & \text{Ob}_A \longleftarrow \text{Arr}_f \longrightarrow \text{Ob}_B \\
& & \downarrow \mathbf{R}^- \quad \downarrow g^- \\
& & \text{Ob}_B
\end{array} \tag{2.6}$$

using the bifibrational structure of q , very similarly to how we defined the interpretation functor $\llbracket - \rrbracket_\theta : \Lambda_p \rightarrow \mathcal{E}$ from the free bifibration in Section 1.4. For example, the function $f^+ : \text{Ob}_A \rightarrow \text{Ob}_B$ maps any object $S \in \mathcal{E}$ lying over A to (the object component of) its pushforward $f^+ S$, the span morphism $\mathbf{L}^+ : \text{Arr}_{fg} \rightarrow \text{Arr}_g$ is interpreted using the universal property of this pushforward, and $\mathbf{R}^+ : \text{Arr}_{f'} \rightarrow \text{Arr}_{f'f}$ by postcomposition with the associated $+$ -cartesian arrow $S \rightarrow f^+ S$.

Conversely, any such lax double functor $F : \mathbb{Z}(\mathcal{C})^\top \rightarrow \text{Span}(\text{Set})$ determines a category $\int_{\mathcal{C}} F$ together with a functor $\pi_{\mathcal{C}} : \int_{\mathcal{C}} F \rightarrow \mathcal{C}$ equipped with the structure of a bifibration. Indeed, the category $\int_{\mathcal{C}} F$ is constructed in the usual way as a displayed category: objects of $\int_{\mathcal{C}} F$ are given by pairs (A, S) of an object $A \in \mathcal{C}$ and an element $S \in F(A)$, arrows are given by pairs (f, α) of an arrow $f \in \mathcal{C}$ and an element $\alpha \in F(f)$ (meaning an element of the apex of the span $F(A) \leftarrow F(f) \rightarrow F(B)$), and composition and identity arrows are defined using the lax structure of F .

In turn, the tight component of the lax double functor $F : \mathbb{Z}(\mathcal{C})^\top \rightarrow \text{Span}(\text{Set})$ enables us to go further and equip the projection functor $\pi_{\mathcal{C}} : \int_{\mathcal{C}} F \rightarrow \mathcal{C}$ with the structure of a bifibration. For example, the pushforward along an arrow $f : A \rightarrow B$ is defined using the functions $F(f^+ A) : F(A) \rightarrow F(B)$ to compute pushforward objects, and applying $F(\mathbf{R}^+) : F(\text{id}_A) \rightarrow F(f)$ to an identity arrow to compute the associated cartesian arrow. The universal property of this pushforward is derived from the interpretation of $F(\mathbf{L}^+)$ and the defining equations of $\mathbb{Z}(\mathcal{C})$.

The constructions ∂ and \int witness an equivalence

$$\begin{array}{ccc}
\mathcal{D} & & \\
\downarrow \text{bifibration} & \iff & \mathbb{Z}(\mathcal{C})^\top \xrightarrow{\text{lax}} \text{Span}(\text{Set}) \\
\mathcal{C} & &
\end{array}$$

and one might therefore reasonably refer to a lax double functor $F : \mathbb{Z}(\mathcal{C})^\top \rightarrow \text{Span}(\text{Set})$ as a “displayed bifibration”.

We will not pursue these reflections further in this paper, although we think they could eventually be helpful for developing a more algebraic perspective on the sequent calculus.

2.4. A dagger structure on zigzags. One advantage of considering free bifibrations in terms of the double category of zigzags is that the latter clearly has a large degree of symmetry. Of particular relevance is a dagger-type involution

$$(-)^\dagger : \mathbb{Z}(\mathcal{C}) \rightarrow \mathbb{Z}(\mathcal{C})^v$$

that reverses the direction of vertical arrows and double cells. Formally, this corresponds to an endofunctor on $\mathcal{Z}(\mathcal{C})$ swapping source and target while leaving the underlying arrows in \mathcal{C} invariant:

$$\begin{array}{ccc}
 & (-)^\dagger & \\
 \mathcal{Z}(\mathcal{C}) & \xrightarrow{\quad} & \mathcal{Z}(\mathcal{C}) \\
 & \sim & \\
 & (-)^\dagger & \\
 \text{src} \swarrow & & \searrow \text{tgt} \\
 & \mathcal{C} &
 \end{array}$$

At the level of formulas, $(-)^{\dagger}$ turns pushforwards into pullbacks and vice versa while reversing the sequence of connectives applied (e.g., $(h^+ g^- f^+ \square)^{\dagger} = f^- g^+ h^- \square$). More geometrically, the transformation may be visualized as reflecting the generating cells (2.2) across the x -axis (or equivalently reading them from bottom to top), which swaps Lf^+ with Lf^- and Rg^+ with Rg^- . The $(-)^{\dagger}$ transformation extends to arbitrary stacks of double cells by $(\gamma_1 \odot \cdots \odot \gamma_n)^{\dagger} = \gamma_n^{\dagger} \odot \cdots \odot \gamma_1^{\dagger}$. The following is immediate.

Proposition 2.6. $(-)^{\dagger}$ is an involutive isomorphism.

As a corollary, we can derive that the pairing of the source and target functors is a bifibration.

Theorem 2.7 (Lamarche [33, Theorem 3]). $(\text{src}, \text{tgt}) : \mathcal{Z}(\mathcal{C}) \rightarrow \mathcal{C} \times \mathcal{C}$ is a bifibration.

Proof. By Proposition 2.6, $\text{src} : \mathcal{Z}(\mathcal{C}) \rightarrow \mathcal{C}$ is a bifibration since it inherits a bifibrational structure from tgt via the $(-)^{\dagger}$ involution. As we saw earlier by Proposition 1.24, the pairing $(\text{src}, \text{tgt}) : \mathcal{Z}(\mathcal{C}) \rightarrow \mathcal{C} \times \mathcal{C}$ is a morphism of bifibrations from tgt to π_2 , which means that tgt -cartesian arrows in $\mathcal{Z}(\mathcal{C})$ are mapped by src to identity arrows in \mathcal{C} . Symmetrically, applying the $(-)^{\dagger}$ involution, it is also a morphism of bifibrations from src to π_1 . This ensures that the pairing is a bifibration $(\text{src}, \text{tgt}) : \mathcal{Z}(\mathcal{C}) \rightarrow \mathcal{C} \times \mathcal{C}$, since the two bifibrational structures act independently. \square

2.5. Recovering the free adjoint construction. Given a double category \mathbb{E} seen as an internal category $\mathbb{E} = (\mathcal{C}, \mathcal{E}, \text{src}_{\mathcal{E}}, \text{tgt}_{\mathcal{E}}, U_{\mathcal{E}}, \odot_{\mathcal{E}})$, we write $\mathcal{H}(\mathbb{E}) = \mathcal{C}$ for its underlying *horizontal category*. Any double category \mathbb{E} also induces a *vertical 2-category* denoted $\mathcal{V}(\mathbb{E})$, with same set of objects, 1-cells given by vertical arrows $F : A \dashrightarrow B$, and 2-cells $\alpha : F \Longrightarrow G$ given by globular double cells

$$\begin{array}{ccc}
 A & \xlongequal{\quad} & A \\
 F \downarrow & \alpha & \downarrow G \\
 B & \xlongequal{\quad} & B
 \end{array}$$

that is, cells whose source and target are identities.

A double category \mathbb{E} is said to be *fibrant* (or to be a *framed bicategory* [49]) just in case the pairing of its source and target functors $(\text{src}_{\mathcal{E}}, \text{tgt}_{\mathcal{E}}) : \mathcal{E} \rightarrow \mathcal{C} \times \mathcal{C}$ is a fibration, or equivalently an opfibration, and therefore a bifibration [49, Theorem 4.1]. This is equivalent to asking that the double category has all *companions and conjoints* meaning that for every

horizontal arrow $f : A \rightarrow B$ there is both a vertical arrow $F_f : A \dashrightarrow B$ and a vertical arrow $G_f : B \dashrightarrow A$ equipped with a pair of cells

$$\begin{array}{ccc} A & \xlongequal{\quad} & A \\ \parallel & \alpha & \downarrow F_f \\ A & \xrightarrow{f} & B \end{array} \quad \begin{array}{ccc} B & \xlongequal{\quad} & B \\ G_f \downarrow & \beta & \parallel \\ A & \xrightarrow{f} & B \end{array}$$

satisfying a few equations (see [16] and [49, Theorem A.2]; note that the conventions for the roles of horizontal and vertical arrows are often swapped in the literature).

Fibrant double categories are moreover essentially equivalent to *proarrow equipments* [49, Appendix C]. In particular, any fibrant double category \mathbb{E} canonically determines an identity-on-objects pseudofunctor $(-)^+ : \mathcal{C} \rightarrow \mathcal{K}$ from its underlying horizontal category $\mathcal{C} = \mathcal{H}(\mathbb{E})$ to its vertical 2-category $\mathcal{K} = \mathcal{V}(\mathbb{E})$ that sends every horizontal arrow f to its vertical companion F_f , which has a right adjoint $F_f \dashv G_f$ given by the vertical conjoint. Conversely, given a 2-category \mathcal{K} equipped with a pseudofunctor $I : \mathcal{C} \rightarrow \mathcal{K}$ such that every arrow $f \in \mathcal{C}$ is mapped to an arrow $I f \in \mathcal{K}$ with a right adjoint, one can define an associated double category $\mathbb{E}_{\mathcal{K}}$ whose objects and horizontal arrows are the objects and arrows of \mathcal{C} , whose vertical arrows $M : A \rightarrow B$ are 1-cells $M : I A \rightarrow I B$ in \mathcal{K} , and whose double cells

$$\begin{array}{ccc} A & \xrightarrow{f} & C \\ M \downarrow & \alpha & \downarrow N \\ B & \xrightarrow{g} & D \end{array}$$

are 2-cells $\alpha : M \cdot I g \Rightarrow I f \cdot N$ in \mathcal{K} . Then $\mathbb{E}_{\mathcal{K}}$ is a fibrant double category [49, Prop. C.3].

The correspondence between fibrant double categories and proarrow equipments means that we can recover Dawson, Paré, and Pronk's free adjoint construction $\Pi_2(\mathcal{C})$ up to equivalence as the vertical 2-category of $\mathbb{Z}(\mathcal{C})$.


Theorem 2.8. $\mathbb{Z}(\mathcal{C})$ is the free fibrant double category over \mathcal{C} .

Proof. By Theorem 2.7, $\mathbb{Z}(\mathcal{C})$ is a fibrant double category. Moreover, given any other fibrant double category $\mathbb{E} = (\mathcal{C}, \mathcal{E}, \text{src}_{\mathcal{E}}, \text{tgt}_{\mathcal{E}}, U_{\mathcal{E}})$ with underlying horizontal category $\mathcal{H}(\mathbb{E}) = \mathcal{C}$, there exists a unique double functor $\mathbb{Z}(\mathcal{C}) \rightarrow \mathbb{E}$ whose underlying functor $\mathbb{Z}(\mathcal{C}) \rightarrow \mathcal{E}$ is a morphism of bifibrations from (src, tgt) to $(\text{src}_{\mathcal{E}}, \text{tgt}_{\mathcal{E}})$. Indeed, any double functor $\mathbb{Z}(\mathcal{C}) \rightarrow \mathbb{E}$ is determined by the image of the generators (2.2), and the image of the generators is uniquely determined by the bifibrational structure of $\text{src}_{\mathcal{E}} : \mathcal{E} \rightarrow \mathcal{C}$ and $\text{tgt}_{\mathcal{E}} : \mathcal{E} \rightarrow \mathcal{C}$. \square


In particular, there is an identity-on-objects pseudofunctor $(-)^+ : \mathcal{C} \rightarrow \mathcal{V}(\mathbb{Z}(\mathcal{C}))$ mapping every arrow $f : A \rightarrow B$ of \mathcal{C} to the forward zigzag $f^+ A : A \rightsquigarrow B$, which has a right adjoint given by the backward zigzag $f^- B : B \rightsquigarrow A$. Observe that $(-)^+$ is only a pseudofunctor rather than a strict functor because, as we already observed, the free bifibration is *not* split in the sense that the canonical identities $(g \circ f)^+ S \equiv g^+ f^+ S$ and $\text{id}^+ U \equiv U$ only hold as isomorphisms in the fiber categories rather than as strict equalities. In contrast, Dawson, Paré, and Pronk constructed $\Pi_2(\mathcal{C})$ in a way that admits a strict embedding functor $(-)^+ : \mathcal{C} \rightarrow \Pi_2(\mathcal{C})$. Nevertheless, up to splitting of the free bifibration (see Theorem 3.5 below), the two constructions are equivalent.

Corollary 2.9. $\mathcal{V}(\mathbb{Z}(\mathcal{C})) \simeq \Pi_2(\mathcal{C})$.


2.6. From double cells to string diagrams. Another reward of the double categorical analysis is that it leads to a natural graphical calculus, derived by applying standard conventions of string diagrams to the generators of $\mathbb{Z}(\mathcal{C})$. The idea is simple: we begin by dualizing the generators (2.2) geometrically, so that cells become vertices and objects become regions, while arrows become edges separating the regions which meet at a vertex:



Then we remove the vertices, depicting the cells by simply bending edges appropriately:



The permutation equivalences (1.2)–(1.5) corresponding to the stack relations (2.3) now have natural interpretations as planar isotopies. For example, (1.2) and (1.3) (which permute R^+ with L^- and L^+ respectively) may be depicted as follows:



The equations defining cut-elimination (Definition 1.8) also have natural topological interpretations. Below we illustrate one principal cut case and one commutative cut case:



We recover in this way the calculus of planar arc diagrams that was used by Dawson, Paré, and Pronk [14, §6] to represent 2-cells in $\Pi_2(\mathcal{C})$ for the special case that \mathcal{C} is a free category. More generally, when \mathcal{C} is not free we also need to include vertices representing equations

in the base category. For example, an equation $h = fg$ that is applied implicitly within a sequent proof, as depicted on the left side of Figure 1, may be represented by a trivalent vertex in the corresponding string diagrammatic proof, as depicted on the right side of that figure. As also depicted there, we also need to include terminating vertices corresponding to the initial axioms in \mathcal{D} .

The graphical calculus that we thus obtain is comparable to the string diagram calculus for proarrow equipments of David Jaz Myers [43].

Given the isomorphisms between the three different ways of representing arrows of free bifibrations—as sequent calculus proofs, as stacks of double cells, or as string diagrams—in the sequel we will freely interweave them without comment.

2.7. Related work. Since $\mathcal{Bif}(p)$ and $\Pi_2(\mathcal{C})$ are determined by their universal properties, a detailed comparison with the constructions of Dawson, Paré, and Pronk (DPP) [14] and of Lamarche [32] feels a bit superfluous, and the reader interested in the details of those constructions is referred to those papers. Still, a few points are worth highlighting. Zigzags and morphisms between them are essential in all three constructions, but whereas Lamarche (like us) considers zigzags as not necessarily alternating sequences of signed arrows, DPP restrict to strictly alternating zigzags beginning and ending with a forward arrow. This is motivated by the fact that they wanted the embedding of \mathcal{C} into $\Pi_2(\mathcal{C})$ to define a strict functor $(-)^+ : \mathcal{C} \rightarrow \Pi_2(\mathcal{C})$ rather than a pseudofunctor. Although the free bifibration is not split and hence $(-)^+ : \mathcal{C} \rightarrow \mathcal{V}(\mathbb{Z}(\mathcal{C}))$ defines a proper pseudofunctor, we will also have motivations to consider strictly alternating zigzags beginning in Section 3.3.

All three constructions likewise involve first introducing some notion of “pre-morphism” between zigzags and then defining zigzag morphisms as equivalence classes of pre-morphisms modulo an equivalence relation. Lamarche notes that “these classes can be quite complex, and there is no hope of finding normal forms in general” [32, p.16]. In the next section we will establish that it is possible to obtain normal forms in the case that the base category is *factorization preordered*, a condition that was also identified by DPP as a sufficient condition for decidability.

Interestingly, categories and double categories of zigzags have also been introduced in other contexts. A category of zigzags and zigzag morphisms was defined by Heidemann, Reutter, and Vicary [22] motivated by the formalization of higher-dimensional diagrams and n -categories. Their definition of zigzag map [22, Def. 2.2] is nearly identical to DPP’s definition of *fence* [14, Def. 1.2], without condition (c) which enforces globularity. However, Heidemann, Reutter, and Vicary do not quotient zigzag maps by permutation equivalence. Hence their zigzag category is *not* equivalent to $\mathcal{Z}(\mathcal{C})$, but rather to Lamarche’s category $\mathbb{R}(\mathcal{C})$, from which $\mathbb{P}(\mathcal{C}) \simeq \mathcal{Z}(\mathcal{C})$ is obtained as the quotient by permutation equivalence. Our zigzag double category $\mathbb{Z}(\mathcal{C})$ is equivalent to the “weave double category” by Williams [55, Def. 22], as follows from Williams’ definition as a fibered coproduct of arrow and oparrow double categories, which freely add companions and conjoints respectively. As mentioned, Lamarche first identified the double category structure on $\mathbb{P}(\mathcal{C}) \simeq \mathcal{Z}(\mathcal{C})$, although it was not emphasized in his papers [32, 33].² An alternative double categorical perspective on

²According to personal communication from François Lamarche (November 2025), he noticed the close connection between this double category and the 2-category $\Pi_2(\mathcal{C})$, discussing it in October 2013 by email with Paré, Dawson, and Pronk. Although the originally submitted version of the manuscript [32] did not mention DPP’s work, Lamarche says that this was an oversight and that he had originally intended to include

the $\Pi_2(\mathcal{C})$ construction was also developed by DPP at the end of their paper on the Span construction [16].

Perhaps one of the most interesting open questions in double category theory is an explicit construction that freely adds companions and conjoints to a double category \mathbb{D} . For each category \mathcal{C} , there is a double category $\mathbb{H}(\mathcal{C})$ whose objects and tight morphisms are the objects and morphism of \mathcal{C} , and whose loose morphisms and cells are identities. In Theorem 2.8, we established $\mathcal{Z}(\mathcal{C})$ as the free companion and conjoint completion of the double category $\mathcal{D} = \mathbb{H}(\mathcal{C})$. The free cornering of a strict monoidal category, introduced by Nester [44] (see also Section 1 of [9]) is another example of a free companion and conjoint completion, in this case, of a strict monoidal category viewed as a double category. Coulembier, Street, and van den Bergh likewise considered the closely related problem of freely adjoining duals to specified objects of a monoidal category [13]

3. FOCUSED NORMAL FORMS

In Section 1, we described how to construct $\text{Bif}(p)$, the total category of the free bifibration on a functor $p : \mathcal{D} \rightarrow \mathcal{C}$, by taking its arrows to be equivalence classes of derivations in a sequent calculus modulo a notion of permutation equivalence, with composition defined by the admissible cut rule. In this section we will use additional ideas from proof theory—in particular the notion of *focusing* originating in the proof theory of linear logic—to get a more precise characterization of the relative homsets $\text{Bif}(p)_f(S, T)$, via a series of sequent calculi with progressively more rigid structure. In the case that \mathcal{C} is *factorization preordered* in the sense of Johnstone [25], we derive a normal form theorem that leads to a complete inductive characterization of these relative homsets, without any quotienting. As a corollary, we obtain a simple procedure for deciding permutation equivalence of derivations, as well as for enumerating relative homsets without duplicates.

Lurking in the background is the fact that in the general case, permutation equivalence is undecidable. Dawson, Paré, and Pronk [15] established an analogue of this undecidability result for their freely adjoining adjoints construction $\Pi_2(\mathcal{C})$, and their proof translates readily to our presentation of the zigzag double category $\mathbb{Z}(\mathcal{C})$. Nevertheless, we will see in Sections 4 and 5 that the canonicity result has wide applications.

3.1. Background: invertibility, polarity, and synthetic connectives. Focusing, introduced by Andreoli [2] as an optimization for linear logic proof search and subsequently applied in other settings, combines several observations about the rules of sequent calculus and the way that they determine the behavior of the connectives. To provide some context and motivation for our approach, we take a moment here to briefly recall these observations in the setting of linear logic [18], before explaining how to transfer these ideas to analyze free bifibrations.

We consider a fragment of the sequent calculus for intuitionistic linear logic, containing the inference rules for tensor \otimes (also called *multiplicative conjunction*) and internal hom \multimap (also called *linear implication*):

$$\frac{A, B, \Gamma \vdash C}{A \otimes B, \Gamma \vdash C} \text{L}\otimes \quad \frac{\Gamma \vdash A \quad \Delta \vdash B}{\Gamma, \Delta \vdash A \otimes B} \text{R}\otimes \quad \frac{\Gamma \vdash A \quad B, \Delta \vdash C}{A \multimap B, \Gamma, \Delta \vdash C} \text{L}\multimap \quad \frac{\Gamma, A \vdash B}{\Gamma \vdash A \multimap B} \text{R}\multimap$$

a comparison. The manuscript [32] was submitted for publication but unfortunately ran into some negative referee reports and never made it to print.

Here the basic judgment $\Gamma \vdash B$ takes a single formula B on the right and a list of formulas $\Gamma = A_1, \dots, A_n$ on the left considered up to reordering, and we write a comma for concatenation of lists. The above inference rules are combined with initial axioms of the form $X \vdash X$ for every atomic formula X . Identity derivations for arbitrary formulas as well as the cut rule

$$\frac{}{A \vdash A} \text{id}_A \qquad \frac{\Gamma \vdash A \quad A, \Delta \vdash B}{\Gamma, \Delta \vdash B} \text{cut}$$

are admissible, similarly to how we saw in Section 1.

A basic observation about the inference rules for the logical connectives is that the rules $\text{L}\otimes$ and $\text{R}\multimap$ are *invertible* in the proof-theoretic sense, meaning that if the conclusion is derivable then the premise is derivable. Indeed, we can derive inverses of $\text{L}\otimes$ and $\text{R}\multimap$ using the admissible cut and identity rules, combined with instances of $\text{R}\otimes$ and $\text{L}\multimap$:

$$\frac{\frac{\overline{A \vdash A} \quad \overline{B \vdash B}}{A, B \vdash A \otimes B} \quad A \otimes B, \Gamma \vdash C}{A, B, \Gamma \vdash C} \qquad \frac{\Gamma \vdash A \multimap B \quad \frac{\overline{A \vdash A} \quad \overline{B \vdash B}}{A \multimap B, A \vdash B}}{\Gamma, A \vdash B}$$

One interest of invertible inference rules is that they may be safely applied as *goal transformers* during proof search. For example, if our goal is to prove a sequent of the form $A \otimes B, \Gamma \vdash C$, we can immediately apply the $\text{L}\otimes$ rule and try to prove $A, B, \Gamma \vdash C$, knowing that if there is any proof of the original goal then there is one that goes via $\text{L}\otimes$. Moreover, if a goal matches the conclusion of more than one invertible rule, then the order in which these rules are invoked is irrelevant. For example, a goal of the form $A \otimes B, \Gamma \vdash C \multimap D$ can be transformed into the equivalent goal $A, B, \Gamma, C \vdash D$ by applying the $\text{L}\otimes$ and $\text{R}\multimap$ rules in either order.

On the other hand, the rules $\text{R}\otimes$ and $\text{L}\multimap$ are not invertible in this sense. Intuitively, this is due to the fact that list concatenation is not an invertible operation, meaning that passing from the conclusion to the premises requires finding an appropriate splitting of the formulas on the left-hand side of the turnstile. As a simple counterexample to invertibility of $\text{R}\otimes$, consider the sequent $X \otimes Y \vdash X \otimes Y$. It is derivable, but attempting to prove it using the rule $\text{R}\otimes$ produces two unprovable premises—no matter how we decide to distribute the unique formula on the left of the sequent:

$$\frac{X \otimes Y \vdash X \quad \cdot \vdash Y}{X \otimes Y \vdash X \otimes Y} \qquad \frac{\cdot \vdash X \quad X \otimes Y \vdash Y}{X \otimes Y \vdash X \otimes Y}$$

Focusing originally arose from the idea that proof search could be structured to alternate between an *inversion phase*, during which invertible rules are applied eagerly as goal transformers, and a *focus phase*, during which non-invertible rules are applied to decompose a single formula and its subformulas. It was found that this greatly reduces the overall non-determinism of the sequent calculus, leading to more efficient proof search. Moreover, an interesting duality emerges, with half of the connectives being invertible on the *left* and half being invertible on the *right*.

This left or right bias was referred to as “polarity”—with \otimes being a *positive* connective and \multimap a *negative* connective—and helped to explain some of the distinctive associativity, distributivity, and non-distributivity properties of linear logic. For example, like tensor, additive disjunction \oplus is also invertible on the left, and distributes through \otimes by the

equivalence $A \otimes (B \oplus C) \equiv (A \otimes B) \oplus (A \otimes C)$. Similarly, additive conjunction $\&$ is invertible on the right, like internal hom, and distributes through the codomain of \multimap by the equivalence $A \multimap (B \& C) \equiv (A \multimap B) \& (A \multimap C)$.

Related to these properties is the possibility of considering certain compositions of the logical connectives as “synthetic connectives” [3]. For example, consider the composite formula $A_1 \multimap (A_2 \multimap B)$. By applying the left rule twice and the right rule twice, we can build the following open derivations:

$$\frac{\Gamma_1 \vdash A_1 \quad \frac{\Gamma_2 \vdash A_2 \quad B, \Delta \vdash C}{A_2 \multimap B, \Gamma_2, \Delta \vdash C} \text{L}\multimap}{A_1 \multimap (A_2 \multimap B), \Gamma_1, \Gamma_2, \Delta \vdash C} \text{L}\multimap} \quad \frac{\frac{\Gamma, A_1, A_2 \vdash B}{\Gamma, A_1 \vdash A_2 \multimap B} \text{R}\multimap}}{\Gamma \vdash A_1 \multimap (A_2 \multimap B)} \text{R}\multimap}$$

If we remove the intermediate steps, then these derivations could be seen as inference rules for a new ternary connective $(A_1, A_2) \multimap' B$:

$$\frac{\Gamma_1 \vdash A_1 \quad \Gamma_2 \vdash A_2 \quad B, \Delta \vdash C}{(A_1, A_2) \multimap' B, \Gamma_1, \Gamma_2, \Delta \vdash C} \text{L}\multimap'} \quad \frac{\Gamma, A_1, A_2 \vdash B}{\Gamma \vdash (A_1, A_2) \multimap' B} \text{R}\multimap'}$$

Adding this new connective and these inference rules preserves all the important properties of the sequent calculus, in particular admissibility of identity and cut. On the other hand, not every combination of basic connectives can be viewed as a first-class, synthetic connective. For example, there is no good way to define a right rule for the composite connective $A \multimap (B_1 \otimes B_2)$, intuitively because after applying the invertible rule $\text{R}\multimap$ to decompose the outer implication, we may need to examine the formula A on the left before applying the non-invertible rule $\text{R}\otimes$ to decompose the tensor.

3.2. Invertibility and polarity in the bifibrational calculus. Returning to bifibrational logic, we begin by observing that pushforward behaves like a positive connective and pullback like a negative connective, since f^+ has an invertible left rule and non-invertible right rule, while g^- has an invertible right rule and non-invertible left rule. We already established the invertibility of $\text{L}f^+$ and $\text{R}g^-$ with Lemma 1.11, in the strong sense that they define bijections between proofs of their premises and proofs of their conclusions up to permutation equivalence. Indeed, for any derivation $\alpha : f^+ S \Longrightarrow_g T$ there is a permutation equivalent derivation ending in the $\text{L}f^+$ rule, namely the one represented by the term $f \setminus_g (f_S \cdot \alpha)$, which is $\text{L}f^+$ applied to the derivation of $S \Longrightarrow_{f_S} T$ obtained by cutting α with the cartesian lifting $f_S : S \Longrightarrow_f f^+ S$. Dually, for any derivation $\beta : S \Longrightarrow_f g^- T$ there is a permutation equivalent derivation ending in the $\text{R}g^-$ rule, namely $(\beta \cdot \bar{g}_T) \setminus_f \bar{g}$.

To see why the $\text{R}f^+$ rule is *not* invertible in general, consider the free bifibration on the functor depicted below

$$\begin{array}{ccc} \mathcal{D} & X & \xrightarrow{\delta} Y \\ \downarrow p & & \\ \mathcal{C} & E & \begin{array}{c} \xrightarrow{e_1} A \\ \xleftarrow{e_2} A \end{array} \xrightarrow{f} B \end{array}$$

where the base category \mathcal{C} contains two parallel arrows e_1, e_2 and a third arrow f such that $e_1 f = e_2 f = h$, and where \mathcal{D} is generated by a single arrow δ such that $p(\delta) = e_1$. Observe that the sequent $X \Longrightarrow_h f^+ Y$ is derivable by applying $\text{R}f^+$ to the atomic derivation

$\delta : X \Longrightarrow_{e_1} Y$. Since $h = e_2 f$ is another valid factorization, if the Rf^+ rule were invertible then there should also be some derivation of $X \Longrightarrow_{e_2} Y$. But by conservativity (Prop. 1.23) we can immediately tell that there isn't, contradicting invertibility of Rf^+ . Dualizing the example establishes non-invertibility of the Lg^- rule. Of course these counterexamples do not preclude the possibility of Rf^+ and Lg^- being invertible in specific cases—for example, they will be when \mathcal{C} is a groupoid. Nevertheless, the behavior of the connectives in the general case justifies our adopting the polarity terminology to bifibrational formulas, and from now on we will refer to pushforwards as positive and pullbacks as negative, with atomic formulas having neutral polarity.

Continuing the analogy with linear logic, compositions of like-polarity connectives may be grouped into synthetic connectives, in the sense that $g^+ f^+$ can be represented by $(g \circ f)^+$ and $f^- g^-$ can be represented by $(f g)^-$, corresponding to the pseudofunctoriality laws of a bifibration. We already explained how these laws are witnessed by logical equivalences in the bifibrational calculus (Prop. 1.21). Here we simply point out that, similarly to the linear logic situation discussed in Section 3.1, the inference rules for the synthetic connectives are operationally equivalent to compositions of basic inference rules, in the sense that, for example, by composing the left rules and right rules for pushforward along f and g ,

$$\frac{\frac{S \Longrightarrow T}{fgh} \text{L}f^+}{\frac{f^+ S \Longrightarrow T}{gh} \text{L}g^+} \text{L}f^+ \qquad \frac{\frac{S' \Longrightarrow S}{f'} \text{R}f^+}{\frac{S' \Longrightarrow f^+ S}{f'f} \text{R}g^+} \text{R}f^+$$

we obtain derivations with identical premises to the rules for the pushforward along $g \circ f$:

$$\frac{S \Longrightarrow T}{\frac{(g \circ f)^+ S \Longrightarrow T}{h} \text{L}(g \circ f)^+} \text{L}(g \circ f)^+ \qquad \frac{S' \Longrightarrow S}{\frac{S' \Longrightarrow (g \circ f)^+ S}{f'fg} \text{R}(g \circ f)^+} \text{R}(g \circ f)^+$$

A difference with the linear logic situation is that such “synthetic” connectives are already contained in the logic, since we assume pushforward and pullback along arbitrary arrows, and the arrows of a category are closed under composition. On the other hand, in general it is not possible to reduce opposite-polarity composites like $f^+ g^-$ or $g^- f^+$ to a single pushforward or pullback, just like alternating quantifiers $\exists\forall$ or $\forall\exists$ in predicate logic cannot be reduced to a single quantifier in general.

3.3. Strictly alternating formulas, splittings, and weak focusing. Motivated by the preceding observations, in this section we introduce a restriction of the sequent calculus in which sequences of invertible rules and sequences of non-invertible rules must be performed all at once, on the left or right, but without (yet) placing any restrictions on their interleavings. This calculus may be said to be “weakly focused”, and will serve as an intermediate point before our introduction of a strongly focused calculus in the next section.

3.3.1. Strictly alternating formulas. Before defining the calculus itself, we motivate it logically and categorically by observing that every bifibrational formula is canonically isomorphic to a *strictly alternating formula*.

Definition 3.1. A bifibrational formula is **strictly alternating** if it contains no subformula of the form $g^+ f^+ S$ or $f^- g^- T$.

It is easy to see that by repeatedly applying the equivalences $g^+ f^+ S \equiv (g \circ f)^+ S$ and $f^- g^- T \equiv (f g)^- T$ (Prop. 1.21), any formula S may be normalized to a logically equivalent strictly alternating formula $[S]$. Let us write $\theta_S : S \equiv [S]$ for this family of equivalences, and define a function on derivations

$$\begin{aligned} [-] & : (S \Longrightarrow_f T) \rightarrow ([S] \Longrightarrow_f [T]) \\ [-] & := \alpha \mapsto \theta_S^{-1} \cdot \alpha \cdot \theta_T \end{aligned}$$

which we call the *strictification* map. Observe that by the subformula property (Prop. 1.1), the strictification map transforms any derivation into a derivation involving only strictly alternating formulas, since every subformula of a strictly alternating formula is strictly alternating.

Now let $sa\mathcal{Bif}(p)$ denote the full subcategory of $\mathcal{Bif}(p)$ spanned by the strictly alternating formulas, let $i : sa\mathcal{Bif}(p) \rightarrow \mathcal{Bif}(p)$ be the full and faithful inclusion functor, and let us write $\Lambda_p|_{sa} = \Lambda_p \circ i : sa\mathcal{Bif}(p) \rightarrow \mathcal{C}$ for the restriction of the domain of the free bifibration to strictly alternating formulas.

Theorem 3.2. *Strictification extends to a functor $[-] : \mathcal{Bif}(p) \rightarrow sa\mathcal{Bif}(p)$ realizing an equivalence of categories over \mathcal{C} :*

$$\begin{array}{ccc} & [-] & \\ & \curvearrowright & \\ \mathcal{Bif}(p) & & sa\mathcal{Bif}(p) \\ & \curvearrowleft & \\ & i & \\ & \Lambda_p & \Lambda_p|_{sa} \\ & \searrow & \swarrow \\ & \mathcal{C} & \end{array}$$

Proof. It is immediate from the definition of the $[-]$ map on derivations that

- (1) it respects permutation equivalence (since, by Lemma 1.9, cut does so);
- (2) it satisfies the functor laws $[\alpha \cdot \beta] \sim [\alpha] \cdot [\beta]$ and $[\text{id}_S] \sim \text{id}_{[S]}$ (since the derivations θ_S and θ_S^{-1} cancel out);
- (3) it leaves the underlying objects and arrows in \mathcal{C} unchanged;
- (4) it is invariant on derivations in the strictly alternating fragment.

Hence strictification defines a functor $[-] : \mathcal{Bif}(p) \rightarrow sa\mathcal{Bif}(p)$ that commutes with the projection functors and such that $[-] \circ i = \text{id}_{sa\mathcal{Bif}(p)}$. Finally, the family of logical equivalences $\theta_S : S \equiv [S]$ realize an isomorphism $i \circ [-] \cong \text{id}_{\mathcal{Bif}(p)}$ that is natural by construction. Indeed, it suffices to show that the square

$$\begin{array}{ccc} S & \xrightarrow{\alpha} & T \\ \theta_S \downarrow & & \downarrow \theta_T \\ i[S] & \xrightarrow{i[\alpha]} & i[T] \end{array}$$

commutes for every arrow $\alpha : S \rightarrow T$ in $\mathcal{Bif}(p)$, which is equivalent to asserting that every derivation $\alpha : S \Longrightarrow_f T$ may be factored up to permutation equivalence as

$$S \xrightarrow{\alpha} T \sim S \xrightarrow{\theta_S} [S] \xrightarrow{[\alpha]} [T] \xrightarrow{\theta_T^{-1}} T . \quad (3.1)$$

This holds by definition of the strictification map $[-]$, as we have

$$\theta_S \cdot [\alpha] \cdot \theta_T = \theta_S \cdot (\theta_S^{-1} \cdot \alpha \cdot \theta_T^{-1}) \cdot \theta_T \sim \alpha$$

□

Corollary 3.3. *The functor $\Lambda_p|_{sa} : sa\mathcal{Bif}(p) \rightarrow \mathcal{C}$ is a bifibration.*

Remark 3.4. The bifibrational structure of $\Lambda_p|_{sa}$ may be described explicitly, with the pushforward and pullback of strictly alternating formulas in $sa\mathcal{Bif}(p)$ along arrows of \mathcal{C} defined by cases depending on the polarity of the formulas, so as to ensure strict alternation. For instance, the pushforward of a strictly alternating formula $S \sqsubset A$ along an arrow $f : A \rightarrow B$ is defined as follows:

- If S is negative or atomic, then the pushforward is given by $f^+ S$ equipped with the “free” cartesian lifting $f_S : S \rightarrow f^+ S$.
- If S is positive, i.e., of the form $S = f'^+ S'$ for some $S' \sqsubset A'$ and $f' : A' \rightarrow A$, then $f^+ S$ is defined as the pushforward along the composite $f^+ S := (f \circ f')^+ S'$, equipped with the cartesian arrow given by composing $f_S : f'^+ S' \rightarrow f^+ f'^+ S'$ with the logical equivalence $f^+ f'^+ S' \equiv (f \circ f')^+ S'$.

The pullback is defined symmetrically. □

3.3.2. *Aside: splitting bifibrations.* It is worth observing that the bifibrational structure on the subcategory of strictly alternating formulas described in Remark 3.4 is *nearly* split: the pushforward and pullback functors preserve composition on the nose, but identity is only preserved up to isomorphism witnessed by the logical equivalences $\text{id}^+ U \equiv U$ and $\text{id}^- U \equiv U$. Thus the associated pseudofunctor $\mathcal{C} \rightarrow \text{Adj}$ representing the bifibration $sa\mathcal{Bif}(p) \rightarrow \mathcal{C}$ is not fully strict (it is not “normalized”), although it is a strict functor of semicategories.

It is clear that a split bifibration could be obtained, if desired, by adjusting the construction to deal explicitly with identity arrows. Indeed, it suffices to restrict further to the full subcategory of $\mathcal{Bif}(p)$ spanned by the formulas that strictly alternate along nonidentity arrows, which we denote $nsa\mathcal{Bif}(p)$. The equivalence $\mathcal{Bif}(p) \simeq nsa\mathcal{Bif}(p)$ can be established just like Theorem 3.2, applying the full suite of pseudofunctoriality laws (1.10)–(1.12) to normalize any formula to an equivalent nonidentity strictly alternating formula. The split bifibrational structure on $nsa\mathcal{Bif}(p)$ is then almost identical to what was described for $sa\mathcal{Bif}(p)$ in Remark 3.4, but taking care to never push or pull along an identity arrow. For example, when defining the pushforward of $f'^+ S'$ along f , if $f \circ f' = \text{id}$ then we choose the pushforward object to be S' . We can derive the following as a consequence.

Theorem 3.5 (cf. [28, Theorem 7.1]). *Every bifibration is equivalent to a split bifibration.*

Proof. Let $q : \mathcal{E} \rightarrow \mathcal{C}$ be a bifibration. The free bifibration $\Lambda_q : \mathcal{Bif}(q) \rightarrow \mathcal{C}$ is equipped with a canonical morphism of bifibrations $\llbracket - \rrbracket_{\mathcal{E}} : \mathcal{Bif}(q) \rightarrow \mathcal{E}$, which restricts to a functor $nsa\mathcal{Bif}(q) \rightarrow \mathcal{E}$. Now consider the *full image* of this functor. This defines a category $\llbracket nsa\mathcal{Bif}(q) \rrbracket_{\mathcal{E}}$ with same objects as $nsa\mathcal{Bif}(q)$ (that is, whose objects are nonidentity strictly alternating formulas) and whose arrows are given by arbitrary arrows $\llbracket S \rrbracket_{\mathcal{E}} \rightarrow \llbracket T \rrbracket_{\mathcal{E}}$ in \mathcal{E} . Since every object of \mathcal{E} is equal to its interpretation as an atomic formula $\llbracket X \rrbracket_{\mathcal{E}} = X$, the full and faithful functor $\llbracket nsa\mathcal{Bif}(q) \rrbracket_{\mathcal{E}} \rightarrow \mathcal{E}$ is surjective on objects and hence an equivalence of categories. Finally, the composite $q \circ \llbracket - \rrbracket_{\mathcal{E}} : \llbracket nsa\mathcal{Bif}(q) \rrbracket_{\mathcal{E}} \rightarrow \mathcal{C}$ may be equipped with the structure of a split bifibration by transferring the structure of the free split bifibration $\Lambda_q|_{nsa} : nsa\mathcal{Bif}(q) \rightarrow \mathcal{C}$ along the morphism $\llbracket - \rrbracket_{\mathcal{E}}$. □

The above proof is similar to the proof of the analogous fact for fibrations described by Anders Kock [28, Theorem 7.1] as Giraud’s “left adjoint splitting” construction, which is also analyzed (together with the “right adjoint splitting”) by Emmenegger, Mesiti, Rosolini, and Streicher [17].

In the rest of the paper we are not particularly concerned with obtaining fully split bifibrations, and so unless otherwise stated we do not forbid pushing and pulling along identity arrows.

3.3.3. A syntax of alternating formulas. We now introduce a new sequent calculus, whose purpose is to force strict alternation at the level of derivations without restricting the language of formulas. It does this by performing multiple inference steps at once, in one large step, precisely where the $[-]$ translation would force them to be grouped together.

Strictly alternating formulas $[S]$ admit a simple inductive characterization: a formula is strictly alternating just in case it is either atomic X , or the pushforward $f^+[N]$ of a non-positive strictly alternating formula, or the pullback $g^- [P]$ of a non-negative strictly alternating formula.

$$\begin{aligned} [S] &::= [P] \mid [N] \\ [P] &::= X \mid f^+[N] \\ [N] &::= X \mid g^- [P] \end{aligned}$$

Here we use “non-positive” and “non-negative” in the expected sense, to mean that N is either atomic or a pullback formula, and P is either atomic or a pushforward formula. Note that our definition of strictly alternating allows for pushing and pulling along identity arrows (cf. §3.3.2).

In fact, *any* bifibrational formula can be described in a similar alternating way, if instead of pushing or pulling applying a single arrow f or g , we push or pull along a *sequence* π , where π denotes a non-empty sequence of composable arrows (f_0, \dots, f_n) . We introduce an explicit strictly-alternating grammar for arbitrary formulas S_{sa} , non-positive formulas N_{sa} and non-negative formulas P_{sa} .

$$\begin{aligned} S_{sa} &::= P_{sa} \mid N_{sa} \\ P_{sa} &::= X \mid \pi^+ N_{sa} \\ N_{sa} &::= X \mid \pi^- P_{sa} \end{aligned}$$

Formulas described in this way can be interpreted in two different ways, depending on the interpretation of sequences $\pi = (f_0, \dots, f_n)$:

- (1) If we interpret π^+ as an *iterated* push $f_n^+ \cdots f_0^+$, a formula S_{sa} in alternating syntax is interpreted as an arbitrary formula (not strictly alternating) that we denote $[S_{sa}]$.
- (2) If we interpret π^+ as a push along the composite arrow $[\pi] := f_0 \cdots f_n$, a formula S_{sa} is interpreted as a strictly alternating formula that we denote $[S_{sa}]$.

For example if we take $N_{sa} = (f, g)^+ (h, i)^- X$, we have $[N_{sa}] = g^+ f^+ h^- i^- X$ and $[N_{sa}] = (f, g)^+ (h, i)^- X$. All bifibrational formulas are in the image of $[-]$, all strictly-alternating bifibrational formulas are in the image of $[-]$, and for any S_{sa} we have $[S_{sa}] = [[S_{sa}]]$.

To keep notations lighter, we write S to mean either a bifibrational formula S or its alternating presentation (the unique S_{sa} determined by $[S_{sa}] = S$).

$$\begin{aligned}
\pi &::= (f_0, \dots, f_n) \\
\alpha_w &::= \delta \mid \bar{\pi}.\alpha_w \mid \pi \setminus_g \alpha_w \mid \alpha_w.\pi \mid \alpha_w f / \bar{\pi}
\end{aligned}$$

$$\frac{\delta : X \rightarrow Y \in \mathcal{D} \quad p(\delta) = f}{\delta : X \xRightarrow[f]{\quad} Y} \delta$$

$$\begin{array}{cc}
\frac{\alpha_w : N \xRightarrow[\pi g]{\quad} T}{\pi \setminus_g \alpha_w : \pi^+ N \xRightarrow[g]{\quad} T} \mathbf{L}\pi^+ & \frac{\alpha_w : S \xRightarrow[f]{\quad} N}{\alpha_w.\pi : S \xRightarrow[f\pi]{\quad} \pi^+ N} \mathbf{R}\pi^+ \\
\frac{\alpha_w : P \xRightarrow[g]{\quad} T}{\bar{\pi}.\alpha_w : \pi^- P \xRightarrow[\pi g]{\quad} T} \mathbf{L}\pi^- & \frac{\alpha_w : S \xRightarrow[f\pi]{\quad} P}{\alpha_w f / \bar{\pi} : S \xRightarrow[f]{\quad} \pi^- P} \mathbf{R}\pi^-
\end{array}$$

Figure 3: Rules of the weakly focused calculus.

3.3.4. *Weak focusing as a syntax of alternating derivations.* We now introduce a weakly focused calculus as a sequent calculus on formulas in the alternating syntax. We present the inference rules in Figure 3 at the same time as a more compact term syntax for derivations, similar to one that we introduced in §1.3.

Just as formulas S_{sa} in alternating syntax admit two natural interpretations, each derivation $\alpha_w : S_{sa} \xRightarrow{f} T_{sa}$ in the weakly focused calculus admits two additional interpretations besides its formal interpretation as a sequence of applications of the above rules:

- (1) As a derivation $[\alpha_w] : [S_{sa}] \xRightarrow{f} [T_{sa}]$ in the unrestricted sequent calculus, by expanding the “big-step” rules $\mathbf{L}\pi^+$, $\mathbf{R}\pi^+$, $\mathbf{L}\pi^-$, and $\mathbf{R}\pi^-$ appropriately in terms of the “small-step” rules $\mathbf{L}f^+$, $\mathbf{R}f^+$, $\mathbf{L}g^-$, and $\mathbf{R}g^-$. This corresponds to translating $\pi \setminus_g \alpha_w$ to $[\pi \setminus_g \alpha_w] = \pi \setminus_g [\alpha_w]$, where $\pi \setminus_g \alpha$ is defined as the iterated division $f_n \setminus_g \dots f_0 \setminus_{f_1 \dots f_n g} \alpha$ translating $\alpha_w.\pi$ to $[\alpha_w.\pi] = [\alpha_w].\pi$, where $\alpha.\pi$ is defined as $\alpha.f_0 \dots f_n$, and so on.
- (2) As a derivation $[\alpha_w] : [S_{sa}] \xRightarrow{f} [T_{sa}]$ in the strictly alternating fragment of the unrestricted sequent calculus, by interpreting each instance of $\mathbf{L}\pi^+$, $\mathbf{R}\pi^+$, etc., as a single instance of $\mathbf{L}f^+$, $\mathbf{R}f^+$, etc., on the composite arrow $[\pi] = f_0 \dots f_n$. This corresponds to translating $\pi \setminus_g \alpha_w$ to a single left division $[\pi \setminus_g \alpha_w] = [\pi] \setminus_g [\alpha_w]$, translating $\alpha_w.\pi$ to a single right multiplication $[\alpha_w.\pi] = [\alpha_w].[\pi]$, and so on.

When we remain ambiguous about whether S denotes an arbitrary bifibrational formula or its alternating syntax, we can write that for any $\alpha_w : S \xRightarrow{f} T$, we have $[\alpha_w] : S \xRightarrow{f} T$ and $[\alpha_w] : [S] \xRightarrow{f} [T]$.

3.3.5. *Completeness of weak focusing.* We will use the following properties of iterated divisions and iterated multiplications on unrestricted derivations.

Lemma 3.6 (Iterated identity). $\pi \setminus (\text{id}_S.\pi) \sim \text{id}_{\pi^+ S}$ and $(\bar{\pi}.\text{id}_T) / \bar{\pi} \sim \text{id}_{\pi^- T}$ for all compatible formulas S, T and sequences of arrows π .

Lemma 3.7 (Iterated cut). $(\alpha.\pi) \cdot (\pi \setminus_h \beta) \sim \alpha \cdot \beta$ and $(\alpha g / \bar{\pi}) \cdot (\bar{\pi}.\beta) \sim \alpha \cdot \beta$ for all compatible derivations α, β and sequences of arrows π .

Proof. Both straightforward by induction on π . □

We can now give explicit derivations for the equivalences $\theta_S : S \equiv [S]$, by defining for every $S \sqsubset A$ a pair of (unrestricted) derivations

$$\theta_S : S \xRightarrow{\text{id}_A} [S] \quad \theta_S^{-1} : [S] \xRightarrow{\text{id}_A} S$$

as follows. (We omit the subscript arrow on the division signs below, which is always id_A .)

$$\begin{aligned} \theta_{\pi^+ N} &= \pi \setminus (\theta_N \cdot [\pi]) & \theta_{\pi^+ N}^{-1} &= [\pi] \setminus (\theta_N^{-1} \cdot \pi) \\ \theta_{\pi^- P} &= (\overline{[\pi]} \cdot \theta_P) / \overline{\pi} & \theta_{\pi^- P}^{-1} &= (\overline{\pi} \cdot \theta_P^{-1}) / \overline{[\pi]} \\ \theta_X &= \theta_X^{-1} = \text{id}_X \end{aligned}$$

By appealing to Lemmas 3.6 and 3.7, we have:

Proposition 3.8. θ_S and θ_S^{-1} are inverses up to permutation equivalence.

Proof. For example in the case $S = \pi^+ N$ we have:

$$\begin{aligned} \theta_{\pi^+ N} \cdot \theta_{\pi^+ N}^{-1} &= (\pi \setminus (\theta_N \cdot [\pi])) \cdot ([\pi] \setminus (\theta_N^{-1} \cdot \pi)) \\ &\sim \pi \setminus (\theta_N \cdot \theta_N^{-1} \cdot \pi) && \text{definition of cut} \\ &\sim \pi \setminus (\text{id}_N \cdot \pi) && \text{induction hypothesis} \\ &\sim \text{id}_{\pi^+ N} && \text{Lemma 3.6} \end{aligned}$$

$$\begin{aligned} \theta_{\pi^+ N}^{-1} \cdot \theta_{\pi^+ N} &= ([\pi] \setminus (\theta_N^{-1} \cdot \pi)) \cdot (\pi \setminus (\theta_N \cdot [\pi])) \\ &\sim [\pi] \setminus (\theta_N^{-1} \cdot \theta_N \cdot [\pi]) && \text{definition of cut and Lemma 3.7} \\ &\sim [\pi] \setminus (\text{id}_N \cdot [\pi]) && \text{induction hypothesis} \\ &\sim \text{id}_{[\pi]^+ N} && \text{definition of identity} \end{aligned}$$

□

On formulas we remarked that we have the relation $[S] = [[S]]$. There is a similar relation at the level of derivations.

Proposition 3.9. For any $\alpha_w : S \xRightarrow{f} T$ we have $[\alpha_w] \sim [[\alpha_w]]$.

Proof. By induction on the weakly focused derivation α_w . For example, in the case that α_w is of the form $\pi \setminus_g \alpha'_w : \pi^+ N \xRightarrow{g} T$, we have

$$\begin{aligned} [[\pi \setminus_g \alpha'_w]] &= ([\pi] \setminus (\theta_N^{-1} \cdot \pi)) \cdot (\pi \setminus_g [\alpha'_w]) \cdot \theta_T && \text{by definition of } [-] \text{ and } [-] \\ &\sim [\pi] \setminus (\theta_N^{-1} \cdot [\alpha_w]) \cdot \theta_T && \text{definition of cut and Lemma 3.7} \\ &\sim [\pi] \setminus [\alpha_w] && \text{induction hypothesis} \\ &= [\pi \setminus_g \alpha_w] && \text{definition of } [-] \end{aligned}$$

□

We remarked above that every bifibrational formula is of the form $[S_{sa}]$, and every strictly alternating formula is of the form $[S_{sa}]$, where S_{sa} is a formula in alternating syntax. The situation for derivations is different: all strictly alternating derivations are of the form $[\alpha_w]$, but *not* all unrestricted derivations are of the form $[\alpha_w]$ —it is really a more restricted

syntax. For example, the derivation on the left below is not in the image of $[-]$, but it is permutation equivalent to the derivation on the right, which is:

$$\begin{array}{c} \frac{S \Longrightarrow T}{fghi} \\ \frac{f^+ S \Longrightarrow T}{ghi} \text{Lf}^+ \\ \frac{f^+ S \Longrightarrow i^- T}{gh} \text{Ri}^- \\ \frac{g^+ f^+ S \Longrightarrow i^- T}{h} \text{Lg}^+ \end{array} \qquad \begin{array}{c} \frac{S \Longrightarrow T}{fghi} \\ \frac{S \Longrightarrow i^- T}{fgh} \text{Ri}^- \\ \frac{f^+ S \Longrightarrow i^- T}{gh} \text{Lf}^+ \\ \frac{g^+ f^+ S \Longrightarrow i^- T}{h} \text{Lg}^+ \end{array}$$

Proposition 3.10. *The mapping $[-]$ is a one-to-one correspondence between weakly focused derivations of $S_{sa} \Longrightarrow_f T_{sa}$ and unrestricted derivations of strictly alternating sequents $[S_{sa}] \Longrightarrow_f [T_{sa}]$.*

Proof. We prove that for any judgment $S_{sa} \Longrightarrow_f T_{sa}$, if α is an unrestricted derivation of $[S_{sa}] \Longrightarrow_f [T_{sa}]$, then there exists a *unique* weakly focused derivation α_w such that $\alpha = [\alpha_w]$. The proof proceeds by induction on the product ordering of the two sides of the sequent (S_{sa}, T_{sa}) with respect to the subformula ordering. We will have more occasions to use this style of induction in the rest of this section, which we refer to as *induction on judgments*. The associated well-founded ordering on judgments is defined explicitly by

$$(S' \xRightarrow{h'} T') < (S \xRightarrow{h} T)$$

just in case S' is a subformula of S , T' is a subformula of T , and either S' is a strict subformula of S or T' is a strict subformula of T . We now proceed to the proof.

If α is an axiom rule then it is already a weakly focused proof.

If α is of the form $f \setminus_g \alpha'$:

$$\frac{\alpha' : S' \xRightarrow{fg} [T_{sa}]}{f \setminus_g \alpha' : f^+ S' \xRightarrow{g} [T_{sa}]}$$

then from $f^+ S' = [S]_{sa}$ we can deduce that S_{sa} must be of the form $\pi^+ N$, so f is uniquely determined to be $[\pi]$ and S' is $[N]$. So in fact the derivation of α is of the form

$$\frac{\alpha' : [N] \xRightarrow{\pi g} [T_{sa}]}{[\pi] \setminus_g \alpha' : [\pi]^+ [N] \xRightarrow{g} [T_{sa}]}$$

By induction hypothesis, there exists a unique α'_w such that $\alpha' = [\alpha'_w]$, so α is necessarily equal to $[\pi] \setminus_g [\alpha'_w]$, and $\alpha_w := \pi \setminus_g \alpha'_w$ is uniquely determined.

The reasoning is symmetric when if α is of the form $\bar{f} \cdot \alpha'$. □

Corollary 3.11 (Completeness of weak focusing). *For any derivation $\alpha : S \Longrightarrow_f T$, there is a weakly focused derivation $\alpha_w : S \Longrightarrow_f T$ such that $\alpha \sim [\alpha_w]$.*

Proof. Let $[\alpha] : [S] \Longrightarrow_f [T]$ be the strictification of α , and let $\alpha_w : S \Longrightarrow_f T$ be the unique weakly focused derivation such that $[\alpha] = [\alpha_w]$, which exists by the previous proposition. In particular, $[\alpha]$ and $[[\alpha_w]]$ are equal as $sa\text{Bif}(p)$ morphisms, so by Theorem 3.2 we have that α and $[\alpha_w]$ are equal as $\text{Bif}(p)$ morphisms, that is, $\alpha \sim [\alpha_w]$. □

3.3.6. *Equivalence of weakly focused proofs.* There are two reasonable notions of equivalence on weakly focused proofs α_w, β_w :

- One can consider the relation $[\alpha_w] \sim [\beta_w]$, which views α_w, β_w as unrestricted derivations and considers permutation equivalence between them. Such a proof of equivalence permutes inference rules one at a time, and in particular it may contain intermediate steps that are not weakly focused derivations.
- One can consider the relation $[\alpha_w] \sim [\beta_w]$, which views α_w and β_w as derivations on strictly alternating formulas. Such proofs of equivalence respect the strictly-alternating structure in the sense that they cannot break sequences of pushes or sequence of pulls apart. They can be seen as permutation equivalences on “synthetic” connectives.

Corollary 3.12. *These two notions of equivalence coincide.*

Proof.

$$[\alpha_w] \sim [\beta_w] \quad \xleftrightarrow{\text{Proposition 3.9}} \quad [[\alpha_w]] \sim [[\beta_w]] \quad \xleftrightarrow{\text{Theorem 3.2}} \quad [\alpha_w] \sim [\beta_w] \quad \square$$

3.3.7. *Presentation.* The results of this section can be reformulated as an alternative presentation of the arrows of $\mathcal{Bif}(p)$ as triples $(S, \hat{\alpha}_w, T)$, where $\hat{\alpha}_w \in sa\mathcal{Bif}(p)(S, T)$ is a permutation-equivalence class of weakly focused derivations $\alpha_w : S \Longrightarrow_f T$. Weak focusing removes some of the inessential non-determinism of the sequent calculus, but it is possible to go further and obtain a still more canonical presentation, as we will now examine.

3.4. Multifocusing.

3.4.1. *Introduction.* The previous calculus was only “weakly” focused in the sense that it allows applying non-invertible rules to derive a judgment even where an invertible rule is applicable. It would be more canonical to enforce that invertible rules are applied eagerly. For instance, if we have a positive formula on the left then we can always apply an invertible left rule, and if we have a negative formula on the right then we can always apply an invertible right rule. In the case of a sequent of the form $\pi^+ N \Longrightarrow_f \rho^- P$ with *both* a positive formula on the left and a negative formula, it seems we have a useless choice to make of whether to work on the left or right first:

$$\frac{\frac{N \Longrightarrow P}{\pi^+ N \Longrightarrow P} \text{L}\pi^+}{\pi^+ N \Longrightarrow_f \rho^- P} \text{R}\rho^- \qquad \frac{\frac{N \Longrightarrow P}{N \Longrightarrow \rho^- P} \text{R}\rho^-}{\pi^+ N \Longrightarrow_f \rho^- P} \text{L}\pi^+$$

The choice of whether to apply $\text{R}\rho^-$ or $\text{L}\pi^+$ first is irrelevant, since the two derivations above are permutation equivalent. However, a natural way of addressing this issue is to add a rule that can perform *both* left and right inversion at once. By combining such a bi-inversion rule with polarity-restricted instances of the left- and right-inversion rules,

$$\frac{N \Longrightarrow_{\pi^+} P}{\pi^+ N \Longrightarrow_f P} \text{L}\pi^+ \qquad \frac{N \Longrightarrow_{\rho^-} P}{\pi^+ N \Longrightarrow_f \rho^- P} \text{L}\pi^+ \text{R}\rho^- \qquad \frac{N \Longrightarrow_{\rho^-} P}{N \Longrightarrow_f \rho^- P} \text{R}\rho^- \qquad (3.2)$$

$$\begin{array}{c}
\alpha_m ::= \delta \\
\left| \begin{array}{c|c|c} \bar{\sigma}.\alpha_m & \alpha_m.\tau & \bar{\sigma}.\alpha_m.\tau \\ \pi \setminus_g \alpha_m & \alpha_m \ g / \bar{\rho} & \pi \setminus \alpha_m / \bar{\rho} \end{array} \right. \\
\\
\frac{\delta : X \rightarrow Y \in \mathcal{D} \quad p(\delta) = f}{\delta : X \Longrightarrow_f Y} \delta \\
\\
\frac{\alpha_m : N \Longrightarrow_{\pi f} P}{\pi \setminus_f \alpha_m : \pi^+ N \Longrightarrow_f P} \mathsf{L}\pi^+ \quad \frac{\alpha_m : N \Longrightarrow_{\pi f \rho} P}{\pi \setminus \alpha_m / \bar{\rho} : \pi^+ N \Longrightarrow_f \rho^- P} \mathsf{L}\pi^+ \mathsf{R}\rho^- \quad \frac{\alpha_m : N \Longrightarrow_{f \rho} P}{\alpha_m \ f / \bar{\rho} : N \Longrightarrow_f \rho^- P} \mathsf{R}\rho^- \\
\\
\frac{\alpha_m : P \Longrightarrow_f Q}{\bar{\sigma}.\alpha_m : \sigma^- P \Longrightarrow_{\sigma f} Q} \mathsf{L}\sigma^- \quad \frac{\alpha_m : P \Longrightarrow_f N}{\bar{\sigma}.\alpha_m.\tau : \sigma^- P \Longrightarrow_{\sigma f \tau} \tau^+ N} \mathsf{L}\sigma^- \mathsf{R}\tau^+ \quad \frac{\alpha_m : N \Longrightarrow_f M}{\alpha_m.\tau : N \Longrightarrow_{f \tau} \tau^+ M} \mathsf{R}\tau^+
\end{array}$$

Figure 4: Rules of the strongly multifocused calculus.

we can enforce that non-invertible rules are only used on polarized sequents of the form $N \Longrightarrow_f P$, with a uniquely determined invertible rule applicable in all other cases. We say that such a system, with a polarity restriction on non-invertible phases, is *strongly focused*.

Once we arrive at a sequent of the form $N \Longrightarrow_f P$ we have to make a potentially significant choice as to whether to try applying a non-invertible rule on the left or right first. However, here a similar phenomenon arises in considering permutable applications of non-invertible rules. Again it is natural to introduce a rule that allows to perform two non-invertible rules in parallel, written in the middle below:

$$\frac{P \Longrightarrow_f Q}{\sigma^- P \Longrightarrow_{\sigma f} Q} \mathsf{L}\sigma^- \quad \frac{P \Longrightarrow_f N}{\sigma^- P \Longrightarrow_{\sigma f \tau} \tau^+ N} \mathsf{L}\sigma^- \mathsf{R}\tau^+ \quad \frac{N \Longrightarrow_f M}{N \Longrightarrow_{f \tau} \tau^+ M} \mathsf{R}\tau^+ \quad (3.3)$$

The collection of six polarized rules (3.2) and (3.3), together with the ordinary axiom rule, together define a system that we call *strongly multifocused*.

Strong focusing and multi-focusing are standard notions from the literature [12, 51]. Going from a weakly-focused to a strongly-focused system (enforcing that invertible rules are applied eagerly) makes the system more canonical. On the other hand, introducing multifocusing (the bi-focusing rule $\mathsf{L}\sigma^- \mathsf{R}\tau^+$ in our case) does not automatically make the system more canonical, since it actually allows more proofs: every time a bi-focused rule is used, it is also possible to use two mono-focused rules in sequence. In Section 3.6 below, we will explain how to restrict this system by limiting the usage of mono-focused rules, resulting in a calculus that is canonical under an assumption on the base category.

In Figure 4 we recapitulate the rules of the strongly multifocused system while introducing a term syntax.

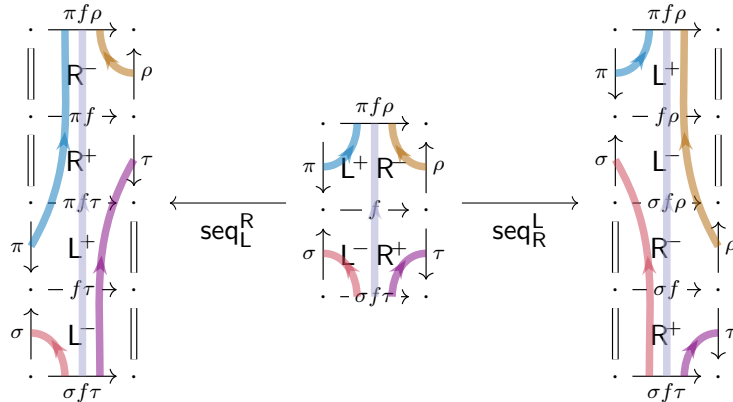
3.4.2. Neutral sequents, bipoles. Following the focusing tradition, we call *neutral* a sequent of the form $N \Longrightarrow_f P$. A proof of a (non-atomic) neutral sequent has to start with a focusing step, either on the left, right or both sides of the sequent.

Following the same tradition, we call *bipole* a two-step (open) derivation of a neutral sequent beginning with a focusing step followed by an inversion step. We use the notation \mathbb{L} for a bipole of $L^+ L^-$ rules, \mathbb{R} for a bipole of $R^- R^+$ rules, and \mathbb{LR} for a bi-focused bipole (L^+R^-) (L^-R^+). In order to avoid having to treat atomic formulas as a special case, we will expand the interpretation of the $L\pi^+$, $R\rho^-$, and $L\pi^+R\rho^-$ rules by allowing π (respectively, ρ) to be empty when the left (respectively right) side of the sequent is atomic. With that convention, every strongly focused proof of a neutral sequent decomposes as a sequence of $\mathbb{L} / \mathbb{R} / \mathbb{LR}$ bipoles ended by an axiom.

3.4.3. *Sequentialization equivalence of strongly multi-focused proofs.* The permutation equivalence that we introduced in Section 1.5 and extended to weakly focused proofs is not directly applicable for strongly focused proofs. Indeed, permutation equivalence breaks the polarity invariants that we described above. However, there is a natural notion of equivalence for strongly multifocused derivations that arises from considering both possible *sequentializations* of a \mathbb{LR} bipole:

$$\begin{array}{ccc}
 & \frac{N \Longrightarrow P}{\frac{\pi f \rho}{\pi^+ N \Longrightarrow \rho^- P} L\pi^+ R\rho^-} L\sigma^- R\tau^+ & \\
 & \frac{\frac{\frac{\frac{N \Longrightarrow P}{\pi f \rho} R\rho^-}{N \Longrightarrow \rho^- P} \pi f}{N \Longrightarrow \tau^+ \rho^- P} R\tau^+}{\pi^+ N \Longrightarrow \tau^+ \rho^- P} L\pi^+}{\sigma^- \pi^+ N \Longrightarrow \tau^+ \rho^- P} L\sigma^- & \\
 \swarrow \text{seq}_L^R & & \searrow \text{seq}_R^L \\
 \frac{N \Longrightarrow P}{\frac{\pi f \rho}{N \Longrightarrow \rho^- P} R\rho^-} R\tau^+ & & \frac{N \Longrightarrow P}{\frac{\pi f \rho}{\pi^+ N \Longrightarrow P} L\pi^+} L\sigma^- \\
 \frac{\frac{\frac{N \Longrightarrow P}{\pi f \rho} R\rho^-}{N \Longrightarrow \rho^- P} \pi f}{N \Longrightarrow \tau^+ \rho^- P} R\tau^+}{\pi^+ N \Longrightarrow \tau^+ \rho^- P} L\pi^+ & & \frac{\frac{\frac{N \Longrightarrow P}{\pi f \rho} L\pi^+}{\pi^+ N \Longrightarrow P} f \rho}{\sigma^- \pi^+ N \Longrightarrow P} L\sigma^- \\
 \frac{\frac{\frac{N \Longrightarrow P}{\pi f \rho} R\rho^-}{N \Longrightarrow \rho^- P} \pi f}{N \Longrightarrow \tau^+ \rho^- P} R\tau^+}{\pi^+ N \Longrightarrow \tau^+ \rho^- P} L\pi^+}{\sigma^- \pi^+ N \Longrightarrow \tau^+ \rho^- P} L\sigma^- & & \frac{\frac{\frac{\frac{N \Longrightarrow P}{\pi f \rho} L\pi^+}{\pi^+ N \Longrightarrow P} f \rho}{\sigma^- \pi^+ N \Longrightarrow P} L\sigma^-}{\sigma^- \pi^+ N \Longrightarrow \rho^- P} R\rho^-}{\sigma^- \pi^+ N \Longrightarrow \tau^+ \rho^- P} R\tau^+
 \end{array}$$

Or equivalently in stack-and-string diagrammatic form:



Or equivalently again in proof term syntax:

$$\bar{\sigma}.\left(\pi \setminus_{f\tau} \left(\left(\beta \ \pi_f / \bar{\rho}\right).\tau\right)\right) \xleftarrow{\text{seq}_L^R} \bar{\sigma}.\left(\pi \setminus_f \beta / \bar{\rho}\right).\tau \xrightarrow{\text{seq}_R^L} \left(\left(\bar{\sigma}.\left(\pi \setminus_{f\rho} \beta\right)\right)_{\sigma_f} / \bar{\rho}\right).\tau \quad (3.4)$$

Recall that we allow π and ρ to be empty, $\pi = \epsilon$ or $\rho = \epsilon$ (in which case, e.g., $\epsilon \setminus \alpha = \alpha$).

Let us write $(\rightarrow_{\text{seq}})$ for this sequentialization order, and (\sim_{seq}) for the corresponding equivalence relation.

3.4.4. Relating the weak-monofocused and strong-multifocused systems.

The $(\rightarrow_{\text{seq}})$ -normal forms of this relation are the strongly *monofocused* proofs—in other words proofs that do not use the bi-focusing rule, although they can use bi-inversion. Such proofs can then be turned into derivations in the weakly focused calculus by selecting an ordering for each bi-inversion rule. We define $\text{Seq}(\alpha_m)$ as the set of weakly monofocused sequentializations of a strongly multifocused proof α_m .

Proposition 3.13. *If $\alpha_m \sim_{\text{seq}} \alpha'_m$ then for all $\alpha_w \in \text{Seq}(\alpha_m)$ and $\alpha'_w \in \text{Seq}(\alpha'_m)$ we have $\alpha_w \sim \alpha'_w$.*

Proof. The idea of the proof is that for any non-deterministic choice in the sequentialization process, the two choices will sequentialize into permutation-equivalent derivations.

The first choice comes from the definition of the sequentialization order between multifocused proofs. For each α with two direct sequentializations α_1, α_2 as depicted in (3.4) above, we have that $\text{Seq}(\alpha) = \text{Seq}(\alpha_1) \uplus \text{Seq}(\alpha_2)$, and we prove that for any $\alpha_{1w} \in \text{Seq}(\alpha_1)$ and $\alpha_{2w} \in \text{Seq}(\alpha_2)$ we have $\alpha_{1w} \sim \alpha_{2w}$. Indeed, the two sequentializations sets are

$$\begin{aligned} \text{Seq}(\alpha_1) &= \{ \bar{\sigma}.\left(\pi \setminus \left(\left(\beta_w \ \rho_f / \bar{\rho}\right).\tau\right)\right) \mid \beta_w \in \text{Seq}(\beta) \} \\ \text{Seq}(\alpha_2) &= \{ \left(\left(\bar{\sigma}.\left(\pi \setminus_f \beta_w\right)\right) / \bar{\rho}\right).\tau \mid \beta_w \in \text{Seq}(\beta) \} \end{aligned}$$

and for any $\beta_w \in \text{Seq}(\beta)$, we do have as expected:

$$\begin{aligned} &\bar{\sigma}.\left(\pi \setminus \left(\left(\beta_w \ \rho_f / \bar{\rho}\right).\tau\right)\right) \\ &\sim \bar{\sigma}.\left(\left(\pi \setminus \left(\beta_w \ \rho_f / \bar{\rho}\right)\right).\tau\right) \\ &\sim \left(\bar{\sigma}.\left(\pi \setminus \left(\beta_w \ \rho_f / \bar{\rho}\right)\right)\right).\tau \\ &\sim \left(\bar{\sigma}.\left(\left(\pi \setminus_f \beta_w\right) / \bar{\rho}\right)\right).\tau \\ &\sim \left(\left(\bar{\sigma}.\left(\pi \setminus_f \beta_w\right)\right) / \bar{\rho}\right).\tau \end{aligned}$$

The second choice is in the ordering of inversion rules when sequentializing bi-inversion rules, and the two choices are immediately permutation-equivalent. \square

Corollary 3.14. *All elements of $\text{Seq}(\alpha_m)$ are permutation-equivalent: for any α_m we have $\alpha_m \sim_{\text{seq}} \alpha_m$, so $\alpha_w \in \text{Seq}(\alpha_m)$ and $\alpha'_w \in \text{Seq}(\alpha_m)$ imply $\alpha_w \sim \alpha'_w$.*

Conversely, given a weakly focused derivation $\alpha_w : S \Longrightarrow_h T$, let us write $\mathbf{S}(\alpha_w)$ for the *strengthening* of α_w , a uniquely-determined strongly (mono-)focused proof defined as follows, by induction on judgments:

$$\begin{aligned} \mathbf{S}(\alpha : \pi^+ N \Longrightarrow_f \rho^- P) &:= \pi \setminus \mathbf{S}(\pi_N \cdot \alpha \cdot \bar{\rho}_P : N \Longrightarrow_{\pi_f \rho} P) / \bar{\rho} \\ \mathbf{S}((\bar{\sigma}.\alpha) : \sigma^- P \Longrightarrow_{\sigma_f} Q) &:= \bar{\sigma}.\mathbf{S}(\alpha : P \Longrightarrow_f Q) \\ \mathbf{S}((\alpha.\tau) : N \Longrightarrow_{f\tau} \tau^+ M) &:= \mathbf{S}(\alpha : N \Longrightarrow_f M).\tau \\ \mathbf{S}(\delta : X \Longrightarrow_f Y) &:= \delta \end{aligned}$$

In other words, to strengthen α_w , whenever we encounter a judgment that admits invertible steps, we do them (both at the same time if applicable), and otherwise we imitate the non-invertible steps.

This strengthening transformation defines a mapping from weakly focused derivations to strongly focused derivations of the same judgment. We now establish that it sends permutation equivalent weakly focused derivations to sequentially equivalent strongly focused derivations.

Proposition 3.15. *If $\alpha_w \sim \alpha'_w$ then $S(\alpha_w) \sim_{\text{seq}} S(\alpha'_w)$.*

Proof. By induction on the underlying judgment derived by both derivations $\alpha_w \sim \alpha'_w$. There are two cases to consider, depending on whether or not the judgment is invertible.

- Case $\pi^+ N \Longrightarrow_f \rho^- P$. (Recall we allow $\pi = \epsilon$ or $\rho = \epsilon$.) We have

$$S(\alpha_w) = \pi \setminus S(\pi_N \cdot \alpha_w \cdot \bar{\rho}_P) / \bar{\rho} \sim_{\text{seq}} \pi \setminus S(\pi_N \cdot \alpha'_w \cdot \bar{\rho}_P) / \bar{\rho} = S(\alpha'_w)$$

where the first and last steps are by definition and the middle step is by induction on $\pi_N \cdot \alpha_w \cdot \bar{\rho}_P \sim \pi_N \cdot \alpha'_w \cdot \bar{\rho}_P$ at the smaller judgment $N \Longrightarrow_{\pi f \rho} P$.

- Case $N \Longrightarrow_h P$. We consider the derivation of the permutation equivalence $\alpha_w \sim \alpha'_w$.
 - Subcase $\alpha_w \sim \alpha'_w$ is derived by the generator (1.2), with $\alpha_w = (\bar{f}.\beta).h$ and $\alpha'_w = \bar{f}.(\beta.h)$. In this case the two strengthenings are sequentializations of the same bi-focused derivation $\bar{f}.S(\beta).h$.
 - Subcase $\alpha_w \sim \alpha'_w$ is derived by a generator (1.3)–(1.5): impossible at the judgment $N \Longrightarrow_h P$.
 - Subcase $\alpha_w = \bar{f}.\beta$ and $\alpha'_w = \bar{f}.\beta'$ for some $\beta \sim \beta' : Q \Longrightarrow_{fh} P$ where $N = f^- Q$. Then $S(\beta) \sim_{\text{seq}} S(\beta')$ by induction and so $S(\alpha_w) = \bar{f}.S(\beta) \sim_{\text{seq}} \bar{f}.S(\beta') = S(\alpha'_w)$.
 - Subcase $\alpha_w = \beta.g$ and $\alpha'_w = \beta'.g$ for some $\beta \sim \beta' : N \Longrightarrow_{hg} M$ where $P = g^+ M$: same as above.

□

The relation between permutation equivalence and sequentialization equivalence is captured by the following diagram:

$$\begin{array}{ccc}
 & \text{S} & \\
 \text{(WeakMono}(S, f, T), \sim) & \xrightarrow{\quad} & \text{(StrongMono}(S, f, T), \sim_{\text{seq}}) \sim \text{(StrongMulti}(S, f, T), \sim_{\text{seq}}) \\
 & \xleftarrow{\quad} & \\
 & \text{InvSeq} & \text{FocSeq}
 \end{array}$$

where we have factored sequentialization $\text{Seq}(-)$ in two different nondeterministic / multivalued functions: $\text{FocSeq}(\alpha_m)$ is the sequentialization of the bi-focused rules of a multifocused proof, and $\text{InvSeq}(\alpha_m)$ is the sequentialization of the bi-invertible rules of a strongly mono-focused proof. This diagram expresses two equivalences of setoids (sets equipped with an equivalence relation), in the following sense:

- (1) S is a morphism of setoids, per Proposition 3.15. (This gives us completeness and relational completeness of strongly focused proof with respect to weakly focused proofs).
- (2) InvSeq and FocSeq are nondeterministic morphisms of setoids, per Proposition 3.13.
- (3) The equivalence on the right corresponds to the following statement:

$$\begin{array}{c}
 \alpha_m \sim \text{FocSeq}(\alpha_m) \\
 \forall \alpha'_m \in \text{FocSeq}(\alpha_m), \alpha'_m \sim_{\text{seq}} \alpha_m
 \end{array}$$

It is immediate from the fact that $\text{FocSeq}(\alpha_m)$ is defined as the set of $(\rightarrow_{\text{seq}})$ -normal forms of α_m , and (\sim_{seq}) as the least congruence containing $(\rightarrow_{\text{seq}})$.

(4) The equivalence on the left corresponds to the following statements:

Proposition 3.16.

$$\begin{aligned} & \forall \alpha_w, \quad \alpha_w \sim \text{InvSeq}(\mathbf{S}(\alpha_w)) \\ \forall \alpha_w, \forall \alpha'_w \in \text{InvSeq}(\mathbf{S}(\alpha_w)), \quad & \alpha_w \sim \alpha'_w \end{aligned}$$

Proposition 3.17.

$$\begin{aligned} & \forall \alpha_m, \quad \alpha_m \sim_{\text{seq}} \mathbf{S}(\text{InvSeq}(\alpha_m)) \\ \forall \alpha_m, \forall \alpha_w \in \text{InvSeq}(\alpha_m), \quad & \alpha_m \sim_{\text{seq}} \mathbf{S}(\alpha_w) \end{aligned}$$

For each property, on the first line we wrote the statement with InvSeq seen as a multi-valued function nondeterministically returning a weakly focused proof, and on the second line we wrote it with InvSeq seen directly as a function into sets of weakly focused proofs. Both statements can be established by induction on derivations; we detail the proofs in Appendix A.2.

The intermediate system **StrongMono** is useful to describe the fine-grained properties of strengthening and sequentialization, but note that the equivalence relation (\sim_{seq}) has no natural definition that is internal to **StrongMono**, since the source derivation α_m of a sequentialization step $\alpha_m \rightarrow_{\text{seq}} \alpha'_m$ uses a bi-focused rule.

3.4.5. *Horizontal composition (cut) of strongly multi-focused proofs.* Now that we have defined strengthening, sequentialization and their properties, we can easily define horizontal composition (cut) of strongly multifocused proofs modulo permutation equivalence: $\alpha_m \cdot \beta_m$ can be defined as $\mathbf{S}(\alpha_w \cdot \beta_w)$ for some $\alpha_w \in \text{Seq}(\alpha_m)$, $\beta_w \in \text{Seq}(\beta_m)$. The choice of α_w and β_w is irrelevant: they are all permutation equivalent, so the resulting strengthening is unique modulo sequentialization equivalence. This definition is easily seen to be associative (lifted from the associativity of composition of non-focused proofs) and unital—for identity derivations defined as the $\mathbf{S}(\text{id}_A)$. Finally, we also directly get that $\mathbf{S}(-)$ respects composition and preserves identity derivations.

We can thus define a category of strongly multifocused derivations quotiented by (\sim_{seq}), with $\mathbf{S}(-)$ a functor from $\mathcal{Bif}(p)$. But this category is just $\mathcal{Bif}(p)$ itself, given that our equivalence of setoids gives an equality of quotiented sets.

It is also possible to define cut directly on strongly multifocused derivations. We detail this construction in Appendix B.

3.5. **Normal forms via rewriting.** In this section we introduce a strongly normalizing rewriting relation between multifocused proofs that, under a certain condition on the base category, is locally confluent and thus admits unique normal forms: all proofs in each sequentialization equivalence class eventually reduce to the same normal form.

An intuition is that these normal forms will be “maximally parallel”, in the sense that \mathbb{L} and \mathbb{R} bipoles should be combined into an \mathbb{LR} bipole whenever possible. We therefore begin by considering the *parallelization* order (\rightarrow_{par}), which is just the opposite of the sequentialization order ($\alpha \rightarrow_{\text{par}} \beta$ iff $\beta \rightarrow_{\text{seq}} \alpha$). Diagrammatically it is defined by the pair of rules $\text{par}_{\mathbb{L}}^{\mathbb{R}}$ and $\text{par}_{\mathbb{R}}^{\mathbb{L}}$ shown below:

$$\begin{array}{ccccc}
 \cdot \xrightarrow{\pi f \rho} \cdot & & \cdot \xrightarrow{\pi f \rho} \cdot & & \cdot \xrightarrow{\pi f \rho} \cdot \\
 \parallel & & \parallel & & \parallel \\
 \mathbb{R}^- & \uparrow \rho & \mathbb{L}^+ & & \mathbb{L}^+ \\
 \cdot \xrightarrow{-\pi f} \cdot & & \cdot \xrightarrow{-f \rho} \cdot & & \cdot \xrightarrow{-f \rho} \cdot \\
 \parallel & & \parallel & & \parallel \\
 \mathbb{R}^+ & \downarrow \tau & \mathbb{L}^- & & \mathbb{L}^- \\
 \cdot \xrightarrow{-\pi f \tau} \cdot & & \cdot \xrightarrow{-f} \cdot & & \cdot \xrightarrow{-\sigma f \rho} \cdot \\
 \parallel & & \parallel & & \parallel \\
 \mathbb{L}^+ & & \mathbb{R}^- & & \mathbb{R}^- \\
 \pi \downarrow & & \sigma \uparrow & & \sigma \uparrow \\
 \cdot \xrightarrow{-f \tau} \cdot & & \cdot \xrightarrow{-\sigma f \tau} \cdot & & \cdot \xrightarrow{-\sigma f} \cdot \\
 \parallel & & \parallel & & \parallel \\
 \mathbb{L}^- & & \mathbb{R}^+ & & \mathbb{R}^+ \\
 \sigma \uparrow & & \downarrow \tau & & \downarrow \tau \\
 \cdot \xrightarrow{\sigma f \tau} \cdot & & \cdot \xrightarrow{\sigma f \tau} \cdot & & \cdot \xrightarrow{\sigma f \tau} \cdot
 \end{array}$$

Let us now make a few observations about the above rules. While the sequentialization order is non-deterministic in the sense that a given multifocused proof can have many different sequentializations, each rule seq_L^R and seq_R^L is deterministic: for a given \mathbb{LR} bipole α , there are unique β and γ such that $\alpha \rightarrow_{\text{seq}_L^R} \beta$ and $\alpha \rightarrow_{\text{seq}_R^L} \gamma$. Dually, the rules par_L^R and par_R^L are *co-deterministic* in the sense that for any given \mathbb{LR} bipole α there are unique β and γ such that $\beta \rightarrow_{\text{par}_L^R} \alpha$ and $\gamma \rightarrow_{\text{par}_R^L} \alpha$. On the other hand, in general the parallelization rules are not deterministic. Indeed, consider a rewriting diagram of the following form:

$$\begin{array}{ccccc}
 \cdot \xrightarrow{\pi f \rho} \cdot & & \cdot \xrightarrow{\pi f \rho} \cdot & & \cdot \xrightarrow{\pi f \rho} \cdot \\
 \parallel & & \parallel & & \parallel \\
 \mathbb{R}^- & \uparrow \rho & \mathbb{R}^- & & \mathbb{L}^+ \\
 \cdot \xrightarrow{-\pi f} \cdot & & \cdot \xrightarrow{-\pi f} \cdot & & \cdot \xrightarrow{-\pi f} \cdot \\
 \parallel & & \parallel & & \parallel \\
 \mathbb{R}^+ & \downarrow \tau & \mathbb{R}^+ & & \mathbb{L}^+ \\
 \cdot \xrightarrow{-\pi f \tau} \cdot & & \cdot \xrightarrow{-\pi f \tau} \cdot & & \cdot \xrightarrow{-\pi f \tau} \cdot \\
 \parallel & & \parallel & & \parallel \\
 \mathbb{L}^+ & & \mathbb{L}^+ & & \mathbb{L}^+ \\
 \pi \downarrow & & \pi \downarrow & & \pi \downarrow \\
 \cdot \xrightarrow{-f \tau} \cdot & & \cdot \xrightarrow{-f \tau} \cdot & & \cdot \xrightarrow{-f \tau} \cdot \\
 \parallel & & \parallel & & \parallel \\
 \mathbb{L}^- & & \mathbb{L}^- & & \mathbb{L}^- \\
 \sigma \uparrow & & \sigma \uparrow & & \sigma \uparrow \\
 \cdot \xrightarrow{\sigma f \tau} \cdot & & \cdot \xrightarrow{\sigma f \tau} \cdot & & \cdot \xrightarrow{\sigma f \tau} \cdot
 \end{array}
 \tag{3.5}$$

Under the assumption that $f\tau = g\tau = h\tau$ and $\pi f = \pi g = \pi h$, both rewrites in (3.5) are valid instances of the par_L^R rule. However, the two resulting bipoles will be distinct unless $g = h$. The same argument shows that the par_R^L rule is not deterministic in general.

A sufficient condition for the parallelization rules to be deterministic is that the underlying base category is *factorization preordered* in the sense of Johnstone [25].

Definition 3.18. A category \mathcal{C} is **factorization preordered** if for any diagram of composable arrows of the form

$$\cdot \xrightarrow{f} \cdot \xrightarrow[g]{h} \cdot \xrightarrow{k} \cdot$$

if both $fg = fh$ and $gk = hk$ then necessarily $g = h$. Equivalently, \mathcal{C} is factorization preordered just in case every commuting square has at most one diagonal filler:

$$\begin{array}{ccc} \cdot & \xrightarrow{x} & \cdot \\ f \downarrow & \nearrow g=h & \downarrow k \\ \cdot & \xrightarrow{y} & \cdot \end{array}$$

Equivalently again, \mathcal{C} is factorization preordered just in case the *category of factorizations* of every arrow is a preorder, where the category of factorizations of an arrow $w : A \rightarrow B$ has objects given by composable pairs of arrows $A \xrightarrow{u} X \xrightarrow{v} B$ such that $w = uv$, and morphisms $(A \xrightarrow{u} X \xrightarrow{v} B) \rightarrow (A \xrightarrow{s} Y \xrightarrow{t} B)$ given by arrows $h : X \rightarrow Y$ such that $uh = s$ and $ht = v$.

Example 3.19. Any category that has only monic arrows or only epic arrows is factorization preordered. In particular, every preorder category is factorization preordered, as is every free category.

Let us call ‘‘FP condition’’ the hypothesis that the base category \mathcal{C} of the free bifibration $\Lambda_p : \text{Bif}(p) \rightarrow \mathcal{C}$ is factorization preordered.

Proposition 3.20. *Under the FP condition, the parallelization rules par_L^R and par_R^L are deterministic, and the sequentialization rules seq_L^R and seq_R^L are co-deterministic.*

Determinism is a special case of confluence, ensuring that a single application of a rule does not induce a critical pair with itself. On the other hand, general confluence fails for the parallelization order, as one can exhibit a divergent critical pair from a stack of three bipoles. For example, there is a critical pair whose apex is a stack of bipoles $\mathbb{L} \mathbb{R} \mathbb{L}$,

$$\begin{array}{ccc} \begin{array}{c} \cdot \xrightarrow{afc} \cdot \\ a \downarrow \mathbb{L}^+ \parallel \\ \cdot \xrightarrow{-fc} \cdot \\ b \uparrow \mathbb{L}^- \parallel \\ \cdot \xrightarrow{-bfc} \cdot \\ \parallel \mathbb{R}^- \uparrow c \\ \cdot \xrightarrow{bf=b'g} \cdot \\ \parallel \mathbb{R}^+ \downarrow d \\ \cdot \xrightarrow{-b'gd} \cdot \\ b' \downarrow \mathbb{L}^+ \parallel \\ \cdot \xrightarrow{-gd} \cdot \\ e \uparrow \mathbb{L}^- \parallel \\ \cdot \xrightarrow{egd} \cdot \end{array} & \xrightarrow{\text{par}_R^L} & \begin{array}{c} \cdot \xrightarrow{afc} \cdot \\ a \downarrow \mathbb{L}^+ \mathbb{R}^- \uparrow c \\ \cdot \xrightarrow{-f} \cdot \\ b \uparrow \mathbb{L}^- \mathbb{R}^+ \downarrow d \\ \cdot \xrightarrow{bfd=b'gd} \cdot \\ b' \downarrow \mathbb{L}^+ \parallel \\ \cdot \xrightarrow{-gd} \cdot \\ e \uparrow \mathbb{L}^- \parallel \\ \cdot \xrightarrow{egd} \cdot \end{array} \end{array} \quad (3.6)$$

as well as a symmetric critical pair whose apex is a stack of bipoles $\mathbb{R} \mathbb{L} \mathbb{R}$. Here we clarify that the condition that $bf = b'g$ is necessary for the well-formedness of the stack in the middle

of (3.6), and implies the conditions $bfc = b'gc$ and $bfd = b'gd$ ensuring well-formedness of the stacks on the left and right.

Following the principles of Knuth–Bendix completion [27], we can obtain a confluent system by adding two new rewriting rules to orient each of these critical pairs in one direction:

$$\begin{array}{ccc}
 \begin{array}{c} \cdot \xrightarrow{afc} \cdot \\ a \downarrow \text{L}^+ \text{R}^- \uparrow c \\ \cdot \xrightarrow{-f} \cdot \\ b \uparrow \text{L}^- \text{R}^+ \downarrow d \\ \cdot \xrightarrow{bfd=b'gd} \cdot \\ b' \downarrow \text{L}^+ \parallel \\ \cdot \xrightarrow{-gd} \cdot \\ e \uparrow \text{L}^- \parallel \\ \cdot \xrightarrow{egd} \cdot \end{array} & \xrightarrow[\text{gra}_L]{bf = b'g} & \begin{array}{c} \cdot \xrightarrow{afc} \cdot \\ a \downarrow \text{L}^+ \parallel \\ \cdot \xrightarrow{-fc} \cdot \\ b \uparrow \text{L}^- \parallel \\ \cdot \xrightarrow{bfc=b'gc} \cdot \\ b' \downarrow \text{L}^+ \text{R}^- \uparrow c \\ \cdot \xrightarrow{-g} \cdot \\ e \uparrow \text{L}^- \text{R}^+ \downarrow d \\ \cdot \xrightarrow{egd} \cdot \end{array} & \xrightarrow[\text{gra}_R]{fd = gd'} & \begin{array}{c} \cdot \xrightarrow{afc} \cdot \\ \parallel \text{R}^- \uparrow c \\ \cdot \xrightarrow{-af} \cdot \\ \parallel \text{R}^+ \downarrow d \\ \cdot \xrightarrow{afd=agd'} \cdot \\ a \downarrow \text{L}^+ \text{R}^- \uparrow d' \\ \cdot \xrightarrow{-g} \cdot \\ b \uparrow \text{L}^- \text{R}^+ \downarrow e \\ \cdot \xrightarrow{bge} \cdot \end{array} \\
 & & & &
 \end{array} \tag{3.7}$$

Here we annotate the rules with the side conditions that are necessary for applying them. Observe that these conditions, which correspond to the equation appearing in the middle of (3.6) as well as in the symmetric critical pair, imply that both sides of each rewrite rule are well-formed.

We call these “gravitation” rules because they have the effect of pulling bi-focused bipoles down towards the root of the derivation. The choice of moving downwards towards the root is arbitrary, and we could have taken the opposite convention to move them up. (In the multifocusing literature the derivations generally have the shape of a tree rather than a stack, and then the convention to move multifocusing steps down towards the root is natural.)

It is clear that the gravitation rules are both deterministic and co-deterministic, without making any assumptions about the base category, since the arrows a, b, c, d, e, f appear in the source and target of each rule.

Theorem 3.21 (Local confluence). *The rewrite system $(\text{par} \cup \text{gra})$ is locally confluent: if $\beta_m \leftarrow \alpha_m \rightarrow \beta'_m$ then there exists γ_m such that $\beta_m \rightarrow^* \gamma_m \leftarrow^* \beta'_m$.*

Proof sketch. The source of each rewrite rule must be a stack of two bipoles. Thus all critical pairs (where the source of two rewrite rules overlap) can be found by considering stacks of two or three bipoles.

The proof (detailed in Appendix A.2) enumerates all possible critical pairs, and shows that each of them can be resolved. There are two non-trivial pairs (modulo symmetries), one with three bipoles




Figure 5: Resolution of a critical pair.

$\mathbb{L} \mathbb{R} \mathbb{L}$ which is exactly the critical pair (3.6) which is resolved by the rule $\mathbf{gra}_{\mathbb{L}}$, and one with three bipoles $\mathbb{L}\mathbb{R} \mathbb{L} \mathbb{R}$ involving the rules $\mathbf{gra}_{\mathbb{L}}$ on one side and $\mathbf{par}_{\mathbb{R}}^{\perp}$ on the other, which is resolved by applying $\mathbf{par}_{\mathbb{R}}^{\perp}$ on one side and $\mathbf{gra}_{\mathbb{R}}$ on the other (see Figure 5). \square

Lemma 3.22. *$\mathbf{par} \circ \mathbf{gra}$ is terminating: any valid rewriting sequence is finite.*

Proof. We provide a non-negative measure for derivations that $\mathbf{par} \circ \mathbf{gra}$ decreases strictly.

We measure each derivation by summing the weight of all bipoles. If a bipole is at position i starting from the bottom, the root of the derivation, we define its weight as i if it is a \mathbb{L} or \mathbb{R} bipole, and $2 \times i$ if it is a $\mathbb{L}\mathbb{R}$ bipole: the weight of $\mathbb{L}\mathbb{R}$ is the weight of \mathbb{L} plus the weight of \mathbb{R} .

Consider a \mathbf{gra} rewrite ($\mathbf{gra}_{\mathbb{L}}$ for example), which is of the following shape annotated with vertical positions:

$$\begin{array}{ccc} \beta & & \beta \\ \mathbb{L}\mathbb{R} & (i+1) & \mathbb{L} & (i+1) \\ \mathbb{L} & (i) & \mathbb{L}\mathbb{R} & (i) \\ \alpha & & \alpha \end{array} \xrightarrow{\mathbf{gra}_{\mathbb{L}}}$$

The starting positions of the sub-stacks α and β do not matter in this rewrite as their weight is unchanged. The weight of the middle of the derivation decreases strictly from $2 \times (i+1) + i$ to $(i+1) + 2 \times i$.

Consider a \mathbf{par} rewrite ($\mathbf{par}_{\mathbb{R}}^{\perp}$ for example), it is of the shape

$$\begin{array}{ccc} \beta & (i+2) & & \\ \mathbb{L} & (i+1) & \beta & (i+1) \\ \mathbb{R} & (i) & \mathbb{L}\mathbb{R} & (i) \\ \alpha & & \alpha \end{array} \xrightarrow{\mathbf{par}_{\mathbb{R}}^{\perp}}$$

The weight of α is unchanged by the rewrite. The combined weight of the \mathbb{L} and \mathbb{R} bipoles decreases from $i + (i+1)$ to $2 \times i$. The weight of β also decreases strictly as all its bipoles are shifted one position down. \square

From local confluence and termination, we derive our main canonicity theorem.


Theorem 3.23. *Under the FP condition, any multifocused derivation α_m admits a unique normal form for $\mathbf{par} \circ \mathbf{gra}$, which we write $\mathbf{NF}(\alpha_m)$. Moreover, for any pair of derivations α_m and β_m we have*

$$\alpha_m \sim_{\text{seq}} \beta_m \iff \mathbf{NF}(\alpha_m) = \mathbf{NF}(\beta_m).$$

3.6. Maximal multifocusing. We finally introduce a restriction of the multifocused sequent calculus that precisely captures the normal forms for $\mathbf{par} \circ \mathbf{gra}$. Under the FP condition, proofs in this restricted calculus give canonical representatives for the arrows of $\mathcal{Bif}(p)$.


An intuition is that, when building the proof from the root to the axioms, we only need information about the last invertible rule that was applied to conclude a bipole. In the immediately following bipole, we can use this information to check that the non-invertible rule does not introduce a \mathbf{par} or \mathbf{gra} redex.

Suppose for example that the last invertible rule was $\mathbb{L}\pi^+$ as part of a \mathbb{L} bipole with intermediate base arrow f , and suppose that—assuming πf is of the form $g\tau$ for some g —we want to continue the proof by applying $\mathbb{R}\tau^+$ as part of a \mathbb{R} bipole, as depicted on the left:



This creates a $\text{par}_{\mathbb{L}}^{\mathbb{R}}$ redex exactly when there is an arrow h such that $g = \pi h$ and $f = h\tau$, that is when the square to the right has a diagonal filler, corresponding to a morphism $h : (\pi, f) \rightarrow (g, \tau)$ in the category of factorizations of the composite arrow $\pi f = g\tau$.

If instead we try to extend the \mathbb{L} bipole with a $\mathbb{L}\mathbb{R}$ bipole, we should check that we don't create a gra redex. Such an extension takes the shape on the left:



This forms a gra redex exactly when f is of the form $h\tau$ for some h such that $\pi h = \sigma g$, that is when the square to the right has a diagonal filler corresponding to a morphism $h : (\pi, f) \rightarrow (\sigma g, \tau)$ in the category of factorizations of the composite arrow $\pi f = \sigma g\tau$.

Notice that in both cases, in order to decide whether or not there is a redex we need to know the intermediate arrow f of the \mathbb{L} bipole. By keeping track of the last rule applied as well as the intermediate arrow, we can thus rule out the creation of any par or gra redexes.

In Figure 6 we present a refinement of the multifocused sequent calculus that we call *maximally multifocused*, in which judgments $S \Longrightarrow_f^q T$ are annotated with a state q that can be either a *pre-inversion state* $q = q_i$ or a *pre-focus state* $q = q_f$. The rules are best read operationally as goal transformers from conclusion to premise, following the conventions of proof search. The idea is that a pre-inversion state forces an inversion rule to be applied in the next step (recall that we allow π and ρ to be empty), while a pre-focus state forces one of the focusing rules or an atomic rule to be applied. After applying a focusing rule, the state transitions to a pre-inversion state of the form \mathbb{L} , \mathbb{R} , or \perp_i , indicating the start of a \mathbb{L} -, \mathbb{R} -, or $\mathbb{L}\mathbb{R}$ -bipole respectively. The subsequent inversion rule supplements the state with additional information about the arrows involved, which is then used at the start of the next bipole to rule out any par or gra redices: this is done with the side conditions for the focusing rules, which are indicated to the side in a smaller font in Figure 6.

$$\begin{array}{l} \text{pre-inversion states } q_i ::= \perp_i \mid \mathbb{L} \mid \mathbb{R} \\ \text{pre-focus states } q_f ::= \perp_f \mid \mathbb{L}[\pi, f] \mid \mathbb{R}[f, \rho] \end{array} \quad \perp_i[\pi, f] = \perp_f = \perp_i[f, \rho]$$

$$\frac{\delta : X \rightarrow Y \in \mathcal{D} \quad p(\delta) = f}{X \xrightarrow[f]{q_f} Y} \delta$$

$$\begin{array}{ccc} \frac{N \xrightarrow[q_i[\pi, f]]{\pi f} P}{\pi^+ N \xrightarrow[f]{q_i} P} \mathbb{L}\pi^+ & \frac{N \xrightarrow[\pi f \rho]{\perp_f} P}{\pi^+ N \xrightarrow[f]{\perp_i} \rho^- P} \mathbb{L}\pi^+ \mathbb{R}\rho^- & \frac{N \xrightarrow[f \rho]{q_i[f, \rho]} P}{N \xrightarrow[f]{q_i} \rho^- P} \mathbb{R}\rho^- \\ \\ q_f = \mathbb{R}[f, \rho] \Rightarrow (\sigma, g) \not\leq (f, \rho) \frac{P \xrightarrow[g]{\mathbb{L}} Q}{\sigma^- P \xrightarrow[\sigma g]{q_f} Q} \mathbb{L}\sigma^- & q_f = \mathbb{L}[\pi, f] \Rightarrow (\pi, f) \not\leq (g, \tau) \frac{N \xrightarrow[g]{\mathbb{R}} M}{N \xrightarrow[g \tau]{q_f} \tau^+ M} \mathbb{R}\tau^+ \\ \\ q_f = \mathbb{L}[\pi, f] \Rightarrow (\pi, f) \not\leq (\sigma g, \tau) \frac{P \xrightarrow[g]{\perp_i} N}{\sigma^- P \xrightarrow[\sigma g \tau]{q_f} \tau^+ N} \mathbb{L}\sigma^- \mathbb{R}\tau^+ & & \end{array}$$

where $(a, b) \leq (c, d) := \exists e. c = ae \wedge b = ed$.

Figure 6: Maximally multifocused derivations (factorization preordered base category)

For example, suppose we are in a “left-locked” state $q_f = \mathbb{L}[\pi, f]$. That means that in the current stage of proof search, we just performed a left-focus rule to initiate a \mathbb{L} -bipole followed by a $\mathbb{L}\pi^+$ rule to terminate it. We now have the following choices are available:

- any left-focus rule can be applied freely;
- a right-focus rule $\mathbb{R}\tau^+$ beginning a \mathbb{R} bipole with intermediate arrow g can only be applied if $(\pi, f) \not\leq (g, \tau)$ in the factorization order;
- a bi-focus rule $\mathbb{L}\sigma^- \mathbb{R}\tau^+$ beginning a $\mathbb{L}\mathbb{R}$ bipole with intermediate arrow g can only be applied if $(\pi, f) \not\leq (\sigma g, \tau)$ in the factorization order.

Dual restrictions apply in a “right-locked” state $q_f = \mathbb{R}[f, \rho]$.

Our use of lock states is inspired by formulations of focusing and maximal multifocusing using tag annotations by Uustalu, Veltri, and Wan [53, 54].

By the preceding analysis we immediately have the following result.

Proposition 3.24. *Maximally multifocused proofs of unlocked judgments $S \xRightarrow{f} T$ are in one-to-one correspondence with $(\text{par} \cup \text{gra})$ -normal multifocused proofs of $S \xRightarrow{f} T$.*

Proof. The side conditions on the rules $\mathbb{L}\sigma^-$ and $\mathbb{R}\tau^+$ ensure that they never create a par redex, while the side condition on the $\mathbb{L}\sigma^- \mathbb{R}\tau^+$ rule ensures that it never creates a gra redex. Conversely, given a multifocused proof with no par or gra redex, we can always annotate it

with states following the locking discipline of Figure 6 in a uniquely determined way, starting from an unlocked state $q_i = \perp_i$. \square

Combining Proposition 3.24 with Theorem 3.23 then implies the following.

Corollary 3.25. *Under the FP condition, there is a one-to-one correspondence between arrows $\alpha : S \rightarrow T$ in $\mathcal{Bif}(p)$ such that $\Lambda_p(\alpha) = f$ and maximally multifocused proofs of $S \Longrightarrow_f^{\perp_i} T$.*

This correspondence means that if the base category \mathcal{C} is factorization preordered then maximally multifocused proofs give an equation-free inductive definition of arrows of $\mathcal{Bif}(p)$. An important special case is when $\mathcal{C} = \mathcal{FG}$ is the free category over a graph \mathbb{G} . In that case, the factorization order reduces to a linear order (a path of length n has $n + 1$ linearly ordered factorizations) and the calculus can be simplified somewhat: states only need to keep track of the type of the last bipole and the intermediate arrow f (but not π or ρ), and the factorization order test can be replaced by a single prefix or suffix test (since for any two factorizations of a word $w = uv = u'v'$, u is a prefix of u' just in case v' is a suffix of v).

On the other hand, in both the free case and general FP case, our definition of cut-elimination or horizontal composition of two multifocused proofs in Section 3.4.5 does not preserve maximality—to obtain a maximal multifocused proof one needs to renormalize the result of the composition. We leave open the question of whether there is a more direct definition of horizontal composition on maximal derivations.

3.7. Decidability and enumeration. In studying their free adjoint construction, Dawson, Paré, and Pronk [15] established the surprising result that, in general, equality of 2-cells in $\Pi_2(\mathcal{C})$ is undecidable—or more precisely that there is no decision procedure for determining whether two representatives of 2-cells (what they call “fences”) denote the same 2-cell. They gave two proofs of this result, first by exhibiting a very concrete reduction from the halting problem for 2-register machines, and second by giving a somewhat simpler and more abstract reduction from the undecidable problem of determining whether two vertices of an infinite (but locally finite) bipartite graph are connected by a path. Their latter proof translates cleanly to an analogous undecidability result for sequentialization equivalence.

Definition 3.26. We say that a category \mathcal{C} has **locally finite factorizations** just in case every commuting square in \mathcal{C} has at most finitely many diagonal fillers, or equivalently if the category of factorizations of every arrow is locally finite (i.e., has finite homsets).

By a *locally finite bipartite graph*, we mean a span of sets $V \xleftarrow{p_1} E \xrightarrow{p_2} W$ such that the fibers of p_1 and p_2 are finite. (Every edge $e \in E$ connects a vertex $v \in V$ to a vertex $w \in W$, where $v = p_1(e)$ and $w = p_2(e)$. So local finiteness says that every vertex of either class is involved in only finitely many edges, i.e., has finite degree.) To any locally finite bipartite graph G , we can associate a category \mathcal{C}_G with locally finite factorizations, defined as the subcategory of Set generated by the five sets

$$0, \quad 1, \quad V, \quad W, \quad E$$

(where 0 and 1 are respectively the empty set and singleton set), the projection functions

$$V \xleftarrow{p_1} E \xrightarrow{p_2} W$$

describing the incidence relation of G , the constant functions

$$1 \xrightarrow{v} V \quad 1 \xrightarrow{w} W \quad 1 \xrightarrow{e} E$$

associated to every vertex $v \in V$ or $w \in W$ or edge $e \in E$, and finally the unique functions from the empty set

$$0 \xrightarrow{a_X} X$$

for $X \in \{1, V, W, E\}$. To see that local finiteness of G implies that \mathcal{C}_G has locally finite factorizations, first note that the only non-degenerate commuting squares are of the form

$$\begin{array}{ccc} 0 & \xrightarrow{a_E} & E \\ a_1 \downarrow & & \downarrow p_1 \\ 1 & \xrightarrow{v} & V \end{array} \quad \begin{array}{ccc} 0 & \xrightarrow{a_E} & E \\ a_1 \downarrow & & \downarrow p_2 \\ 1 & \xrightarrow{w} & W \end{array}$$

and then observe that diagonal fillers for these squares correspond to edges incident to the respective vertices:

$$\begin{array}{ccc} 0 & \xrightarrow{a_E} & E \\ a_1 \downarrow & \nearrow e & \downarrow p_1 \\ 1 & \xrightarrow{v} & V \end{array} \quad \begin{array}{ccc} 0 & \xrightarrow{a_E} & E \\ a_1 \downarrow & \nearrow e & \downarrow p_2 \\ 1 & \xrightarrow{w} & W \end{array}$$

Now, to any vertex or edge x of G let us associate a stack α_x of the zigzag double category $\mathbb{Z}(\mathcal{C}_G)$, as follows:

$$\alpha_v := \begin{array}{c} 0 \xrightarrow{a_V} V \\ a_1 \downarrow \quad \mathbf{L}^+ \quad \parallel \\ 1 \xrightarrow{v} V \\ a_1 \uparrow \quad \mathbf{L}^- \quad \parallel \\ 0 \xrightarrow{a_V} V \\ \parallel \quad \mathbf{R}^- \quad \uparrow p_1 \\ 0 \xrightarrow{a_E} E \\ \parallel \quad \mathbf{R}^+ \quad \downarrow p_2 \\ 0 \xrightarrow{a_W} W \end{array} \quad \alpha_e := \begin{array}{c} 0 \xrightarrow{a_V} V \\ a_1 \downarrow \quad \mathbf{L}^+ \quad \mathbf{R}^- \quad \uparrow p_1 \\ 1 \xrightarrow{e} E \\ a_1 \uparrow \quad \mathbf{L}^- \quad \mathbf{R}^+ \quad \downarrow p_2 \\ 0 \xrightarrow{a_W} W \end{array} \quad \alpha_w := \begin{array}{c} 0 \xrightarrow{a_V} V \\ \parallel \quad \mathbf{R}^- \quad \uparrow p_1 \\ 0 \xrightarrow{a_E} E \\ \parallel \quad \mathbf{R}^+ \quad \downarrow p_2 \\ 0 \xrightarrow{a_W} W \\ a_1 \downarrow \quad \mathbf{L}^+ \quad \parallel \\ 1 \xrightarrow{w} W \\ a_1 \uparrow \quad \mathbf{L}^- \quad \parallel \\ 0 \xrightarrow{a_W} W \end{array}$$

Observe that α_e sequentializes to α_v and α_w , that is, $\alpha_v \leftarrow_{\text{seq}} \alpha_e \rightarrow_{\text{seq}} \alpha_w$, just in case $v = p_1(e)$ and $w = p_2(e)$.

Theorem 3.27 (cf. Theorems 2 and 3 of Dawson, Paré, and Pronk [15]). *There is a category \mathcal{C} with locally finite factorizations for which the following equivalent problems are both undecidable:*

- given two stacks of double cells in $\mathbb{Z}(\mathcal{C})$, determine whether they vertically compose to the same double cell;
- given two proofs in the bifibrational calculus generated by $p = \text{id}_{\mathcal{C}}$, determine whether they are permutation equivalent.

Proof. We take $\mathcal{C} = \mathcal{C}_{G_M}$ where G_M is a bipartite version of the configuration graph of some universal Turing machine M . Without loss of generality we will assume that M is deterministic and has a unique halting configuration τ_{halt} . Let us write Conf_M for the configuration space of M , that is, the set of tuples (q, u, v) of a state q and finite strings u and v over the working alphabet denoting the contents of the tape to the left and right of the head. We define G_M by taking $V = \text{Conf}_M \times \{0\}$ and $W = \text{Conf}_M \times \{1\}$ to be two disjoint copies of the configuration space, and $E \subseteq V \times W$ as the union

$$E = \{((\sigma, 0), (\tau, 1)) \mid \sigma \rightarrow \tau \text{ is a valid transition of } M\} \\ \cup \{((\sigma, 0), (\sigma, 1)) \mid \sigma \in \text{Conf}_M\}$$

By the preceding observation, it is clear that any directed chain of configurations

$$\sigma_1 \rightarrow \sigma_2 \rightarrow \sigma_3 \rightarrow \dots$$

induces a corresponding chain of sequentialization equivalences:

$$\alpha_{(\sigma_1, 0)} \sim_{\text{seq}} \alpha_{(\sigma_2, 1)} \sim_{\text{seq}} \alpha_{(\sigma_2, 0)} \sim_{\text{seq}} \alpha_{(\sigma_3, 1)} \sim_{\text{seq}} \alpha_{(\sigma_3, 0)} \sim_{\text{seq}} \dots$$

Conversely, the assumption that M is deterministic and has a unique halting state ensures that any chain of sequentialization equivalences $\alpha_{(\sigma, 0)} \sim_{\text{seq}} \dots \sim_{\text{seq}} \alpha_{(\tau_{\text{halt}}, 1)}$ corresponds to a halting execution $\sigma \rightarrow^* \tau_{\text{halt}}$. We have thus reduced the halting problem for M to deciding sequentialization equivalence $\alpha_{(\sigma, 0)} \sim_{\text{seq}} \alpha_{(\tau_{\text{halt}}, 1)}$, or equivalently permutation equivalence $\alpha_{(\sigma, 0)} \sim \alpha_{(\tau_{\text{halt}}, 1)}$. \square

Dawson, Paré, and Pronk proved decidability of (fence equivalence for) the $\Pi_2(\mathcal{C})$ construction separately in the cases that \mathcal{C} is locally finite or factorization preordered. We now prove analogues of these two results for deciding equality of morphisms in $\mathcal{Bif}(p)$, or more precisely for deciding whether two proofs of the same judgment represent the same morphism. The first is already a consequence of the results of Section 1, while the second is a corollary of our normal form theorem.

Theorem 3.28 (cf. Propositions 2 and 3 of Dawson, Paré, and Pronk [15]). *Let $p : \mathcal{D} \rightarrow \mathcal{C}$ and suppose that \mathcal{C} satisfies either of the following conditions:*

- (1) \mathcal{C} is locally finite, or
- (2) \mathcal{C} is factorization preordered;

then for any two proofs $\alpha_1, \alpha_2 : S \Longrightarrow_f T$ of the bifibrational calculus generated by p , it is decidable whether $\alpha_1 \sim \alpha_2$.

Proof. (1) Two proofs are equivalent just in case there is some sequence of permutations transforming one to the other. Any such permutation leaves invariant the length of the proof, the choice of terminating axiom $\delta : X \rightarrow Y$ in \mathcal{D} , and the selection of left and right rules applied, merely changing the order of the rules and the intermediate base arrows that annotate the derivation. Since \mathcal{C} is locally finite, there are therefore only finitely many proofs permutation equivalent to any given proof, and we can reduce $\alpha_1 \sim \alpha_2$ to a finite test $\alpha_2 \in [\alpha_1]$.

(2) This is a corollary of our normal theorem, Theorem 3.23, which states that each multifocused proof has a unique maximal normal form, under the FP condition. We can thus decide whether two proofs denote the same morphism of $\mathcal{Bif}(p)$ simply by computing their normal forms and checking that these are equal. \square

Remark 3.29. Technically, we need \mathcal{C} to satisfy computable versions of either local finiteness or factorization preorder to get a decision procedure. For instance, we need to be able to effectively enumerate the finite homset $\mathcal{C}(A, B)$ between any pair of objects. Or, for any given square, we need to be able to decide whether that square is empty or else compute a filler. We also need the equality of morphisms in \mathcal{C} and \mathcal{D} to be decidable to check that two proofs are identical. We omit these precisions from the statement of the theorem, since it is clear from context that the conditions should be interpreted constructively.

Besides entailing decidability of the word problem, the normal form theorem has the equally important application of allowing us to *enumerate* homsets of free bifibrations without duplicates, assuming the FP condition. We are most interested in the situation that $\mathcal{Bif}(p)$ is relatively locally finite in the following sense.

We say that a category \mathcal{D} equipped with a functor $p : \mathcal{D} \rightarrow \mathcal{C}$ is **relatively locally finite** (relative to p) if the *relative homsets* over f

$$\mathcal{D}_f(X, Y) := \{ \delta \mid \delta : X \rightarrow Y, p(\delta) = f \}$$

are finite, for every arrow $f : A \rightarrow B$ of \mathcal{C} and pair of objects X and Y of \mathcal{D} lying above A and B respectively. (Thus a category \mathcal{C} is locally finite iff the unique functor $!_{\mathcal{C}} : \mathcal{C} \rightarrow \mathbf{1}$ is relatively locally finite.)

Theorem 3.30. *Let $p : \mathcal{D} \rightarrow \mathcal{C}$ and suppose \mathcal{C} satisfies the FP condition. If \mathcal{D} is relatively locally finite (relative to p) and \mathcal{C} has locally finite slices and coslices, then $\mathcal{Bif}(p)$ is relatively locally finite (relative to Λ_p). In particular, the fiber category $\mathcal{Bif}(p)_A$ of every object $A \in \mathcal{C}$ is locally finite.*

Proof. Suppose given an arrow $f : A \rightarrow B$ of \mathcal{C} and a pair of bifibrational formulas $S \sqsubset A$ and $T \sqsubset B$. By Corollary 3.25, we can enumerate canonical representatives of arrows $\alpha \in \mathcal{Bif}(p)_f(S, T)$ using the sequent calculus of Figure 6 to search for maximally multifocused proofs of $S \Longrightarrow_f^{\perp_i} T$. There are only three potential sources of nondeterminism:

- (1) the choice of whether to initiate a L, R, or LR bipole when proving a neutral sequent;
- (2) the choice of an arrow g in the premise of a left-/right-/bi-focus rule;
- (3) the choice of an axiom $\delta : X \rightarrow Y \in \mathcal{D}$ to terminate the proof.

But each of these choice points are finite branching:

- (1) at most three rules to try (some may be forbidden by the lock state q_f or if one side of the sequent is atomic);
- (2) finite by the assumption that \mathcal{C} has locally finite slice and coslice categories (for example, to derive a sequent of the form $\sigma^- P \Longrightarrow_h^q Q$ using the $L\sigma^-$ rule, we must pick an arrow g such that $h = \sigma g$; observe that such an arrow corresponds to a morphism $\sigma \rightarrow h$ in the coslice category A/\mathcal{C} , where $A = \text{dom } \sigma = \text{dom } h$);
- (3) finite by the assumption that \mathcal{D} is relatively locally finite.

We conclude that there are only finitely many maximally multifocused proofs of any given judgment, and hence that the relative homsets are finite. In particular, for every object $A \in \mathcal{C}$, the homset of the fiber category $\mathcal{Bif}(p)_A$, corresponding to the relative homset over id_A , is finite. \square

Corollary 3.31. *Under the assumptions of Theorem 3.30, logical equivalence of bifibrational formulas is decidable.*

Proof. Recall that by definition, two formulas $S_1, S_2 \sqsubset A$ are logically equivalent if there exist derivations $\alpha : S_1 \Longrightarrow_{\text{id}_A} S_2$ and $\beta : S_2 \Longrightarrow_{\text{id}_A} S_1$ whose compositions are permutation equivalent to identities (Defn. 1.19), or equivalently just in case they are isomorphic in the fiber category $\text{Bif}(p)_A$ (Prop. 1.20). It suffices to enumerate canonical representatives of all such derivations α and β using Theorem 3.30, compute their cuts $\alpha \cdot \beta$ and $\beta \cdot \alpha$, and decide whether $\alpha \cdot \beta \sim \text{id}_{S_1}$ and $\beta \cdot \alpha \sim \text{id}_{S_2}$ using Theorem 3.28. \square

As already remarked, free categories satisfy the FP condition because their factorization categories are linear orders. Moreover, although free categories $\mathcal{F}\mathbb{G}$ are not locally finite unless the underlying graph \mathbb{G} is acyclic, their slices and coslices are posetal and hence trivially locally finite. Thus the enumeration and decision procedures described above are applicable (and even become slightly simpler since the sequent calculus of Figure 6 can be simplified a bit, as remarked at the end of Section 3.6).

Finally, Theorem 3.30 can clearly be generalized to derive any infinite cardinality bound on the relative homsets of $\text{Bif}(p)$, assuming the same bound on the relative homsets of \mathcal{D} and the homsets of the slices and coslices of \mathcal{C} . We can also use the sequent calculus to effectively enumerate the (potentially infinite) relative homsets of $\text{Bif}(p)$, given procedures for doing the same for \mathcal{D} and for the slices and coslices of \mathcal{C} .

4. FREE $(\mathcal{P}, \mathcal{N})$ -FIBRATIONS

As previously mentioned, Dawson, Paré, and Pronk's $\Pi_2(\mathcal{C})$ construction may be seen as a weaker version of the free groupoid construction, which turns a category into a groupoid by freely adding an inverse to every arrow. More often, one only wants to add formal inverses to a class of arrows \mathcal{W} , a process known as *localization* (see, e.g., Borceaux [10, Ch.5]), which freely builds a category $\mathcal{C}[\mathcal{W}^{-1}]$ equipped with a functor $\mathcal{C} \rightarrow \mathcal{C}[\mathcal{W}^{-1}]$ sending every arrow $f \in \mathcal{W}$ to an isomorphism. DPP [14, §4.1] discuss how to adapt their construction of $\Pi_2(\mathcal{C})$ to define a 2-category $\mathcal{C}[\mathcal{W}^*]$ in which every arrow $f \in \mathcal{W}$ is freely equipped with a right adjoint. In this section we quickly explain how our construction of the free bifibration on a functor $p : \mathcal{D} \rightarrow \mathcal{C}$ may be similarly adapted when the operations of pushing and pulling are restricted to certain classes of arrows of \mathcal{C} .

Definition 4.1. Let $\mathcal{P} \subset \mathcal{C}$ and $\mathcal{N} \subset \mathcal{C}$ be two subsets of arrows (not necessarily subcategories) of \mathcal{C} . We say that a functor $\pi : \mathcal{E} \rightarrow \mathcal{C}$ is a $(\mathcal{P}, \mathcal{N})$ -**fibration** if it has $+$ -cartesian liftings of all arrows in \mathcal{P} and $-$ -cartesian liftings of all arrows in \mathcal{N} .

As special cases, a (\emptyset, \mathcal{C}) -fibration is a Grothendieck fibration in the ordinary sense, a (\mathcal{C}, \emptyset) -fibration is an opfibration, and a $(\mathcal{C}, \mathcal{C})$ -fibration is a bifibration.

A case of particular interest is when $(\mathcal{P}, \mathcal{N})$ constitute a factorization system for the base category. Such $(\mathcal{P}, \mathcal{N})$ -fibrations relative to a factorization system have been termed *ambifibrations* by Joachim Kock and André Joyal. One interesting property of ambifibrations, noted by Kock [30], is that they give rise to ternary factorization systems in which every arrow of the total category factors as a $+$ -cartesian arrow, followed by an arrow lying over

an identity, followed by a \dashv -cartesian arrow:

$$\begin{array}{ccc}
 \mathcal{E} & & \\
 \downarrow \pi & & \\
 \mathcal{C} & &
 \end{array}
 \begin{array}{ccc}
 S & \xrightarrow{\alpha} & T \\
 \searrow e_S & & \nearrow \bar{m}_T \\
 e^+ S & \xrightarrow{e \setminus \alpha / \bar{m}} & m^- T
 \end{array}
 \begin{array}{ccc}
 A & \xrightarrow{f} & B \\
 \searrow e \in \mathcal{P} & & \nearrow m \in \mathcal{N}
 \end{array}$$

We will consider an example of a free ambifibration in Section 5.3.

Unless otherwise stated we do not assume anything about \mathcal{P} and \mathcal{N} in this section. We write $\Lambda_{p, \mathcal{P}, \mathcal{N}} : \mathcal{Bif}(p, \mathcal{P}, \mathcal{N}) \rightarrow \mathcal{C}$ for the free $(\mathcal{P}, \mathcal{N})$ -fibration on a functor $p : \mathcal{D} \rightarrow \mathcal{C}$.

Sequent calculus. The presentation of the free bifibration that we gave via the sequent calculus in Section 1 adapts directly to a presentation of the free $(\mathcal{P}, \mathcal{N})$ -bifibration. For this one simply has to restrict the formation rules on bifibrational formulas:

$$\frac{S \sqsubset A \quad f : A \rightarrow B \in \mathcal{P}}{f^+ S \sqsubset B} \qquad \frac{g : B \rightarrow C \in \mathcal{N} \quad T \sqsubset C}{g^- T \sqsubset B}$$

Otherwise, everything remains the same: the definition of the inference rules, permutation equivalence, cut-elimination, et cetera. This presentation makes clear that $\mathcal{Bif}(p, \mathcal{P}, \mathcal{N})$ is a full subcategory of $\mathcal{Bif}(p)$.

The zigzag double category. In Section 2 we defined the double category of zigzags $\mathbb{Z}(\mathcal{C})$ by first constructing the free bifibration on the identity functor $\text{id}_{\mathcal{C}}$. We can similarly define a *double category of restricted zigzags* $\mathbb{Z}(\mathcal{C}, \mathcal{P}, \mathcal{N})$, whose vertical arrows are zigzags moving positively along arrows in \mathcal{P} and negatively along arrows in \mathcal{N} . In particular, we can recover DPP's $\mathcal{C}[\mathcal{W}^*]$ in this way as the vertical 2-category of $\mathbb{Z}(\mathcal{C}, \mathcal{C}, \mathcal{W})$.

Weak focusing. The syntax of weakly focused derivations given in Section 3.3 works just as well for $(\mathcal{P}, \mathcal{N})$ -fibrations, although the proof of correctness based on the equivalence $\mathcal{Bif}(p) \simeq \text{sa}\mathcal{Bif}(p)$ has to be slightly adapted. The equivalence with strictly alternating formulas relied on the pseudofunctoriality laws $(g \circ f)^+ S \equiv g^+ f^+ S$ and $(f g)^- T \equiv f^- g^- T$, whereas \mathcal{P} and \mathcal{N} are not necessarily closed under composition. However, Theorem 3.2 can be directly adapted to prove the following.

Proposition 4.2. *If \mathcal{P} and \mathcal{N} are closed under binary composition, then $\mathcal{Bif}(p, \mathcal{P}, \mathcal{N}) \simeq \text{sa}\mathcal{Bif}(p, \mathcal{P}, \mathcal{N})$.*

The restrictions on \mathcal{P} and \mathcal{N} are not problematic, because we can always close them under finite compositions and obtain an equivalent free $(\mathcal{P}, \mathcal{N})$ -fibration. Given a subset of arrows $\mathcal{W} \subset \mathcal{C}$, let us write \mathcal{W}^* for the subcategory of \mathcal{C} generated by adding identity arrows id_A and id_B for every $f : A \rightarrow B \in \mathcal{W}$, and by closing under binary composition. The pseudofunctoriality laws immediately imply the following.

Proposition 4.3. $\mathcal{Bif}(p, \mathcal{P}, \mathcal{N}) \simeq \mathcal{Bif}(p, \mathcal{P}^*, \mathcal{N}^*)$.

Corollary 4.4. *If $\mathcal{C} = \mathcal{F}\mathbb{G}/\sim$ is freely generated by some graph \mathbb{G} modulo some equations, then $\mathcal{Bif}(p) \simeq \mathcal{Bif}(p, \mathbb{G}, \mathbb{G})$.*

Maximal multifocusing. Finally, the results of Sections 3.5–3.7 provided a strong notion of normal form for free bifibrations as well as decidability results in the case that the base category is factorization preordered (Definition 3.18). One benefit of considering free $(\mathcal{P}, \mathcal{N})$ -fibrations is that we can relax this requirement to the a priori weaker condition that \mathcal{C} is both \mathcal{P} -factorization preordered and \mathcal{N} -factorization preordered.

Definition 4.5. Let $\mathcal{W} \subset \mathcal{C}$ be a subset of arrows of a category \mathcal{C} . We say that \mathcal{C} is **\mathcal{W} -factorization preordered** (or **\mathcal{W} -FP**) if for any diagram of the form

$$\cdot \xrightarrow{f \in \mathcal{W}} \cdot \begin{array}{c} \xrightarrow{g} \\ \xrightarrow{h} \end{array} \cdot \xrightarrow{k \in \mathcal{W}} \cdot$$

if both $fg = fh$ and $gk = hk$ then necessarily $g = h$. Equivalently, \mathcal{C} is factorization preordered just in case every commuting square of the following form has at most one diagonal filler:

$$\begin{array}{ccc} \cdot & \xrightarrow{x} & \cdot \\ f \in \mathcal{W} \downarrow & \nearrow g=h & \downarrow k \in \mathcal{W} \\ \cdot & \xrightarrow{y} & \cdot \end{array}$$

Proposition 4.6. *Every category \mathcal{C} is both \mathcal{C}_{epi} -FP and $\mathcal{C}_{\text{mono}}$ -FP, where \mathcal{C}_{epi} and $\mathcal{C}_{\text{mono}}$ consist respectively of the epic arrows and the monic arrows of \mathcal{C} .*

The analysis of Sections 3.5–3.7 carries over and we obtain the following.

Theorem 4.7. *If \mathcal{C} is both \mathcal{P} -FP and \mathcal{N} -FP, then arrows of $\text{Bif}(p, \mathcal{P}, \mathcal{N})$ are canonically represented by maximally multifocused derivations, and permutation equivalence is decidable. If moreover \mathcal{D} is relatively locally finite and \mathcal{C} has locally finite slices and coslices, then $\text{Bif}(p, \mathcal{P}, \mathcal{N})$ is relatively locally finite and logical equivalence is decidable.*

Remark 4.8. As mentioned, fibrations and opfibrations arise as particular cases of $(\mathcal{P}, \mathcal{N})$ -fibration, taking $\mathcal{P} = \emptyset, \mathcal{N} = \mathcal{C}$ for fibrations and $\mathcal{P} = \mathcal{C}, \mathcal{N} = \emptyset$ for opfibrations. Even if \mathcal{C} is not factorization preordered, strongly focused derivations do provide normal forms for the arrows of $\text{Bif}(p, \emptyset, \mathcal{C})$ or $\text{Bif}(p, \mathcal{C}, \emptyset)$, for the simple reason that there are no $\mathbb{L}\mathbb{R}$ -bipoles and so sequentialization equivalence reduces to discrete equality. The usual constructions of the free fibration and free opfibration as comma categories may be recovered in this way if one restricts to strictly alternating formulas of depth 1, avoiding formulas of depth 0 so as to obtain a split fibration (cf. Remark 3.4).

5. THE COMBINATORICS OF FREE BIFIBRATIONS AND AMBIFIBRATIONS

We conclude with a deeper analysis of the specific examples of free bifibrations and ambifibrations that were mentioned in the Introduction. In all of these examples, some rich combinatorial structure emerges out of very simple generating data, and the exploration of this structure is facilitated by the bifibrational calculus.

5.1. Reconstructing the simplex category. Recall that we write Δ for the category of finite ordinals and order-preserving maps, which is sometimes referred to as the *augmented* simplex category to emphasize that it includes the empty ordinal 0, but which we will simply call the simplex category. We write Δ_{\perp} for the category of non-empty finite ordinals and order-and-least-element preserving maps. We write $L : \Delta \rightarrow \Delta_{\perp}$ for the functor sending an ordinal n to a non-empty ordinal n' by adding a \perp element, and an order-preserving map $\theta : m \rightarrow n$ to the order-and-least-element-preserving map $\theta' : m' \rightarrow n'$ defined by $\theta'(\perp) = \perp$ and $\theta'(x) = \theta(x)$ for $x > \perp$. This functor has a right adjoint given by the forgetful functor $R : \Delta_{\perp} \rightarrow \Delta$, which interprets n' as $1+n$ and forgets that a map is \perp -preserving.

Proposition 5.1. *Let $F : 2 \rightarrow \text{Adj}$ and $p_2 : 1 \rightarrow 2$ be the functors shown below:*

$$F0 = \Delta \quad F1 = \Delta_{\perp} \quad Ff = \Delta \begin{array}{c} \xrightarrow{L} \\ \perp \\ \xleftarrow{R} \end{array} \Delta_{\perp} .$$

$$\begin{array}{ccc} 1 & * & \\ p_2 \downarrow & & \\ 2 & 0 \xrightarrow{f} 1 & \end{array}$$

Then the unique morphism of bifibrations $\text{Bif}(p_2) \rightarrow \int_2 F$ sending the atomic formula $*$ to the empty ordinal 0 is an equivalence of categories. In particular, Δ is equivalent to $\text{Bif}(p_2)_0$.

Proposition 5.1 is a corollary of Theorem 5.3 below, so we do not give a detailed proof here, but we sketch the correspondence. Under the equivalence, each ordinal n in Δ is interpreted as the bifibrational formula

$$\ulcorner n \urcorner := \underbrace{f^- f^+ \dots f^- f^+}_{n \text{ times}} * \sqsubset 0$$


while each non-empty ordinal n' in Δ_{\perp} is interpreted as the bifibrational formula

$$\ulcorner n' \urcorner := f^+ \underbrace{f^- f^+ \dots f^- f^+}_{n \text{ times}} * \sqsubset 1 .$$

Conversely, since $*$ is the only atom and f is the only non-identity arrow, every bifibrational formula on p_2 is equivalent to one of these two forms. To be precise, we note that pushing or pulling along an identity arrow id_0 or id_1 is also possible, but logically equivalent to doing nothing (Prop. 1.21)—hence why we only state an equivalence $\int_2 F \equiv \text{Bif}(p_2)$ rather than an isomorphism of categories. Clearly we obtain an isomorphism if we restrict to the full subcategory $\text{Bif}(p_2, f, f) \subset \text{Bif}(p_2)$.

Observe that ordinal formulas $\ulcorner n \urcorner$ and $\ulcorner n' \urcorner$ are strictly alternating. The sequent calculus of Figure 6 becomes particularly simple when restricted to such ordinal formulas: since there are no non-trivial factorizations in the base category, the bi-focusing and bi-inversion rules can be removed (they are impossible to apply), as can the lock conditions on the left- and right-focusing rules (they are always validated). Thus any morphism in $\text{Bif}(p_2)$ may be represented canonically as a strongly monofocused derivation, which in turn is essentially just a shuffle of \mathbb{L} and \mathbb{R} bipoles.

As an illustration, the three order-preserving maps $2 \rightarrow 2$ depicted on the left correspond to the following three derivations of $f^- f^+ f^- f^+ * \Longrightarrow_{\text{id}_0} f^- f^+ f^- f^+ *$ depicted on the right (drawn as stacks with overlaid string diagrams):



The graphical correspondence should be suggestive. In general, any morphism $m \rightarrow n$ in $\mathcal{Bif}(p_2)_0$ can be canonically represented as a derivation that opens (reading from the bottom) with a R^- rule, followed by an arbitrary shuffle of m \mathbb{L} bipoles with $n - 1$ \mathbb{R} bipoles, followed by a closing R^+ rule. Since there are $\binom{m+n-1}{m}$ such shuffles, we thus derive the well-known formula for the number of order-preserving maps.


Corollary 5.2. *There are $\frac{(m+n-1)!}{m!(n-1)!}$ order-preserving maps $m \rightarrow n$.*

5.2. Categories of trees and walks. A natural generalization of the previous example is obtained by considering the free bifibration generated by the functor p_ω illustrated below which maps the point to the initial object of ω , the totally ordered set of natural numbers considered as a category.

$$\begin{array}{ccc}
 1 & * & \\
 p_\omega \downarrow & & \\
 \omega & 0 \xrightarrow{f_0} 1 \xrightarrow{f_1} 2 \xrightarrow{f_2} \dots &
 \end{array}$$

As we observed in the introduction, up to isomorphism the objects of $\mathcal{Bif}(p_\omega)_0$ may be read as Dyck words, which are in bijection with rooted plane trees, that is, trees embedded in the plane with one vertex marked as the root, considered up to planar isotopy. We draw

such trees growing upwards from the root, as depicted on the left:




To encode a plane tree by a Dyck word, a canonical procedure is to traverse the tree (say) from right to left starting from the root, recording a $+1$ every time you move up towards a child, and a -1 every time you move down back to a parent. We represent the beginning of this traversal, visually, on the right. (We will use these same colors below to visualize the traversal order.)

The corresponding bifibrational formula may be read off directly:

$$f^- f^- f^- f^+ f^+ f^- f^+ f^+ f^- f^+ f^- f^- f^+ f^+ *$$

Here to keep the notation simple we have elided the indices from the f_i 's, which are uniquely determined. Such a word may also be interpreted as a walk on the upper plane, starting at the x -axis (at the point marked $*$ below), constantly moving up or down while moving leftwards and eventually terminating back at the x -axis:




Observe that all of the information about the walk is contained in its peaks and valleys, and a tree may also therefore be represented more concisely as a strictly alternating zigzag:

$$0 \longrightarrow 3 \longleftarrow 1 \longrightarrow 2 \longleftarrow 0 \longrightarrow 1 \longleftarrow 0 \xrightarrow{\text{orange}} 2 \xleftarrow{\text{green}} 0$$

or again as a formula in strictly alternating syntax: $f^{-3} f^{+2} f^{-1} f^{+2} f^{-1} f^{+1} f^{-2} f^{+2} *$.


Next we recall how plane trees admit an alternative representation as functors. (This idea goes back at least to Joyal [26], but may be earlier folklore.) Define the *height* of a vertex as its distance from the root, and the height of a tree as the maximum height of any of its vertices. Let $[k] = k+1$ be the linear order $\{0 < \dots < k\}$ seen as a posetal category. It is clear from the definitions that any rooted plane tree of height k may be represented by a functor $T : [k]^{op} \rightarrow \Delta$ such that $T(0) = 1$: each linearly ordered set $T(i)$ defines the vertices of height i , and the parents of the non-root vertices are specified by the functions $T(i+1) \rightarrow T(i)$, which must be order-preserving by the constraints of planarity. For instance, the height-3 tree given in the example above corresponds to the functor $T : [3] \rightarrow \Delta$ depicted

below.



This tree also has *width* 3, defined as the maximum number of vertices at any given height. More generally, a tree of unbounded height and width may be represented as a functor $T : \omega^{op} \rightarrow \Delta$ such that $T(0) = 1$. In other words, the only requirement is that T sends the initial object of ω to the terminal object of Δ . Finiteness is ensured if we moreover ask that there is some natural number k such that $T(k + 1) = 0$ (and hence $T(j) = 0$ for all $j > k$).

We define PTree as the full subcategory of $[\omega^{op}, \Delta]$ spanned by the finite rooted plane trees, that is, the category whose objects are functors $T : \omega^{op} \rightarrow \Delta$ such that $T(0) = 1$ and $T(k + 1) = 0$ for some k , and all natural transformations between them. For example, below we depict two different morphisms of plane trees $S \rightarrow T$ between the same pair of trees, using different color lines to indicate the mappings $S(k) \rightarrow T(k)$ of the natural transformations at each height k .



Readers familiar with the “topos of trees” $[\omega^{op}, \text{Set}]$ may recognize that PTree is defined in essentially the same way, but replacing Set by Δ to reflect planarity, and taking a full subcategory of $[\omega^{op}, \Delta]$ in order to restrict to finite trees rather than potentially infinite forests. The combination of order-preservation and naturality imposes strong constraints on the shape of plane tree morphisms. The reader may check that there are exactly 11 morphisms $S \rightarrow T$ between the trees S and T shown above, including the two in (5.1).

We should note that PTree is strongly related to but different from a category of plane trees defined by Joyal [26]. Joyal considered a more liberal notion of morphism $S \rightarrow T$ between plane trees encoded in the same way as functors $S, T : \omega^{op} \rightarrow \Delta$. Instead of considering natural transformations in $[\omega^{op}, \Delta]$, he took more general natural transformations between the underlying functors in $[\omega^{op}, \text{Set}]$ so long as they respect the linear orders on the *fibers* of the maps $S(n + 1) \rightarrow S(n)$. (In other words, respect the ordering of siblings, but not necessarily of cousins.) Every morphism in PTree is a morphism in Joyal’s category


of plane trees (i.e., it is a wide subcategory), but the converse is not true as the following example illustrates:

$$\begin{array}{ccc}
 S(2) = 2 \rightarrow T(2) = 2 & \begin{array}{ccccc}
 0 & & 1 & & 0 & & 1 \\
 | & & | & & / & & \backslash \\
 0 & & 1 & & 0 & & 0 \\
 / & & \backslash & & | & & / \\
 0 & & 0 & & 0 & & 0
 \end{array} & \\
 \downarrow & & \downarrow & & & & \\
 S(1) = 2 \rightarrow T(1) = 1 & & & & & & \\
 \downarrow & & \downarrow & & & & \\
 S(0) = 1 \rightarrow T(0) = 1 & & & & & &
 \end{array} \tag{5.2}$$

Joyal used this more relaxed definition of tree morphism in order to axiomatize a notion of ω -category, and his category of trees, which is equivalent to the category Ω studied by Batanin and Street [5], satisfies some important universal properties (see Theorems 1 and 2 of [5]). Despite being a non-full subcategory of the Joyal-Batanin-Street category of trees, the category of plane trees PTree defined above as a full subcategory of $[\omega^{op}, \Delta]$ also seems to us very natural.

Let us now prove that PTree corresponds to the fiber over 0 of the free bifibration generated by p_ω . In order to prove this we give a more general characterization of the fiber categories $\text{Bif}(p_\omega)_k$ for arbitrary $k \geq 0$. We write $\text{PTree}_{\perp(k)}$ for the category whose objects are rooted plane trees with a marked vertex of height k on their leftmost branch, and whose arrows are morphisms of plane trees preserving the mark. It is clear that $\text{PTree} \cong \text{PTree}_{\perp(0)}$, since every tree morphism $S \rightarrow T$ sends the root of S to the root of T , and also that every tree in $\text{PTree}_{\perp(k)}$ must have height at least k . But not every morphism between trees of height $\geq k$ is a morphism in $\text{PTree}_{\perp(k)}$. For example, one of the two tree morphisms $S \rightarrow T$ depicted in (5.1) lifts to a morphism in $\text{PTree}_{\perp(2)}$ (the one in blue), but the other (in red) only lifts to $\text{PTree}_{\perp(1)}$.

Objects of $\text{Bif}(p_\omega)_k$ may be interpreted as left-marked trees: the idea is that a bifibrational formula in $S \sqsubset k$, corresponding to a walk that starts at height 0 and ends at height k , may be read as describing a walk from right to left along the tree that stops at the marked node. Under this correspondence, extending a walk by an up-step $k \rightarrow k + 1$ corresponds to growing a fresh marked leaf as a left child of the marked node, while extending a walk by a down-step $k + 1 \rightarrow k$ corresponds to moving the marked node down towards the root. We call these respective operations L_k and R_k , and illustrate them below on an example for the case $k = 2$:



It is easy to see that L_k and R_k extend to functors defining a family of adjunctions

$$\begin{array}{ccc}
 & L_k & \\
 & \curvearrowright & \\
 \text{PTree}_{\perp(k)} & \perp & \text{PTree}_{\perp(k+1)} \\
 & \curvearrowleft & \\
 & R_k &
 \end{array}$$

between the categories of left-marked plane trees.

Theorem 5.3. *Let $F : \omega \rightarrow \text{Adj}$ be the functor sending k to $\text{PTree}_{\perp(k)}$ and each arrow $f_k : k \rightarrow k + 1$ to the adjunction $L_k \dashv R_k$ defined above. Let $p_\omega : \mathbf{1} \rightarrow \omega$ be the functor mapping $*$ to 0 . Then the unique morphism of bifibrations $\text{Bif}(p_\omega) \rightarrow \int_\omega F$ sending the atomic formula $*$ to the one-vertex tree \bullet is an equivalence of categories. In particular, PTree is equivalent to $\text{Bif}(p_\omega)_0$.*

Proof. Since $\pi_F(\bullet) = 0 = p_\omega(*)$, the universal property of the free bifibration determines a unique morphism of bifibrations $\llbracket - \rrbracket_\bullet : \text{Bif}(p_\omega) \rightarrow \int_\omega F$ mapping $*$ to \bullet .

On objects, $\llbracket - \rrbracket_\bullet$ interprets a bifibrational formula T as the left-marked tree $\llbracket T \rrbracket_\bullet$ encoded by the corresponding walk. This mapping is surjective since every such tree is described by a unique walk. (It is many-to-one, since many formulas isomorphic by the pseudofunctoriality laws represent the same walk.)

As a morphism of bifibrations, the mapping on arrows is determined by its action on the homsets of the fiber categories. We now give an explicit description of a family of bijections

$$(\psi_k)_{S,T} : \text{Bif}(p_\omega)_k(S, T) \xrightarrow{\sim} \text{PTree}_{\perp(k)}(\llbracket S \rrbracket_\bullet, \llbracket T \rrbracket_\bullet)$$

mapping zigzag morphisms to plane tree morphisms. Since these bijections are natural in S and T , they uniquely determine $\llbracket - \rrbracket_\bullet : \text{Bif}(p_\omega) \rightarrow \int_\omega F$ as a morphism of bifibrations. We conclude that $\llbracket - \rrbracket_\bullet$ is both surjective on objects and fully faithful, hence an equivalence.

To define ψ_k , we rely on an inductive characterization of plane tree morphisms that we show can be reflected in both $\text{PTree}_{\perp(k)}$ and $\text{Bif}(p_\omega)_k$. Let us begin by illustrating this bijection in the case $k = 0$, before explaining how to generalize it to arbitrary $k \geq 0$.

Bijection for the case $k = 0$. A plane tree can be described inductively as a root node with a list of plane trees as children. In other words, a tree is built by taking a (possibly empty) forest of trees, and attaching the root of each tree in the forest to a new root vertex. The following grammar summarizes this inductive structure, where T ranges over trees and Φ over forests:

$$T ::= \bullet \Phi \quad \Phi ::= \cdot \mid T, \Phi' \tag{5.3}$$

Order-preserving morphisms of plane trees can similarly be defined inductively. We find it convenient to introduce two auxiliary judgments $\Phi \rightarrow T$ and $\Phi' \rightarrow \Phi$, describing forest-to-tree morphisms and forest-to-forest morphisms, respectively; we then consider tree morphisms as a special case of forest-to-tree morphisms with a singleton forest as input. The two kinds of morphisms can be built inductively using the following rules:

$$\frac{\Phi'_1, \dots, \Phi'_i \rightarrow \Phi}{\bullet \Phi'_1, \dots, \bullet \Phi'_i \rightarrow \bullet \Phi} \quad \frac{\Phi'_1 \rightarrow T_1 \quad \dots \quad \Phi'_n \rightarrow T_n}{\Phi'_1, \dots, \Phi'_n \rightarrow T_1, \dots, T_n} \tag{5.4}$$


The forest-to-tree rule (on the left) expresses that the root nodes of the input forest are necessarily mapped to the unique root of the output tree, so it suffices to provide a forest-to-forest morphism at height 1. The forest-to-forest rule (on the right) expresses that a morphism into a forest with n trees is determined by an ordered partition of the input forest into n subforests together with a list of forest-to-tree morphisms.

It is clear that morphisms in PTree may be constructed inductively in this way—to see this, it is helpful to begin by observing that there is an isomorphism

$$\begin{array}{ccc}
 & T(-) \mapsto T(1+-) & \\
 & \curvearrowright & \\
 \text{PTree} & \sim & \text{PForest} \\
 & \curvearrowleft & \\
 & \Phi \mapsto \bullet\Phi &
 \end{array}$$


where PForest is the category of (finite rooted) plane forests, that is, the full subcategory of $[\omega^{op}, \Delta]$ spanned by the functors $\Phi : \omega^{op} \rightarrow \Delta$ that eventually vanish $\Phi(k) = 0$ (with no constraint on $\Phi(0)$). Indeed, any natural transformation $\theta : S \rightarrow T$ in PTree is uniquely determined at height 0 (since $S(0) = T(0) = 1$) and induces a natural transformation $\theta' : \Phi_S \rightarrow \Phi_T$ in PForest between the underlying forests, where $S = \bullet\Phi_S$ and $T = \bullet\Phi_T$. The height 0 component of such a natural transformation $\theta'_0 : \Phi_S(0) \rightarrow \Phi_T(0)$ determines an ordered partition of the roots of Φ_S into n components, where $\Phi_T = T_1, \dots, T_n$. By naturality of θ' , this extends to an ordered partition of the forest itself $\Phi_S = \Phi'_1, \dots, \Phi'_n$, together with a list of natural transformations $\theta_1 : \Phi'_1 \rightarrow T_1, \dots, \theta_n : \Phi'_n \rightarrow T_n$.

Let us now explain how to similarly decompose morphisms in $\text{Bif}(p_\omega)_0$, illustrating with an example. Consider the maximally multifocused proof below, which we will see corresponds to the blue morphism in (5.1):




We have rotated the proof 90 degrees clockwise to better fit on the page. It should be read as a morphism $S \rightarrow T$ from the zigzag S at the top to the zigzag T at the bottom. Remember that each zigzag describes a walk along the contour of a tree.

The idea is to begin by splitting the proof along $0 = 0$ boundaries:




This splits the proof into exactly n components, where $T = \bullet(T_1, \dots, T_n)$. Next, in each of these components, we “gray out” the chords that are incident to 0:




Each component now implicitly describes a forest-to-tree morphism $\Phi_i \rightarrow T_i$, where $S = \bullet\Phi_S$ and Φ_1, \dots, Φ_n is an ordered partition of Φ_S . If we remove the grayed-out chords and shift levels—in terms of the sequent calculus, this corresponds to removing the pair of R^- and R^+ steps bracketing each component, as well as a list of pairs of L^- and L^+ steps—then we can iterate the process on smaller proofs. Eventually we obtain a full recursive decomposition of the original zigzag morphism $S \rightarrow T$ using the inductive rules in (5.4), which we can play back as a morphism $[[S]]_\bullet \rightarrow [[T]]_\bullet$ in PTree . This defines the bijection ψ_k for the case $k = 0$.

The bijection may be better visualized by imagining the string diagram for the zigzag morphism as a kind of topographic relief map:



The source and target trees are rooted at the North and South ends of this map, with the morphism exhibiting a continuous deformation between them.



Bijection in the general case. The inductive procedure we have sketched may be adapted to the general case $k \geq 0$ by extending the rules (5.3) and (5.4) to represent trees with a marked node on their leftmost branch. We begin by introducing a grammar of formulas for marked trees ${}^\circ T$ and marked forests ${}^\circ \Phi$:


$${}^\circ T ::= \circ\Phi \mid \bullet^\circ\Phi \quad {}^\circ\Phi ::= {}^\circ T, \Phi' \tag{5.5}$$

A marked forest ${}^\circ\Phi$ starts with a marked tree ${}^\circ T$ followed by an unmarked forest Φ' . A marked tree ${}^\circ T$ can be built either by taking an unmarked forest Φ and attaching the root of each tree to a marked root vertex, $\circ\Phi$, or by taking a marked forest ${}^\circ\Phi$ and attaching each of its roots to a new unmarked root vertex, $\bullet^\circ\Phi$. We introduce two additional judgments ${}^\circ\Phi \rightarrow_k {}^\circ T$ and ${}^\circ\Phi' \rightarrow_k {}^\circ\Phi$ defining marked-forest-to-marked-tree and marked-forest-to-marked-forest morphisms, where the index k specifies the height of the mark:

$$\frac{\Phi'_0, \Phi'_1, \dots, \Phi'_i \rightarrow \Phi}{\circ\Phi'_0, \bullet\Phi'_1, \dots, \bullet\Phi'_i \rightarrow_0 \circ\Phi} \quad \frac{{}^\circ\Phi'_0, \Phi'_1, \dots, \Phi'_i \rightarrow_k \Phi}{\bullet({}^\circ\Phi'_0), \bullet\Phi'_1, \dots, \bullet\Phi'_i \rightarrow_{k+1} \bullet^\circ\Phi} \quad \frac{{}^\circ\Phi_0 \rightarrow_k {}^\circ T \quad \Phi \rightarrow \Phi'}{{}^\circ\Phi_0, \Phi \rightarrow_k {}^\circ T, \Phi'} \tag{5.6}$$

The key invariant enforced by these rules is that the marked node of the forest on the left is sent to the marked node of the tree or forest on the right. Just as we argued for $\mathbf{PTree}_\perp(k)$, it is clear that any morphism in $\mathbf{PTree}_\perp(k)$ may be built inductively using the above rules.

Similarly, for $k > 0$, the string diagrams of morphisms in $\mathbf{Bif}(p_\omega)_k$ may be decomposed recursively by first splitting along $0 = 0$ boundaries essentially as before, with the difference that the bottom boundary (corresponding to the left side in our rotated diagrams) is of the form $k = k$. For example, the following diagram represents a morphism in $\mathbf{Bif}(p_\omega)_2$:



After splitting along $0 = 0$ boundaries, we obtain as before a list of components that may be interpreted as morphisms in $\mathbf{Bif}(p_\omega)_0$, except for the leftmost component which may




Figure 7: An exhaustive listing of maximally multifocused derivations representing all morphisms $\alpha : S \rightarrow T$ between a pair of plane trees. The derivations corresponding to the two natural transformations in (5.1) are highlighted in blue and red.

considered earlier. These were generated by applying the proof search procedure described in Section 3.6, which we implemented as a short Haskell program.³

Before moving on to the last example, let us make one more easy observation about $\text{Bif}(p_\omega)$. The ordinal ω may be seen as the category of elements of the regular action of the monoid of natural numbers $(\mathbb{N}, +, 0)$ on itself. In particular, there is an evident forgetful functor $\delta : \omega \rightarrow \mathbf{B}_{\mathbb{N}}$ into the monoid of natural numbers seen as the one-object category freely generated on a loop $f \curvearrowright *$. This functor sends every pair of natural numbers $m \leq n$

³ Available at github.com/noamz/free-bifibrations. This implementation was useful to us for running experiments, formulating conjectures, and generating diagrams. Counting proofs using this procedure even helped to reveal a bug in a previous incorrect formulation of the lock conditions for maximally multifocused proofs, before we arrived at the correct version in Section 3.6.

to the unique arrow $f^k : * \rightarrow *$ such that $k = n - m$; since ω is a poset, δ is trivially a faithful functor.

Proposition 5.4. *If $\delta : \mathcal{C} \rightarrow \mathcal{B}$ is a faithful functor then the induced morphism of bifibrations*

$$\begin{array}{ccc} \mathcal{Bif}(p) & \xrightarrow{\Lambda_\delta} & \mathcal{Bif}(\delta \circ p) \\ \Lambda_p \downarrow & & \downarrow \Lambda_{\delta \circ p} \\ \mathcal{C} & \xrightarrow{\delta} & \mathcal{B} \end{array}$$

is full and faithful relative to δ , in the sense that the induced mapping on relative homsets $\mathcal{Bif}(p)_f(S, T) \rightarrow \mathcal{Bif}(\delta \circ p)_{\delta f}(\Lambda_\delta S, \Lambda_\delta T)$ is a bijection for all compatible S, f, T . (Equivalently, the induced functor from $\mathcal{Bif}(p)$ into the fiber product $\mathcal{C} \times_{\Lambda_{\delta \circ p}} \mathcal{Bif}(\delta \circ p)$ is fully faithful in the ordinary sense.)

Proof. The existence of the morphism of bifibrations $\Lambda_\delta : \Lambda_p \rightarrow \Lambda_{\delta \circ p}$ comes from the extended universal property (0.5) of the free bifibration Λ_p and corresponds to functoriality of the Λ construction. It has a simple inductive definition:

$$\Lambda_\delta X = X \quad \Lambda_\delta(f^+ S) = (\delta f)^+ \Lambda_\delta S \quad \Lambda_\delta(g^- T) = (\delta g)^- \Lambda_\delta T$$

$$\Lambda_\delta(f \setminus_g \alpha) = \delta f \setminus_{\delta g} \Lambda_\delta \alpha \quad \Lambda_\delta(\alpha \cdot f) = \Lambda_\delta \alpha \cdot \delta f \quad \Lambda_\delta(\bar{g} \cdot \alpha) = \bar{\delta g} \cdot \Lambda_\delta \alpha \quad \Lambda_\delta(\alpha \cdot f / \bar{g}) = \Lambda_\delta \alpha \cdot \delta f / \bar{\delta g}$$

The action of Λ_δ may be even easier to see in terms of the double category of zigzags: every arrow of every generator is simply hit with the functor δ . For example, every $\mathbf{L}f^+$ -cell of $\mathbb{Z}(\mathcal{C})$ is mapped to a $\mathbf{L}(\delta f)^+$ -cell of $\mathbb{Z}(\mathcal{B})$, and every $\mathbf{L}g^-$ -cell to a $\mathbf{L}(\delta g)^-$ -cell:

$$\begin{array}{ccc} \begin{array}{ccc} A & \xrightarrow{fg} & C \\ f \downarrow & & \parallel \\ B & \xrightarrow{g} & C \end{array} & \mapsto & \begin{array}{ccc} \delta A & \xrightarrow{\delta(fg)} & \delta C \\ \delta f \downarrow & & \parallel \\ \delta B & \xrightarrow{\delta g} & \delta C \end{array} \end{array} \quad \begin{array}{ccc} \begin{array}{ccc} C & \xrightarrow{g'} & C' \\ g \uparrow & & \parallel \\ B & \xrightarrow{gg'} & C' \end{array} & \mapsto & \begin{array}{ccc} C & \xrightarrow{\delta g'} & C' \\ \delta g \uparrow & & \parallel \\ B & \xrightarrow{\delta(gg')} & C' \end{array} \end{array}$$

In general, we obtain a mapping of relative homsets of the form:

$$\mathcal{Bif}(p)_f(S, T) \longrightarrow \mathcal{Bif}(\delta \circ p)_{\delta f}(\Lambda_\delta S, \Lambda_\delta T)$$


By inspection of the action on the generators, we can see that surjectivity of this mapping reduces to functoriality of δ , and that injectivity reduces to faithfulness. \square

Corollary 5.5. *Let $!_{\mathbb{N}} : 1 \rightarrow \mathbf{B}_{\mathbb{N}}$ be the unique functor of one-object categories corresponding to the unique monoid homomorphism from the trivial monoid into $(\mathbb{N}, +, 0)$. Then the induced morphism of bifibrations $\Lambda_\delta : \Lambda_{p_\omega} \rightarrow \Lambda_{!_{\mathbb{N}}}$ is fully faithful relative to $\delta : \omega \rightarrow \mathbf{B}_{\mathbb{N}}$. In particular, \mathbf{PTree} is equivalent to the full subcategory of the fiber category $\mathcal{Bif}(!_{\mathbb{N}})_*$ spanned by the objects in the image of Λ_δ .*

Proof. We apply Proposition 5.4 substituting $!_{\mathbb{N}} = \delta \circ p_\omega$. Then we instantiate the bijection of relative homsets $\mathcal{Bif}(p_\omega)_f(S, T) \cong \mathcal{Bif}(!_{\mathbb{N}})_{\delta f}(\Lambda_\delta S, \Lambda_\delta T)$ at $f = \text{id}_0$, $\delta f = \text{id}_*$ and conclude using Theorem 5.3. \square


The category $\mathcal{Bif}(!_{\mathbb{N}})$ is in some sense more fundamental than $\mathcal{Bif}(p_\omega)$. Its objects may be interpreted as arbitrary one-dimensional walks, starting at the x -axis and moving up and

down with no restrictions. Below we show an example formula and the associated walk, which crosses below the x -axis:



Every such formula can be assigned an integer corresponding to its final *displacement* relative to the x -axis (in the above case: $+1$). Indeed, displacement corresponds to the object component of a functor $\llbracket - \rrbracket_0 : \mathcal{Bif}(!_{\mathbb{N}}) \rightarrow (\mathbb{Z}, \leq)$ obtained from the universal property of $\mathcal{Bif}(!_{\mathbb{N}})$ by considering the poset (\mathbb{Z}, \leq) as a discrete bifibration over $\mathbf{B}_{\mathbb{N}}$ (where $f^+ n = n + 1$ and $f^- n = n - 1$), equipped with the origin point 0 .

$\mathcal{Bif}(!_{\mathbb{N}})$ contains \mathbf{PTree} as a full subcategory, per Corollary 5.5, but it also contains morphisms between objects that are not Dyck walks, as the below example illustrates:



Here we have indicated the running totals of the displacements of the walks as a guide even though $\mathbf{B}_{\mathbb{N}}$ really only has one object $*$. Speaking literally, the diagram above can be interpreted as a double cell of $\mathbb{Z}(\mathbb{Z}, \leq)$ mapping canonically to an arrow of $\mathcal{Bif}(!_{\mathbb{N}})$.

5.3. A free ambifibration over Δ . The last example we consider is inspired by the so-called *fat Delta* category $\mathbf{\Delta}$, introduced by Joachim Kock [29] to axiomatize a notion of higher category with weak identity arrows. One way of defining $\mathbf{\Delta}$ is as a subcategory of the arrow category of Δ , with objects given by epis $n \twoheadrightarrow k$ and with arrows given by commutative squares whose upper leg is mono:

$$\begin{array}{ccc}
 n & \hookrightarrow & p \\
 \downarrow & & \downarrow \\
 k & \longrightarrow & j
 \end{array}
 \tag{5.8}$$

Observe that there is a canonical embedding $j : \mathbb{N} \rightarrow \mathbf{\Delta}$ of the set of natural numbers viewed as a discrete category into fat Delta, defined by $j(n) = \text{id}_n$. Moreover, there is an evident projection functor $\text{cod} : \mathbf{\Delta} \rightarrow \Delta$, which is known to be an ambifibration relative to the (epi, mono) factorization system on Δ (see Sattler [47] and Pradal [45]).

As well-known (see Mac Lane [38, VII.5]), the simplex category admits an (epi, mono) factorization system, with the epis being generated by the maps $\sigma_i^n : n + 1 \rightarrow n$ ($0 \leq i < n$) and the monos by the maps $\delta_i^n : n \rightarrow n + 1$ ($0 \leq i \leq n$), where σ_i^n is the unique order-preserving surjection sending i and $i + 1$ to i , and δ_i^n is the unique order-preserving injection

whose image avoids element i . Moreover, these generators satisfy the following standard equations (please recall that we write composition in diagrammatic order):

$$\delta_i^n \delta_{j+1}^{n+1} = \delta_j^n \delta_i^{n+1} \quad \text{if } i \leq j \quad (5.9)$$

$$\sigma_{j+1}^{n+1} \sigma_i^n = \sigma_i^{n+1} \sigma_j^n \quad \text{if } i \leq j \quad (5.10)$$

$$\delta_i^n \sigma_j^n = \sigma_{j-1}^{n-1} \delta_i^{n-1} \quad \text{if } i < j \quad (5.11)$$


$$= \text{id}_n \quad \text{if } i = j \text{ or } i = j + 1 \quad (5.12)$$

$$= \sigma_j^{n-1} \delta_{i-1}^{n-1} \quad \text{if } i > j + 1 \quad (5.13)$$

Now let $\mathcal{P} = \{\sigma_i^n \mid 0 \leq i < n\}$ and $\mathcal{N} = \{\delta_i^n \mid 0 \leq i \leq n\}$ and consider the free $(\mathcal{P}, \mathcal{N})$ -fibration generated by the inclusion $i : \mathbb{N} \rightarrow \Delta$ of the set of natural numbers viewed as a discrete category into the simplex category. Since \mathcal{P} and \mathcal{N} generate the epis and monos respectively, $\mathcal{P}^* = \Delta_{\text{epi}}$, $\mathcal{N}^* = \Delta_{\text{mono}}$, we have that $\text{Bif}(i, \mathcal{P}, \mathcal{N}) \simeq \text{Bif}(i, \Delta_{\text{epi}}, \Delta_{\text{mono}})$ by Prop. 4.3, and so $\Lambda_{i, \mathcal{P}, \mathcal{N}} : \text{Bif}(i, \mathcal{P}, \mathcal{N}) \rightarrow \Delta$ is the free (epi, mono)-ambifibration on i .


The universal property of $\text{Bif}(i, \mathcal{P}, \mathcal{N})$ and the equation $i = \text{cod} \circ j$ imply that there is a canonical morphism of (epi, mono)-ambifibrations $\text{Bif}(i, \mathcal{P}, \mathcal{N}) \rightarrow \Delta$.

Let us now forget about this connection with fat Delta, and just make some combinatorial observations about $\text{Bif}(i, \mathcal{P}, \mathcal{N})$. We refer to objects of $\text{Bif}(i, \mathcal{P}, \mathcal{N})$ as *ambisimplicial formulas*, and will pay particular attention to formulas $S \sqsubset 0$, which we call *closed*. Observe that a closed ambisimplicial formula corresponds to a walk $n \rightsquigarrow 0$ along the graph below:



At each positive ordinal $k+1$, we choose to either pull along some δ_i^k or push along some σ_i^k . Hence there are double factorial $(2n - 1)!! = (2n - 1) \cdot (2n - 3) \cdots 3 \cdot 1$ many such formulas with source $\chi(S) = n$. The double factorial numbers are known to have many different combinatorial interpretations [11], and we state here another: they count n -leaf *increasing binary forests*, by which we mean a planar forest of binary trees where every binary node and every root node is assigned a distinct label in such a way that the labels increase when moving from parent to child, and from a binary node bounding a region to a root contained in that region. See Figure 8.

To interpret a closed ambisimplicial formula S as an increasing binary forest, the idea is to start by drawing the $\chi(S) = n$ edges, and then to progressively add either a binary node (joining two edges into one) or a root node (terminating an edge), labelled m and attached to the i th free edge, depending on whether we pushed along an epi $\sigma_i^m : m + 1 \rightarrow m$ or pulled along a mono $\delta_i^m : m \rightarrow m + 1$ in constructing the formula. For example, here is how we build the increasing binary forest corresponding to the formula $\delta_0^{0-} \sigma_0^{1+} \delta_1^{2-} 3$:






Figure 8: Depiction of all 15 closed ambisimplicial formulas with $\chi(S) = 3$ as increasing binary forests. The labels A–G indicate the split into seven equivalence classes.

Now, let us quickly point out that different forests/formulas can be *isomorphic* in the category $\text{Bif}(i, \mathcal{P}, \mathcal{N})$. In particular, equations (5.9) and (5.10) induce natural isomorphisms

$$\delta_i^{n-} \delta_{j+1}^{n+1-} S \equiv \delta_j^{n-} \delta_i^{n+1-} S \quad (5.14)$$

$$\sigma_i^{n+} \sigma_{j+1}^{n+1+} S \equiv \sigma_j^{n+} \sigma_i^{n+1+} S \quad (5.15)$$

for all $i \leq j$. In Figure 8 we indicate the equivalence classes for the case $\chi(S) = 3$ by the letters A–G. Each of these equivalence classes may be canonically represented by a strictly alternating zigzag in $\mathbb{Z}(\Delta, \Delta_{\text{epi}}, \Delta_{\text{mono}})$, as per the following correspondence for the seven equivalence classes of the figure (we only label the arrows that are not uniquely determined):

$$A = 0 \rightarrow 3 \quad B = 0 \rightarrow 2 \xleftarrow{\sigma_0^0} 3 \quad C = 0 \rightarrow 2 \xleftarrow{\sigma_1^1} 3 \quad D = 0 \rightarrow 1 \leftarrow 2 \xrightarrow{\delta_0^0} 3$$

$$E = 0 \rightarrow 1 \leftarrow 3 \quad F = 0 \rightarrow 1 \leftarrow 2 \xrightarrow{\delta_1^1} 3 \quad G = 0 \rightarrow 1 \leftarrow 2 \xrightarrow{\delta_2^2} 3$$

The sequence counting the number of equivalence classes of closed ambisimplicial formulas,

$$1, 1, 2, 7, 35, 226, 1787, 16717, \dots,$$

is apparently OEIS A014307, known to have various other interpretations [46].

Perhaps most interestingly, the remaining equations of the simplex category (5.11)–(5.13), which mix δ and σ , induce non-invertible *logical entailments* between formulas:

$$\sigma_{j-1}^{n-1+} \delta_i^{n-} S \implies \delta_i^{n-1-} \sigma_j^{n+} S \quad \text{if } i < j \quad (5.16)$$

$$\delta_i^{n-} S \implies \sigma_j^{n+} S \quad \text{if } i = j \text{ or } i = j + 1 \quad (5.17)$$

$$\sigma_j^{n-1+} \delta_i^{n-} S \implies \delta_{i-1}^{n-1-} \sigma_j^{n+} S \quad \text{if } i > j + 1 \quad (5.18)$$

The middle entailment is obtained as the composite of cartesian arrows $\overline{\delta}_{iS} \cdot \sigma_{jS}$, which lies over the identity by equation (5.12). For the first and third, recall that in general, any commutative square $ab = cd$ in the base of a bifibration (or in the base of a $(\mathcal{P}, \mathcal{N})$ -fibration with $b, c \in \mathcal{P}$ and $a, d \in \mathcal{N}$) induces a canonical natural transformation $c^+ a^- S \Longrightarrow d^- b^+ S$:


$$\frac{\frac{\frac{\frac{\text{————— } \text{id}_S}{S \Longrightarrow S}}{\frac{a^- S \Longrightarrow b^+ S}{ab}}}{\frac{a^- S \Longrightarrow b^+ S}{cd}}}{c^+ a^- S \Longrightarrow d^- b^+ S} \text{L}a^- Rb^+ \quad \text{L}c^+ Rd^- \tag{abcd}$$

The Beck-Chevalley condition asks that this natural transformation be invertible whenever $ab = cd$ is a pullback square, but we are in a free setting where the BC condition fails.

All of these non-invertible entailments have simple graphical interpretations, either as moving a binary node upwards relative to a root node somewhere on its right/left, or as joining a root node to an edge on its immediate left/right to form a binary node, as drawn below (we only show one version of each type of transformation, eliding the mirror image):

$$\tag{5.19}$$

Now define a family of posets $(F_{k,n})_{k,n \in \mathbb{N}}$ as follows. The carrier of $F_{k,n}$ is the set of ambisimplicial formulas $S \sqsubset k$ such that $\chi(S) = n$ (i.e., forests with k open roots and n open leaves), quotiented by (5.14) and (5.15). The order relation is the least congruence generated by (5.16)–(5.18), where by congruence we mean that $S \leq T$ implies $e^+ S \leq e^+ T$ and $m^- S \leq m^- T$. For example, here is the Hasse diagram of $F_{0,3}$ using the same labels for the equivalence classes from Figure 8:



Theorem 5.6. $\text{Bif}(i, \mathcal{P}, \mathcal{N})_k \simeq \coprod_{n \in \mathbb{N}} F_{k,n}$.

Proof sketch. We have already established that the laws (5.14)–(5.18) hold in the free ambifibration, which means that the $F_{k,n}$ embed into $\text{Bif}(i, \mathcal{P}, \mathcal{N})_k$.

Conversely, we will begin by showing more generally that if $\alpha : S \Longrightarrow_f T$ is an arrow of $\text{Bif}(i, \mathcal{P}, \mathcal{N})$ and $f = em$ is the epi-mono factorization of the underlying arrow in Δ , then $e^+ S \leq m^- T$ in the appropriate poset $F_{k,n}$ with $n = \chi(S) = \chi(T)$ and $k = \text{cod}(e) = \text{dom}(m)$. We prove this by induction on derivations. The interesting case is when α ends in a non-invertible rule, for example suppose it is of the form $\alpha = \bar{g}.\alpha' : g^- S \Longrightarrow_{gh} T$ for some $\alpha' : S \Longrightarrow_h T$ and $g \in \mathcal{N} \subset \Delta_{\text{mono}}$. By the induction hypothesis, we have that $e^+ S \leq m^- T$ where $h = em$, $e \in \Delta_{\text{epi}}$, $m \in \Delta_{\text{mono}}$. Let $f = dk$, $d \in \Delta_{\text{epi}}$, $k \in \Delta_{\text{mono}}$ be the epi-mono factorization of $f = gh$. Thus we have $dk = gem$. By the orthogonal lifting property there is a unique mono ℓ making the diagram below commute:

$$\begin{array}{ccc} \cdot & \xrightarrow{g} & \cdot & \xrightarrow{e} & \cdot \\ d \downarrow & & \swarrow \ell & & \downarrow m \\ \cdot & \xrightarrow{k} & \cdot & & \cdot \end{array}$$

Hence $d^+ g^- S \leq \ell^- e^+ S \leq \ell^- m^- T \equiv k^- T$, where the first entailment is an instance of the (abcd) scheme on the square $ge = d\ell$. (Here we need a lemma that every instance of the scheme is valid, $c^+ a^- S \leq d^- b^+ S$, for $a, b \in \Delta_{\text{epi}}$ and $c, d \in \Delta_{\text{mono}}$, which we can prove by induction on the length of the \mathcal{P} -factorization of c .)

Finally, to establish the equivalence $\text{Bif}(i, \mathcal{P}, \mathcal{N})_k \simeq \coprod_{n \in \mathbb{N}} F_{k,n}$ we need to verify that the fiber categories are preorders. In fact, we claim that the total category $\text{Bif}(i, \mathcal{P}, \mathcal{N})$ is itself a preorder: for any pair of ambisimplicial formulas $S \sqsubset n, T \sqsubset m$ there exists at most one arrow $\alpha : S \Longrightarrow_f T$ lying over a unique order-preserving function $f : n \rightarrow m$. Intuitively, the reason why is that left division by an epi e and right division by a mono m are *deterministic* and *independent* operations:


- *determinism*: given an arrow f , there is at most one g such that $f = eg$ and at most one h such that $f = hm$, so given any derivation α with underlying base arrow f , we can write $e \setminus \alpha$ and α / \bar{m} rather than $e \setminus_g \alpha$ and $\alpha /_h \bar{m}$ without risk of ambiguity;
- *independence*: if α is divisible on the left by e and divisible on the right by m (i.e., both $e \setminus \alpha$ and α / \bar{m} are defined), then it is simultaneously divisible on the left and right (i.e., $e \setminus \alpha / \bar{m}$ is defined), by the orthogonal lifting property:

$$\begin{array}{ccc} \cdot & \xrightarrow{h} & \cdot \\ e \downarrow & & \swarrow i \\ \cdot & \xrightarrow{g} & \cdot \\ & & \downarrow m \end{array}$$


Determinism and independence imply that any two derivations of arrows $\alpha_1, \alpha_2 : S \rightarrow T$ are permutation equivalent: they necessarily perform the same series of multiplication and division operations, although not necessarily in the same order, but if we examine the first place that they differ — say that α_1 performs a left division $e \setminus \beta$ and α_2 a right division β / \bar{m} — then we can always replace that by a bi-division $e \setminus \beta / \bar{m}$ and continue unifying the rest of the proofs. \square

We remark that since Δ is both \mathcal{P} -FP and \mathcal{N} -FP (like any category relative to a class of epis and monos), every entailment $S \leq T$ in $F_{k,n}$ is witnessed by a unique maximally multifocused proof $\alpha : S \Longrightarrow T$. For example, below on the right we show the maximally multifocused proof corresponding to the forest entailment on the left (we work with strictly alternating zigzags, using graphical conventions for the corresponding forests that should be clear, and we notate order-preserving functions $f : n \rightarrow m$ by tuples of numbers a_0, \dots, a_{m-1} that

indicate the cardinalities of the fibers $a_i = |f^{-1}(i)|$:



The $(F_{k,n})_{n \in \mathbb{N}}$ seem to define an interesting family of posets. For example, we observe that $F_{0,n}$ contains the *lattice of noncrossing partitions* as a quotient. A noncrossing partition of order n is a partition of the ordinal $\mathbf{n} = \{0, \dots, n-1\}$ such that the diagram defined by drawing the elements in order and then connecting all of the elements within the same block by a corolla is planar. For example, of the $\binom{4}{2} = 7$ partitions of 4 into two blocks, all but one are noncrossing:



The restriction of the usual refinement ordering to noncrossing partitions of order n defines a poset K_n known as the *Kreweras lattice*.

Proposition 5.7. K_n is the quotient of $F_{0,n}$ obtained by inverting (5.16) and (5.18), or equivalently by imposing the Beck-Chevalley condition on the free ambifibration $\Lambda_{i,p,N}$.

Proof. The squares corresponding to equations (5.11) and (5.13) are pullback squares in Δ , and hence the respective natural transformations (5.16) and (5.18) become invertible if one assumes the BC condition. The noncrossing partition associated to a closed ambisimplicial formula is obvious if one looks at the corresponding increasing forest, since imposing (5.16) and (5.18) as equivalences in addition to (5.14) and (5.15) has the effect of forgetting the

order in which nodes are created and reducing the forest down to its connected components. For instance, formulas A_1 – A_6 from Figure 8 correspond to the partition $\{\{0\}, \{1\}, \{2\}\}$, while formulas C_1 , C_2 , and D_1 correspond to the partition $\{\{0\}, \{1, 2\}\}$. The only remaining law (5.17) (see right side of (5.19) for its diagrammatic depiction) exactly describes the covering relation for the refinement order on noncrossing partitions. \square

For example, K_3 is obtained from $F_{0,3}$ by identifying two pairs of equivalence classes $C \sim D$ and $B \sim G$. Proposition 5.7 and the seminal results of Kreweras [31] on the lattice of noncrossing partitions as well as related results for other Catalan lattices [7] motivate us to ask two open questions.⁴

Question 5.8. K_n is a lattice for all n : is $F_{0,n}$ (and more generally $F_{k,n}$) likewise a lattice?

Question 5.9. The number of intervals $I[K_n] = \{(x, y) \mid x \leq y\}$ in K_n is given by a simple formula $|I[K_n]| = \frac{1}{2n+1} \binom{3n}{n}$: is there likewise a nice formula for $|I[F_{0,n}]|$ (and more generally for $|I[F_{k,n}]|$)?

It is known that the lattice of noncrossing partitions is self-dual $K_n \cong K_n^{op}$, and even *locally self-dual* in the sense that for every pair of noncrossing partitions $x, y \in K_n$, the restriction of K_n to the interval $[x, y]$ is self-dual [52, V2:465]. It may be interesting to study this duality from a fibrational perspective, in particular relating it to the duality between Δ and the category of finite strict intervals and order-and-boundary-preserving maps [26, §1.1].

6. CONCLUSION AND FUTURE DIRECTIONS

This work originally began as an exercise in categorical proof theory, with our hope being that a systematic understanding of the construction of the free bifibration on a functor of categories $p : \mathcal{D} \rightarrow \mathcal{C}$ could eventually help in analyzing more complex fibrational situations that arise in logic and type theory. This “easy” exercise proved more challenging than we expected (particularly the question of characterizing normal forms to obtain a canonicity result), but was also rewarding as we realized the implicit complexity and richness contained in the problem. We were encouraged by the close connection between the sequent calculus presentation of the free bifibration and the presentation of the double category of zigzags by generators and relations, showing the unity of purpose between two basic tools in the proof-theoretic and category-theoretic repertoires; and we were excited to discover the examples of free bifibrations and ambifibrations in which the combinatorial complexity emerged from much simpler generating data. We think there are many natural paths for further study.

Perhaps the most obvious next step would be to generalize the constructions beyond functors of categories. As mentioned in the Introduction, we were originally motivated by phenomena that arise when considering bifibrations of monoidal closed categories and of generalized multicategories and polycategories. We expect that the construction of the free bifibration on a functor should generalize to functors of polycategories, but the details are by no means obvious, especially the appropriate generalization of the FP condition and whether there is a reasonable notion of maximal multifocusing. Another natural generalization would be to incorporate higher cells in the base category, such as by considering functors

⁴Soon after the original version of this paper was posted, we found out that the answer to Question 5.8 is negative: already $F_{0,4}$ is not a lattice. We thank Wenjie Fang for determining this using SageMath from data provided by our decision procedure for entailment (see Footnote 3).

of 2-categories or double categories. Indeed, we already somewhat implicitly treated the base category as a 2-category in our conventions for drawing string diagrams where we represented equalities by vertices, as in Figure 1; taking the base to be an actual 2-category would give this convention a more formal status, and may also make it possible to relax the FP condition while retaining canonicity. It should be noted that Licata, Shulman, and Riley [37] considered bifibrations of cartesian 2-multicategories, and that 2-cells were used to good effect in their notion of “mode theory”.

We used string diagrams in this paper mainly for their evocative power (particularly in the proof of Theorem 5.3), but we believe there is value in taking them seriously as geometric objects and developing a purely topological or purely combinatorial description of free bifibrations. Dawson, Paré, and Pronk did exactly that in their description of $\Pi_2(\mathcal{C})$ for a free category, defining a bijection between 2-cells in $\Pi_2(\mathcal{C})$ and isotopy classes of planar diagrams presented by matchings on a set of oriented points. It could be worthwhile to extend such a correspondence for free bifibrations beyond the case of a free base category, to general categories or 2-categories.

Although we did not explore the direction in this paper, there is clearly a link to be developed between bifibrations or $(\mathcal{P}, \mathcal{N})$ -fibrations and game semantics [23, 39]. It is noteworthy that the free bifibration on $p_2 : 1 \rightarrow 2$ appeared incognito in a paper on the categorical combinatorics of innocent strategies [21], with the total category $\mathcal{Bif}(p_2)$ of the free bifibration being equivalent to the authors’ *category of schedules* Υ .

The surprising connection between the free bifibration on $p_\omega : 1 \rightarrow \omega$ and the Joyal-Batanin-Street category of plane trees—which contains $\mathcal{Bif}(p_\omega)_0$ as a wide subcategory but is not equivalent to it—is still somewhat mysterious to us, and we are not sure whether it is a mere coincidence. Similarly, we do not know whether the free ambifibration on $i : \mathbb{N} \rightarrow \Delta$ may (like the fat Delta ambifibration) be useful for defining weak higher categories, but it appears to be an interesting object of study in its own right.

Acknowledgments. We warmly thank François Lamarche for stimulating discussions about his work on bifibrations [32, 33] and for his enthusiasm about the project. We also thank Nathanael Arkor, Tom Hirshowitz, Paul-André Melliès, Dorette Pronk, and Gilles Schaeffer for discussions and pointers.

This work began while the three authors were at LIX and the first author was on a postdoc supported by the ANR project Repro (ANR-20-CE48-0016). The third author was partially supported by the ANR project LambdaComb (ANR-21-CE48-0017).

REFERENCES

- [1] Benedikt Ahrens and Peter LeFanu Lumsdaine. Displayed categories. *Logical Methods in Computer Science*, 15(1), 2019. doi:10.23638/LMCS-15(1:20)2019.
- [2] Jean-Marc Andreoli. Logic programming with focusing proofs in linear logic. *Journal of Logic and Computation*, 2(3):297–347, 1992. doi:10.1093/logcom/2.3.297.
- [3] Jean-Marc Andreoli. Focussing and proof construction. *Annals of Pure and Applied Logic*, 107(1):131–163, 2001. doi:10.1016/S0168-0072(00)00032-4.
- [4] Claude Auderset. Adjonctions et monades au niveau des 2-catégories. *Cahiers de topologie et géométrie différentielle*, 15(1):3–20, 1974.
- [5] Michael Batanin and Ross Street. The universal property of the multitude of trees. *Journal of Pure and Applied Algebra*, 154(1-3):3–13, 2000. doi:10.1016/S0022-4049(99)00184-X.
- [6] Jean Bénabou. Distributors at work. Notes from a course at TU Darmstadt in June 2000, taken by Thomas Streicher, 2000. URL: <https://www2.mathematik.tu-darmstadt.de/~streicher/FIBR/DiWo.pdf>.

- [7] O. Bernardi and N. Bonichon. Intervals in Catalan lattices and realizers of triangulations. *Journal of Combinatorial Theory Series A*, 116(1):55–75, 2009. doi:10.1016/j.jcta.2008.05.005.
- [8] Nicolas Blanco. *Bifibrations of polycategories and classical multiplicative linear logic*. PhD thesis, University of Birmingham, 2023. arXiv:2305.15139.
- [9] Guillaume Boisseau, Chad Nester, and Mario Román. Cornering optics. In *Proceedings Fifth International Conference on Applied Category Theory*, 2023. doi:10.48550/arXiv.2307.15519.
- [10] Francis Borceaux. *Handbook of Categorical Algebra 1: Basic Category Theory*. Cambridge University Press, 1994.
- [11] David Callan. A combinatorial survey of identities for the double factorial. arXiv:0906.1317, June 2009.
- [12] Kaustuv Chaudhuri, Dale Miller, and Alexis Saurin. Canonical sequent proofs via multi-focusing. In G. Ausiello, J. Karhumäki, G. Mauri, and L. Ong, editors, *Fifth IFIP International Conference On Theoretical Computer Science*, pages 383–396. Springer, 2008. doi:10.1007/978-0-387-09680-3_26.
- [13] Kevin Coulembier, Ross Street, and Michel van den Bergh. Freely adjoining monoidal duals. *Mathematical Structures in Computer Science*, 31(7):748–768, 2021.
- [14] Robert Dawson, Robert Paré, and Dorette Pronk. Adjoining adjoints. *Advances in Mathematics*, 178(1):99–140, 2003. doi:10.1016/S0001-8708(02)00068-3.
- [15] Robert Dawson, Robert Paré, and Dorette Pronk. Undecidability of the free adjoint construction. *Applied Categorical Structures*, 11:403–419, 2003. doi:10.1023/A:1025712521140.
- [16] Robert Dawson, Robert Paré, and Dorette Pronk. The Span construction. *Theory and Applications of Categories*, 24(13):302–377, 2010. URL: <http://www.tac.mta.ca/tac/volumes/24/13/24-13abs.html>.
- [17] Jacopo Emmenegger, Luca Mesiti, Giuseppe Rosolini, and Thomas Streicher. A comonad for Grothendieck fibrations. *Theory and Applications of Categories*, 40(13):371–379, 2024. URL: <http://tac.mta.ca/tac/volumes/40/13/40-13abs.html>.
- [18] Jean-Yves Girard. Linear logic. *Theoretical Computer Science*, 50:1–102, 1987.
- [19] J. W. Gray. Fibred and cofibred categories. In *Proceedings of the Conference on Categorical Algebra*, pages 21–83. Springer-Verlag, 1966.
- [20] Alexander Grothendieck et al. *Revêtements étales et groupe fondamental*, volume 224 of *Lecture Notes in Mathematics*. Springer-Verlag, 1971. Séminaire de Géométrie Algébrique du Bois Marie 1960–1961 (SGA 1), dirigé par A. Grothendieck, augmenté de deux exposés de M. Raynaud.
- [21] Russell Harmer, Martin Hyland, and Paul-André Mellies. Categorical combinatorics for innocent strategies. In *22nd IEEE Symposium on Logic in Computer Science*, pages 379–388. IEEE Computer Society, July 2007. doi:10.1109/LICS.2007.14.
- [22] Lukas Heidemann, David Reutter, and Jamie Vicary. Zigzag normalisation for associative n-categories. In *37th Annual ACM/IEEE Symposium on Logic in Computer Science*. Association for Computing Machinery, 2022. doi:10.1145/3531130.3533352.
- [23] Michel Hirschowitz, André Hirschowitz, and Tom Hirschowitz. A theory for game theories. In *International Conference on Foundations of Software Technology and Theoretical Computer Science*, pages 192–203. Springer, 2007.
- [24] Bart P. F. Jacobs. *Categorical Logic and Type Theory*, volume 141 of *Studies in logic and the foundations of mathematics*. North-Holland, 2001.
- [25] P. T. Johnstone. A note on discrete Conduché fibrations. *Theory and Applications of Categories*, 5(1):1–11, 1999. doi:10.1017/S03050004108001345.
- [26] André Joyal. Disks, duality, and θ -categories. preprint, 1997. URL: <https://ncatlab.org/nlab/files/JoyalThetaCategories.pdf>.
- [27] Donald E Knuth and Peter B Bendix. Simple word problems in universal algebras. *Automation of Reasoning: 2: Classical Papers on Computational Logic 1967–1970*, pages 342–376, 1983.
- [28] Anders Kock. Fibrations as Eilenberg-Moore algebras. arXiv:1312.1608, December 2013.
- [29] Joachim Kock. Weak identity arrows in higher categories. *International Mathematics Research Papers*, 2006:1–54, 2006.
- [30] Joachim Kock. On ternary factorization systems. Comment on an n-Category Café blog post, 2010. URL: https://golem.ph.utexas.edu/category/2010/07/ternary_factorization_systems.html#c034162.
- [31] Germain Lamarche. Sur les partitions non-croisées d’un cycle. *Discrete Mathematics*, 1(4):333–350, 1971.
- [32] François Lamarche. Path functors in Cat. Unpublished manuscript, 2010. URL: <https://hal.inria.fr/hal-00831430>.

- [33] François Lamarche. Modeling Martin-Löf type theory in categories. *Journal of Applied Logic*, 12:28–44, 2014. doi:10.1016/j.jal.2013.08.003.
- [34] Joachim Lambek. The mathematics of sentence structure. *The American Mathematical Monthly*, 65(3):154–170, 1958. URL: <https://www.jstor.org/stable/2310058>.
- [35] Joachim Lambek. Deductive systems and categories I: Syntactic calculus and residuated categories. *Math. Syst. Theory*, 2(4):287–318, 1968. doi:10.1007/BF01703261.
- [36] F. W. Lawvere. Equality in hyperdoctrines and comprehension schema as an adjoint functor. In *Proceedings of the AMS Symposium on Pure Mathematics XVII*, pages 1–14, 1970.
- [37] Daniel R. Licata, Michael Shulman, and Mitchell Riley. A fibrational framework for substructural and modal logics. In Dale Miller, editor, *2nd International Conference on Formal Structures for Computation and Deduction*, volume 84 of *LIPICs*, pages 25:1–25:22. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, September 2017. doi:10.4230/LIPICs.FSCD.2017.25.
- [38] Saunders Mac Lane. *Categories for the Working Mathematician*. Springer, 1998.
- [39] Paul-André Mellies. Game semantics in string diagrams. In *27th Annual IEEE Symposium on Logic in Computer Science*, pages 481–490. IEEE, 2012. doi:10.1109/LICS.2012.58.
- [40] Paul-André Mellies and Noam Zeilberger. Functors are type refinement systems. In *42nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 3–16, January 2015. doi:10.1145/2676726.2679670.
- [41] Paul-André Mellies and Noam Zeilberger. A bifibrational reconstruction of Lawvere’s presheaf hyperdoctrine. In *31st Annual IEEE Conference on Logic in Computer Science*, pages 555–564, July 2016. doi:10.1145/2933575.2934525.
- [42] Paul-André Mellies and Noam Zeilberger. An Isbell duality theorem for type refinement systems. *Mathematical Structures in Computer Science*, 28(6):736–774, 2018. doi:10.1017/S0960129517000068.
- [43] David Jaz Myers. String diagrams for double categories and equipments. arXiv:1612.02762, March 2018.
- [44] Chad Nester. The structure of concurrent process histories. In Ferruccio Damiani and Ornela Dardha, editors, *Coordination Models and Languages*, pages 209–224. Springer, 2021. doi:10.1007/978-3-030-78142-2_13.
- [45] Stiéphen Pradal. On combinatorial aspects of fat Delta. arXiv:2506.01717, June 2025.
- [46] Qingchun Ren. Ordered partitions and drawings of rooted plane trees. *Discrete Mathematics*, 338(1):1–9, 2015. doi:10.1016/j.disc.2014.08.014.
- [47] Christian Sattler. Kock’s fat Δ is a direct replacement of Δ . Unpublished note, 2017. URL: <https://www.cse.chalmers.se/~sattler/docs/fat-delta.pdf>.
- [48] Stephen Schanuel and Ross Street. The free adjunction. *Cahiers de topologie et géométrie différentielle catégoriques*, 27(1):81–83, 1986. URL: http://www.numdam.org/item/CTGDC_1986__27_1_81_0/.
- [49] Michael Shulman. Framed bicategories and monoidal fibrations. *Theory and Applications of Categories*, 20(18):650–738, 2008. URL: <http://www.tac.mta.ca/tac/volumes/20/18/20-18abs.html>.
- [50] Michael Shulman. LNL polycategories and doctrines of linear logic. *Logical Methods in Computer Science*, 19(2), 2023. doi:10.46298/lmcs-19(2:1)2023.
- [51] Robert J. Simmons. Structural focalization. *ACM Trans. Comput. Log.*, 15(3):21:1–21:33, 2014. doi:10.1145/2629678.
- [52] Richard P. Stanley. *Enumerative Combinatorics: Volumes 1 and 2*. Cambridge University Press, 1999.
- [53] Tarmo Uustalu, Niccolò Veltri, and Cheng-Syuan Wan. Proof theory of skew non-commutative MILL. In Andrzej Indrzejczak and Michal Zawidzki, editors, *10th International Conference on Non-Classical Logics*, volume 358 of *EPTCS*, pages 118–135, March 2022. doi:10.4204/EPTCS.358.9.
- [54] Niccolò Veltri. Maximally multi-focused proofs for skew non-commutative MILL. In Helle Hvid Hansen, Andre Scedrov, and Ruy J. G. B. de Queiroz, editors, *29th International Workshop on Logic, Language, Information, and Computation*, volume 13923 of *Lecture Notes in Computer Science*, pages 377–393. Springer, July 2023. doi:10.1007/978-3-031-39784-4_24.
- [55] Christian Williams. *The Metalanguage of Category Theory*. PhD thesis, UC Riverside, 2023. URL: <https://escholarship.org/uc/item/84j4z67h>.

APPENDIX A. SOME MORE DETAILED PROOFS

A.1. Section 1.

Lemma 1.9. *Cut is compatible with permutation equivalence. If $\alpha_1 \sim \alpha_2$ then $\alpha_1 \cdot \beta \sim \alpha_2 \cdot \beta$, and $\beta \cdot \alpha_1 \sim \beta \cdot \alpha_2$,*

Proof. By symmetry it suffices to prove that $\alpha_1 \cdot \beta \sim \alpha_2 \cdot \beta$. The proof is by case analysis on the proof of $\alpha_1 \sim \alpha_2$: the equivalence and congruence cases are direct induction, the more delicate cases are the generating equations, which require an induction on β .

Case $\bar{f}.\alpha.h \sim (\bar{f}.\alpha).h$. To reason about $((\bar{f}.\alpha).h) \cdot \beta$, we perform a case analysis on β , depending on whether it starts with a left rule or a right rule. (It cannot be an atom as its domain is of the form $h^+ U$.)

If β starts with a left rule and its domain is of the form $h^+ U$, we have $\beta = h \setminus_i \beta'$ and

$$\begin{aligned} & (\bar{f}.\alpha.h) \cdot \beta \\ \sim & \bar{f}.\alpha.h \cdot \beta \\ = & \bar{f}.\alpha.h \cdot (h \setminus_i \beta') \\ = & \bar{f}.\alpha \cdot \beta' \\ \sim & (\bar{f}.\alpha) \cdot \beta' \\ = & ((\bar{f}.\alpha).h) \cdot (h \setminus_i \beta') \\ = & ((\bar{f}.\alpha).h) \cdot \beta \end{aligned}$$

If β starts with a right rule, let us assume that it is a division $\beta = \beta' \setminus_{i/\bar{j}}$ —the multiplication case is similar. Then we have

$$\begin{aligned} & (\bar{f}.\alpha.h) \cdot \beta \\ = & (\bar{f}.\alpha.h) \cdot (\beta' \setminus_{i/\bar{j}}) \\ \sim & ((\bar{f}.\alpha.h) \cdot \beta') \setminus_{fghi/\bar{j}} \\ \sim & (((\bar{f}.\alpha).h) \cdot \beta') \setminus_{fghi/\bar{j}} \quad \text{by induction hypothesis } (\beta' < \beta) \\ = & ((\bar{f}.\alpha).h) \cdot (\beta' \setminus_{i/\bar{j}}) \\ = & ((\bar{f}.\alpha).h) \cdot \beta \end{aligned}$$

Case $f \setminus_g \alpha.h \sim f \setminus_{gh} (\alpha.h)$. This proof is in fact identical to the previous one: the proof above only depends on the fact $\bar{f}.\alpha.h$ has a multiplication on the right, not on the operation performed on the left.

Case $(\bar{f}.\alpha) \setminus_{fg/\bar{h}} \sim \bar{f}.\alpha \setminus_{g/\bar{h}}$. To reason about $(\alpha \setminus_{g/\bar{h}}) \cdot \beta$, we again reason by case analysis on β , depending on whether it is a multiplication by h on the left or a right-rule.

If $\beta = \bar{h}.\beta'$, then we have

$$\begin{aligned} & ((\bar{f}.\alpha) \setminus_{fg/\bar{h}}) \cdot \beta \\ = & ((\bar{f}.\alpha) \setminus_{fg/\bar{h}}) \cdot (\bar{h}.\beta') \\ = & (\bar{f}.\alpha) \cdot \beta' \\ \sim & \bar{f}.\alpha \cdot \beta' \\ = & \bar{f}.\alpha \setminus_{g/\bar{h}} \cdot (\bar{h}.\beta') \\ = & \bar{f}.\alpha \setminus_{g/\bar{h}} \cdot \beta \\ \sim & (\bar{f}.\alpha \setminus_{g/\bar{h}}) \cdot \beta \end{aligned}$$

If β starts with a right rule, for example $\beta = \beta'_{i/\bar{j}}$, then we have

$$\begin{aligned}
& ((\bar{f}.\alpha)_{fg/\bar{h}}) \cdot \beta \\
&= ((\bar{f}.\alpha)_{fg/\bar{h}}) \cdot (\beta'_{i/\bar{j}}) \\
&= (((\bar{f}.\alpha)_{fg/\bar{h}}) \cdot \beta')_{fgi/\bar{j}} \\
&\sim (\bar{f}.\alpha_{g/\bar{h}}) \cdot \beta'_{fgi/\bar{j}} && \text{by induction hypothesis } (\beta' < \beta) \\
&\sim (\bar{f}.\alpha_{g/\bar{h}}) \cdot (\beta'_{i/\bar{j}}) \\
&= (\bar{f}.\alpha_{g/\bar{h}}) \cdot \beta
\end{aligned}$$

Case $(f \setminus_{gh} \alpha)_{g/\bar{h}} \sim f \setminus_g (\alpha_{fg/\bar{h}})$. Again, this proof is identical to the previous case, as they both have a division on the right of α . \square

Lemma 1.11 (η -expansion).

For any $\alpha : f^+ S \Longrightarrow_g T$, for any $\beta : S \Longrightarrow_f g^- T$, we have:

$$\alpha \sim f \setminus_g (f_S \cdot \alpha) \qquad \beta \sim (\beta \cdot \bar{g}_T) f / \bar{g}$$

The proof of this equations is included in the proof of the following lemma.

Lemma 1.10 (identity is neutral). *The identity derivations are neutral elements for cut:*

$$\text{id}_S \cdot \alpha \sim \alpha \sim \alpha \cdot \text{id}_S$$

Proof. By induction on S . We prove the left neutrality case, right neutrality is symmetric.

Case X. If α has an atomic domain or codomain, it must be an atomic derivation δ :

$$\text{id}_X \cdot \alpha = (\text{id}_X) \cdot \delta = (\text{id}_X \delta) = \delta = \alpha$$

Case $f^+ S$.

$$\begin{aligned}
& \text{id}_{f^+ S} \cdot \alpha \\
&= (f \setminus_{\text{id}_B} f_S) \cdot \alpha \\
&\sim f \setminus_g (f_S \cdot \alpha)
\end{aligned}$$

We now prove that any α is permutation equivalent to $f \setminus_g (f_S \cdot \alpha)$ by induction on α .

If α starts with a left rule, we must have $\alpha = f \setminus_g \alpha'$, and then:

$$\begin{aligned}
& f \setminus_g (f_S \cdot \alpha) \\
&= f \setminus_g ((\text{id}_S \cdot f) \cdot (f \setminus_g \alpha')) \\
&= f \setminus_g (\text{id}_S \cdot \alpha') \\
&\sim f \setminus_g \alpha' && \text{by induction hypothesis on } S \\
&= \alpha
\end{aligned}$$

If α starts with a right rule, for example $\alpha = \alpha'_{g/\bar{h}}$, we have

$$\begin{aligned}
& f \setminus_g (f_S \cdot \alpha) \\
&= f \setminus_g (f_S \cdot (\alpha'_{g/\bar{h}})) \\
&= f \setminus_g ((f_S \cdot \alpha')_{g/\bar{h}}) \\
&\sim (f \setminus_g (f_S \cdot \alpha'))_{g/\bar{h}} \\
&\sim \alpha'_{g/\bar{h}} && \text{by induction hypothesis on } \alpha' \\
&= \alpha
\end{aligned}$$

Case $g^- T$.

$$\begin{aligned} & \text{id}_{g^- T} \cdot \alpha \\ &= (\bar{g}_T \text{id}_B / \bar{g}) \cdot \alpha \end{aligned}$$

We prove that any α is permutation equivalent to $(\bar{g}_T \text{id}_B / \bar{g}) \cdot \alpha$ by induction on α . If α starts with a left rule, we must have $\alpha = \bar{g} \cdot \alpha'$, and so

$$\begin{aligned} & (\bar{g}_T \text{id}_B / \bar{g}) \cdot \alpha \\ &= ((\bar{g} \cdot \text{id}_T) \text{id}_B / \bar{g}) \cdot (\bar{g} \cdot \alpha') \\ &= (\bar{g} \cdot \text{id}_T) \cdot \alpha' \\ &\sim \bar{g} \cdot (\text{id}_T \cdot \alpha') \\ &\sim \bar{g} \cdot \alpha' && \text{by induction hypothesis on } T \\ &= \alpha \end{aligned}$$

If α starts with a right rule, for example $\alpha = \alpha'_{h/\bar{i}}$, we have

$$\begin{aligned} & (\bar{g}_T \text{id}_B / \bar{g}) \cdot \alpha \\ &= (\bar{g}_T \text{id}_B / \bar{g}) \cdot (\alpha'_{h/\bar{i}}) \\ &= ((\bar{g}_T \text{id}_B / \bar{g}) \cdot \alpha')_{h/\bar{i}} \\ &\sim \alpha'_{h/\bar{i}} && \text{by induction on } \alpha' \\ &= \alpha \end{aligned}$$

□

Lemma 1.12. *Cut is associative: $(\alpha \cdot \beta) \cdot \gamma \sim \alpha \cdot (\beta \cdot \gamma)$.*

Proof. The proof is by induction on α, β, γ .

If α is an atomic derivation, then β and γ must be as well—and conversely—and then associativity is immediate:

$$(\delta \cdot \epsilon) \cdot \zeta = (\delta \epsilon \zeta) = \delta \cdot (\epsilon \cdot \zeta)$$

If α starts with a left rule, then associativity is immediate. For example if $\alpha = \bar{f} \cdot \alpha'$, we have

$$\begin{aligned} & (\alpha \cdot \beta) \cdot \gamma \\ &= ((\bar{f} \cdot \alpha') \cdot \beta) \cdot \gamma \\ &\sim (\bar{f} \cdot (\alpha' \cdot \beta)) \cdot \gamma \\ &\sim \bar{f} \cdot ((\alpha' \cdot \beta) \cdot \gamma) \\ &\sim \bar{f} \cdot (\alpha' \cdot (\beta \cdot \gamma)) && \text{by induction hypothesis } (\alpha' < \alpha) \\ &\sim (\bar{f} \cdot \alpha') \cdot (\beta \cdot \gamma) \\ &= \alpha \cdot (\beta \cdot \gamma) \end{aligned}$$

Symmetrically, associativity is immediate if γ starts with a right rule. When neither of those two easy cases apply, we know that α starts with a right rule and γ with a left rule. We then reason by case analysis on β .

If β starts with a multiplication on the left, $\beta = \bar{g} \cdot \beta'$, then the right rule of α must end a matching division, $\alpha = \alpha' \cdot_f \bar{g}$, and we have

$$\begin{aligned}
& (\alpha \cdot \beta) \cdot \gamma \\
= & ((\alpha' \cdot_f \bar{g}) \cdot (\bar{g} \cdot \beta')) \cdot \gamma \\
\sim & (\alpha' \cdot \beta') \cdot \gamma \\
\sim & \alpha' \cdot (\beta' \cdot \gamma) && \text{by induction hypothesis } (\alpha' < \alpha, \beta' < \beta) \\
\sim & (\alpha' \cdot_f \bar{g}) \cdot (\bar{g} \cdot (\beta' \cdot \gamma)) \\
\sim & (\alpha' \cdot_f \bar{g}) \cdot ((\bar{g} \cdot \beta') \cdot \gamma) \\
= & \alpha \cdot (\beta \cdot \gamma)
\end{aligned}$$

The proof is similar if β starts with a division on the left, $\beta = g \setminus_h \beta'$ and $\alpha = \alpha' \cdot g$:

$$\begin{aligned}
& (\alpha \cdot \beta) \cdot \gamma \\
= & ((\alpha' \cdot g) \cdot (g \setminus_h \beta')) \cdot \gamma \\
\sim & \alpha' \cdot (\beta' \cdot \gamma) && \text{by induction hypothesis } (\alpha' < \alpha, \beta' < \beta) \\
\sim & (\alpha' \cdot g) \cdot (g \setminus_{hi} (\beta' \cdot \gamma)) \\
\sim & (\alpha' \cdot g) \cdot ((g \setminus_h \beta') \cdot \gamma) \\
= & \alpha \cdot (\beta \cdot \gamma)
\end{aligned}$$

The cases where β starts with a right rule are symmetrical. □

A.2. Section 3.

Proposition 3.16.

$$\begin{aligned}
& \forall \alpha_w, \quad \alpha_w \sim \text{InvSeq}(\mathbf{S}(\alpha_w)) \\
& \forall \alpha_w, \quad \forall \alpha'_w \in \text{InvSeq}(\mathbf{S}(\alpha_w)), \quad \alpha_w \sim \alpha'_w
\end{aligned}$$

Proof. By induction on the underlying judgment of α_w .

- If α_w has a non-neutral judgment $\alpha_w : \pi^+ N \Longrightarrow_f \bar{\rho}^- P$, we have

$$\mathbf{S}(\alpha_w) := \pi \setminus \mathbf{S}(\pi_N \cdot \alpha_w \cdot \bar{\rho}_P) / \bar{\rho}$$

For any $\alpha'_w \in \text{InvSeq}(\mathbf{S}(\alpha_w))$, α'_w is one of the two sequentializations of the bi-inversion $\pi \setminus \beta_w / \bar{\rho}$ for some $\beta_w \in \text{InvSeq}(\mathbf{S}(\pi_N \cdot \alpha_w \cdot \bar{\rho}_P : N \Longrightarrow_{\pi f \rho} P))$. Let us show the case where $\alpha'_w := \pi \setminus (\beta_w / \bar{\rho})$, the other case is similar. By induction hypothesis on the smaller judgment $N \Longrightarrow_{\pi f \rho} P$, we have $\beta_w \sim (\pi_N \cdot \alpha_w \cdot \bar{\rho}_P)$, and therefore

$$\begin{aligned}
\alpha'_w &= \pi \setminus (\beta_w / \bar{\rho}) \\
&\sim \pi \setminus ((\pi_N \cdot \alpha_w \cdot \bar{\rho}_P) / \bar{\rho}) \\
&\sim \pi \setminus (\pi_N \cdot \alpha_w) \\
&\sim \alpha_w
\end{aligned}$$

where the last two equivalences come from Lemma 1.11.

- Otherwise α_w has a neutral judgment $\alpha_w : N \Longrightarrow_f P$, and we reason as in the definition of $\mathbf{S}(\alpha_w)$ by case analysis on α_w , which can be of the form $\bar{\sigma} \cdot \beta_w$ or $\beta_w \cdot \tau$. For example in the case $\bar{\sigma} \cdot \beta_w$ (the other case is symmetric), N is of the form $\sigma^- Q$ and we have

$$\text{InvSeq}(\mathbf{S}(\alpha_w)) = \text{InvSeq}(\mathbf{S}(\bar{\sigma} \cdot \beta_w)) = \text{InvSeq}(\bar{\sigma} \cdot \mathbf{S}(\beta_w))$$

For any $\alpha'_w \in \text{InvSeq}(\mathbf{S}(\alpha_w))$, α'_w is of the form $\bar{\sigma}.\beta'_w$ for some $\beta'_w \in \text{InvSeq}(\mathbf{S}(\beta_w))$. We have $\beta'_w \sim \beta_w$ by induction hypothesis on the smaller judgment $Q \implies P$, and therefore

$$\alpha'_w = \bar{\sigma}.\beta'_w \sim \bar{\sigma}.\beta_w = \alpha_w$$

□

Proposition 3.17.

$$\forall \alpha_m, \quad \alpha_m \sim_{\text{seq}} \mathbf{S}(\text{InvSeq}(\alpha_m)) \\ \forall \alpha_m, \forall \alpha_w \in \text{InvSeq}(\alpha_m), \quad \alpha_m \sim_{\text{seq}} \mathbf{S}(\alpha_w)$$

Proof. The statement can in fact be strengthened into the following:

$$\forall \alpha_w \in \text{InvSeq}(\alpha_m), \quad \alpha_m = \mathbf{S}(\alpha_w)$$

This is proved by induction on the judgment of α_w and α_m :

- At a non-neutral judgment $\pi^+ N \implies_f \rho^- P$, α_m must be of the form $\pi \setminus \beta_m / \bar{\rho}$ with $\beta_m : N \implies_{\pi f \rho} P$, and $\alpha_w \in \text{InvSeq}(\alpha_m)$ must be a sequentialization of the bi-inversion $\pi \setminus \beta_w / \bar{\rho}$ for some $\beta_w \in \text{InvSeq}(\beta_m)$. By induction hypothesis we have $\mathbf{S}(\beta_w) = \beta_m$. If α_w is $\pi \setminus (\beta_w / \bar{\rho})$ (the other case is similar), we have

$$\begin{aligned} & \mathbf{S}(\alpha_w) \\ &= \pi \setminus \mathbf{S}(\pi_N \cdot \alpha_w \cdot \bar{\rho}_P) / \bar{\rho} \\ &= \pi \setminus \mathbf{S}(\pi_N \cdot \pi \setminus (\beta_w / \bar{\rho}) \cdot \bar{\rho}_P) / \bar{\rho} \\ &= \pi \setminus \mathbf{S}(\beta_w / \bar{\rho} \cdot \bar{\rho}_P) / \bar{\rho} \\ &= \pi \setminus \mathbf{S}(\beta_w) / \bar{\rho} \\ &= \pi \setminus \beta_m / \bar{\rho} \\ &= \alpha_m \end{aligned}$$

- At a neutral judgment $N \implies_f P$, we reason by case analysis on α_m . For example if it is $\bar{\sigma}.\beta_m$ (the other case $\beta_m.\tau$ is similar), we have $\alpha_w = \bar{\sigma}.\beta_w$ for $\beta_w \in \text{InvSeq}(\beta_m)$. We have $\mathbf{S}(\beta_w) = \beta_m$ by induction hypothesis, so

$$\mathbf{S}(\alpha_w) = \mathbf{S}(\bar{\sigma}.\beta_w) = \bar{\sigma}.\mathbf{S}(\beta_w) = \bar{\sigma}.\beta_m = \alpha_m$$

□

Theorem 3.21 (Local confluence). *The rewrite system ($\text{par} \cup \text{gra}$) is locally confluent: if $\beta_m \leftarrow \alpha_m \rightarrow \beta'_m$ then there exists γ_m such that $\beta_m \rightarrow^* \gamma_m \leftarrow^* \beta'_m$.*

Proof. The source of each rewrite rule must be a stack of two bipoles. Thus all critical pairs (where the source of two rewrite rules overlap) can be found by considering stacks of two or three bipoles. Let us first enumerate all such stack shapes to determine which may contain critical pairs.

Enumerating all possible critical pairs. We exploit a symmetry that permutes L and R rules to reduce the number of cases: for example, the reduction behavior of $\mathbb{L} \mathbb{R}$ stacks is symmetric to the behavior of $\mathbb{R} \mathbb{L}$ stacks.

Let us consider all stacks of two bipoles. They have one of the following shapes:

- (1) $\mathbb{L} \mathbb{L}$: no reduction.
- (2) $\mathbb{L} \mathbb{R}$: *a par_R^L rule may be applied.*
- (3) $\mathbb{L} \mathbb{L} \mathbb{R}$: no reduction.
- (4) $\mathbb{R} \mathbb{L}$: symmetric to (2).
- (5) $\mathbb{R} \mathbb{R}$: symmetric to (1), no reduction.
- (6) $\mathbb{R} \mathbb{L} \mathbb{R}$: symmetric to (3), no reduction.
- (7) $\mathbb{L} \mathbb{R} \mathbb{L}$: *a gra^L rule may be applied.*
- (8) $\mathbb{L} \mathbb{R} \mathbb{R}$: symmetric to (7).
- (9) $\mathbb{L} \mathbb{R} \mathbb{L} \mathbb{R}$: no reduction.

We next enumerate the stacks of three bipoles. Notice that if a 2-bipole stack has no reduction, then a 3-bipole stack formed by extending the 2-stack with a new bipole above or below cannot form a new critical pair not already in the above enumeration. For example, there is no reduction from $\mathbb{L} \mathbb{L}$, so extending it into $\mathbb{L} \mathbb{L} \mathbb{R}$ (or $\mathbb{L} \mathbb{L} \mathbb{L} \mathbb{R}$, or $\mathbb{R} \mathbb{L} \mathbb{L}$ for example) does not produce critical pairs involving all three bipoles. Any critical pair in an extension is a critical pair of 2-bipole stacks, which we have already considered.


In other words, we only need to consider extensions of the reducible 2-bipole stacks, modulo symmetry, to wit (2) $\mathbb{L} \mathbb{R}$ and (7) $\mathbb{L} \mathbb{R} \mathbb{L}$ above:

- (2.a) $\mathbb{L} \mathbb{R} \mathbb{L}$: *two overlapping par rules may be applied.*
- (2.b) $\mathbb{L} \mathbb{R} \mathbb{R}$: extension of no-reduction (5).
- (2.c) $\mathbb{L} \mathbb{R} \mathbb{L} \mathbb{R}$: extension of no-reduction (6).
- (2.d) $\mathbb{L} \mathbb{L} \mathbb{R}$: extension of no-reduction (1).
- (2.e) $\mathbb{R} \mathbb{L} \mathbb{R}$: symmetric to (2.a).
- (2.f) $\mathbb{L} \mathbb{R} \mathbb{L} \mathbb{R}$: *a gra_L rule may be applied above, and a par_R^L below.*
- (7.a) $\mathbb{L} \mathbb{R} \mathbb{L} \mathbb{L}$: extension of no-reduction (1).
- (7.b) $\mathbb{L} \mathbb{R} \mathbb{L} \mathbb{R}$: same as (2.f).
- (7.c) $\mathbb{L} \mathbb{R} \mathbb{L} \mathbb{L} \mathbb{R}$: extension of no-reduction (3).
- (7.d) $\mathbb{L} \mathbb{L} \mathbb{R} \mathbb{L}$: extension of no-reduction (3).
- (7.e) $\mathbb{R} \mathbb{L} \mathbb{R} \mathbb{L}$: extension of no-reduction (6).
- (7.f) $\mathbb{L} \mathbb{R} \mathbb{L} \mathbb{R} \mathbb{L}$: extension of no-reduction (9).

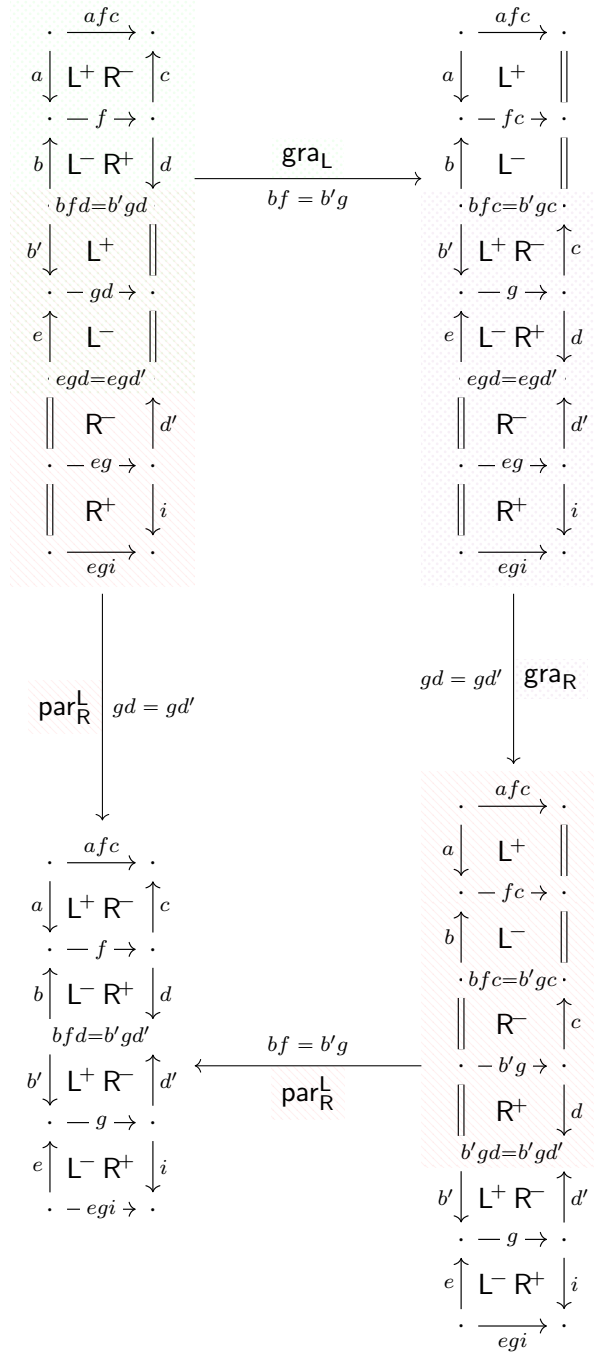
Resolving the critical pairs. We now consider each of the above cases in turn.

- (2) $[\mathbb{L} \mathbb{R}]$ Two reductions from stacks of this shape must be two applications of the rule par_R^L , which is deterministic under the FP condition, so the two right-hand-sides are equal.
- (7) $[\mathbb{L} \mathbb{R} \mathbb{L}]$ As above, the critical pair is trivially resolved because gra_L is deterministic, under no assumptions.

(2.a) $[\mathbb{L} \mathbb{R} \mathbb{L}]$ This is exactly the critical pair (3.6) that we resolved by introducing the gravity rules, in this case $\text{gra}_{\mathbb{L}}$:



(2.f) $[\mathbb{L} \mathbb{R} \mathbb{L} \mathbb{R}]$ This critical pair between $\text{par}_{\mathbb{R}}^{\mathbb{L}}$ and $\text{gra}_{\mathbb{L}}$ is resolved by composing the latter with $\text{gra}_{\mathbb{R}}$ followed by $\text{par}_{\mathbb{R}}^{\mathbb{L}}$:



□

APPENDIX B. CUT ON MULTI-FOCUSED DERIVATIONS, DIRECTLY

In Section 3.4.5 we defined horizontal composition of strongly multifocused proofs, and this definition is modulo permutation equivalence. We can also give a second, more syntactic definition of cut, directly on strongly multifocused derivations, and show that it coincides (modulo permutation equivalence) with the first definition.

Let us first define the operations $(\pi_N : \alpha_m)$, $(\alpha_m : \bar{\rho}_P)$, which are equivalent to compositions $\pi_N \cdot \alpha_m$, $\alpha_m \cdot \bar{\rho}_P$ but can be defined by a simple case analysis—we use a colon to emphasize that this is a more primitive operation:

$$\begin{aligned} (- : \bar{\rho}_P) & : (S \Longrightarrow_f \rho^- P) \rightarrow (S \Longrightarrow_{f\rho} P) \\ (\beta \ /_f \bar{\rho}) : \bar{\rho}_P & := \beta \\ (\pi \ \backslash \ \beta \ / \bar{\rho}) : \bar{\rho}_P & := \pi \ \backslash_{f\rho} \beta \\ \\ (\pi_N : -) & : (\pi^+ N \Longrightarrow_f T) \rightarrow (N \Longrightarrow_{\pi f} T) \\ \pi_N : (\pi \ \backslash \ \beta) & := \beta \\ \pi_N : (\pi \ \backslash \ \beta \ / \bar{\rho}) & := \beta \ \pi_f / \bar{\rho} \end{aligned}$$

In the case of $(\alpha_m : \bar{\rho}_P)$ for example, the definition is by case analysis, but we know by inspection of the judgment that a strongly multifocused derivation of $\alpha_m : S \Longrightarrow_f \rho^- P$ can only start with an inversion or bi-inversion rule.

We can now define the general cut of two strongly multifocused formulas, from $S \Longrightarrow_f U$ and $U \Longrightarrow_f T$ to $S \Longrightarrow_{fg} T$. The definition is first by case analysis on S and T , with a single case if S is positive or T is negative, and then when $S \Longrightarrow_{fg} T$ is a neutral judgment $N \Longrightarrow_{fg} P$ we reason by case analysis on α_m and β_m .

$$\begin{aligned} (- \cdot -) & : (\pi^+ N \Longrightarrow_f U) \times (U \Longrightarrow_g \rho^- P) \rightarrow (\pi^+ N \Longrightarrow_{fg} \rho^- P) \\ \alpha_m \cdot \beta_m & := \pi \ \backslash \ ((\pi_N : \alpha_m) \cdot (\beta_m : \bar{\rho}_P)) \ / \bar{\rho} \end{aligned}$$

$$\begin{aligned} (- \cdot -) & : (N \Longrightarrow_f U) \times (U \Longrightarrow_g P) \rightarrow (N \Longrightarrow_{fg} P) \\ \\ (\bar{\sigma} \cdot \alpha'_m) \cdot (\beta'_m \cdot \tau) & := \bar{\sigma} \cdot (\alpha'_m \cdot \beta'_m \cdot \tau) \\ (\bar{\sigma} \cdot \alpha'_m) \cdot \beta_m & := \bar{\sigma} \cdot (\alpha'_m \cdot \beta_m) & \beta_m \neq \beta'_m \cdot \tau \\ \alpha_m \cdot (\beta'_m \cdot \tau) & := (\alpha_m \cdot \beta'_m) \cdot \tau & \alpha_m \neq \bar{\sigma} \cdot \alpha'_m \\ \\ (\alpha'_m \cdot \pi) \cdot (\pi \ \backslash \ \beta'_m) & := \alpha'_m \cdot \beta'_m \\ (\bar{\sigma} \cdot \alpha'_m \cdot \pi) \cdot (\pi \ \backslash \ \beta'_m) & := (\bar{\sigma} \cdot \alpha'_m) \cdot \beta'_m \\ \\ (\alpha'_m \ /_f \bar{\rho}) \cdot (\bar{\rho} \cdot \beta'_m) & := \alpha'_m \cdot \beta'_m \\ (\alpha'_m \ /_f \bar{\rho}) \cdot (\bar{\rho} \cdot \beta'_m \cdot \tau) & := \alpha'_m \cdot (\beta'_m \cdot \tau) \\ \\ \delta \cdot \epsilon & := \delta \epsilon \end{aligned}$$

For example, if α_m starts with a left-focusing rule, we propagate the left-focusing rule in the output and recursively compute a cut at a strictly simpler judgment. If α_m is a right-focusing or bi-focusing rule, then necessarily β_m starts with a matching invertible rule (thanks to strong focusing) and we have a principal cut. If α_m starts with a left-inversion or bi-inversion rule, then its source formula S is positive and we are in the first case of a

non-neutral judgment. Finally, if α_m starts with a right-inversion rule, then necessarily β_m uses a focusing rule.

Proposition B.1. *The explicit/intensional definition of $\alpha_m \cdot \beta_m$ above agrees with the extensional definition we gave earlier modulo (\sim_{seq}) as $\mathbf{S}(\text{Seq}(\alpha_m) \cdot \text{Seq}(\beta_m))$.*

Proof. For each equation that defines $\alpha_m \cdot \beta_m$, we check that there is a choice of sequentialization of α_m, β_m that gives the expected result.

For example, in the non-neutral case, $\alpha_m : \pi^+ N \Longrightarrow_f U$ must be of the form $\pi \setminus_f \alpha'_m$ or $\pi \setminus_f \alpha'_m / \bar{\tau}$. If we have the latter form, we choose to sequentialize it into $\pi \setminus_f (\alpha'_{\pi f} / \bar{\tau})$,

so that both cases can be sequentialized into a term of the form $\pi \setminus_f \alpha'_w$, where α'_w is a sequentialization of $\pi_N : \alpha_m$. By the same reasoning β_m is sequentialized into a term of the form $\beta'_w g / \bar{\rho}$, with β'_w a sequentialization of $\beta_m : \bar{\rho}_P$. Their composition as weakly focused proofs is $(\pi \setminus_f \alpha'_w) \cdot (\beta'_w g / \bar{\rho})$, which is equal to either of the two permutation-equivalent proofs $\pi \setminus_{fg} ((\alpha'_w \cdot \beta'_w) \pi_{fg} / \bar{\rho})$ or $(\pi \setminus_{fg} (\alpha'_w \cdot \beta'_w))_{fg} / \bar{\rho}$. Both choices strengthen into $\pi \setminus_{fg} (\alpha'_w \cdot \beta'_w) / \bar{\rho}$, which by induction hypothesis is $\pi \setminus_{fg} (\pi_N : \alpha_m) \cdot (\beta_m : \bar{\rho}_P) / \bar{\rho}$ as desired.

In a representative neutral case, $(\bar{\sigma} \cdot \alpha'_m \cdot \pi) \cdot (\pi \setminus_g \beta'_m)$, α'_m starts with a bi-inversion rule, let us assume that it is of the form $\tau \setminus_h \alpha''_m / \bar{\rho}$. We sequentialize α_m into

$$((\bar{\sigma} \cdot (\tau \setminus_{h\rho} \alpha''_w)) \sigma h / \bar{\rho}) \cdot \pi$$

where α''_w is a sequentialization of α''_m , and we have

$$\begin{aligned} & ((\bar{\sigma} \cdot (\tau \setminus_{h\rho} \alpha''_w)) \sigma h / \bar{\rho}) \cdot \pi \cdot (\pi \setminus_h \beta'_w) \\ &= ((\bar{\sigma} \cdot (\tau \setminus_{h\rho} \alpha''_w)) \sigma h / \bar{\rho}) \cdot \beta'_w \\ &\sim (\bar{\sigma} \cdot ((\tau \setminus_{h\rho} \alpha''_w) h / \bar{\rho})) \cdot \beta'_w \\ &\sim \bar{\sigma} \cdot (((\tau \setminus_{h\rho} \alpha''_w) h / \bar{\rho}) \cdot \beta'_w) \end{aligned}$$

where $((\tau \setminus_{h\rho} \alpha''_w) h / \bar{\rho})$ is a valid sequentialization of α'_m , and so this is permutation equivalent to a sequentialization of $\bar{\sigma} \cdot (\alpha_m \cdot \beta'_m)$ as desired. \square

CONTENTS

Introduction	1
1. Sequent calculus	8
1.1. Formulas and unary sequents	8
1.2. Inference rules and derivations	9
1.3. A term syntax for proofs	11
1.4. Interpreting formulas and proofs in any bifibration on p	12
1.5. Permutation equivalence	12
1.6. Cut	15
1.7. The free bifibration	18
1.8. Some proof-theoretic consequences	21
1.9. Related work	23
2. The double category of zigzags	24
2.1. Deriving a double category structure from a right action	24
2.2. From inference rules to double cells	26
2.3. Aside: displayed bifibrations	28
2.4. A dagger structure on zigzags	30
2.5. Recovering the free adjoint construction	31
2.6. From double cells to string diagrams	33
2.7. Related work	34
3. Focused normal forms	35
3.1. Background: invertibility, polarity, and synthetic connectives	35
3.2. Invertibility and polarity in the bifibrational calculus	37
3.3. Strictly alternating formulas, splittings, and weak focusing	38
3.4. Multifocusing	45
3.5. Normal forms via rewriting	50
3.6. Maximal multifocusing	54
3.7. Decidability and enumeration	57
4. Free $(\mathcal{P}, \mathcal{N})$ -fibrations	61
5. The combinatorics of free bifibrations and ambifibrations	63
5.1. Reconstructing the simplex category	64
5.2. Categories of trees and walks	65
5.3. A free ambifibration over Δ	75
6. Conclusion and future directions	81
Acknowledgments	82
References	82
Appendix A. Some more detailed proofs	85
A.1. Section 1	85
A.2. Section 3	88
Appendix B. Cut on multi-focused derivations, directly	93