Herbrand Schemes for First-order Logic

Bahareh Afshari[1], Sebastian Enqvist[2], Graham E. Leigh[1]

[1]Department of Philosophy, Linguistics and Theory of Science, University of Gothenburg, Gothenburg, Sweden.
[2]Department of Philosophy, Stockholm University, Stockholm, Sweden.

Contributing authors: bahareh.afshari@gu.se; sebastian.enqvist@philosophy.su.se; graham.leigh@gu.se;

Abstract

This article provides a language-theoretic rendering of Herbrand's theorem. To each first-order proof is associated a higher-order recursion scheme that abstracts the computation of Herbrand sets obtained through Gentzen-style multicut elimination. The representation extends previous results in this area by lifting the prenex restriction on cut formulas and relaxing the cut-elimination strategies. Features of the new approach are the interpretation of cut as simple composition and contraction as 'call with current continuation'.

Keywords: Sequent Calculus, First-order Logic, Herbrand's Theorem, Cut Elimination, Multicut, Higher-order Recursion Schemes, Computational Content

MSC Classification: 03F05 , 03F07 , 03D05 , 68Q45

1 Introduction

Classical logic, in contrast to intuitionistic logic, does not have the existence property. One may classically prove an existentially quantified statement without necessarily providing an explicit witness. However, Herbrand's theorem implies a weaker form of the existence property for classical logic: if a prenex Σ_1-formula $\exists\vec{x}\varphi(\vec{x})$ is valid then there is a finite set of tuples of terms $\{\vec{t_1}, ..., \vec{t_n}\}$, called a *Herbrand set*, such that the disjunction $\varphi(\vec{t_1}) \vee \cdots \vee \varphi(\vec{t_n})$ is valid.

Herbrand's theorem can be proved via cut elimination, and supplies a form of computational content to classical sequent calculi. A direct representation of this computational content was provided in [1] wherein proofs are associated non-deterministic higher-order recursion schemes, henceforth abbreviated HORS. These so-called *Herbrand schemes* behave as abstract representations of non-deterministic programs which extract witnesses for weak quantifiers from a sequent calculus proof. In particular, when the end sequent is prenex Σ_1, the scheme associated with a proof provides a Herbrand set.

The analysis of LK via formal language theory is a fairly recent development [1–3]. The study of computational content of (classical) proofs, on the other hand, goes a long way back and has received considerable attention and success. There are different ways to ascribe computational content depending on the logic/theory of interest and the application in mind. But these all agree on the view that a classical proof not only serves as a verification tool but contains valuable algorithmic content that can be exploited.

The first line of work in the general direction of finding computational content in classical proofs stems from the idea of translating classical logic validities to intuitionistic realm and extracting the desired algorithmic content thereof. For a general overview see [4–6] and for a comprehensive account [7, 8]. Of note are Kreisel's 'no-counterexample' interpretation, the Gödel–Shoenfield functional interpretation and variant's thereof, negative translations, the Friedman A-translation and combinations of these techniques. Working directly on proofs in classical logic is considered in classical realizability [9], continuation passing style transformations à la [10], and Hilbert's epsilon calculus [11].

Our aim in this article is to give a direct representation of the witnesses implicit in a classical sequent calculus proof, and to provide an analysis of how sequent calculus proofs contribute to the computation of witnesses when composed using the cut rule. To this end we build on the HORS representation to give a refinement of the grammar-theoretic approach that respects the structural rules and inherent symmetry of classical sequent calculi. In particular, the defined Herbrand schemes ascribe computational meaning to the *cut* and *contraction* rules that are not apparent in Hilbert–Frege or natural deduction calculi, nor in common translations of classical logic into intuitionistic logic. Specifically, each Gentzen-style sequent calculus proof π is associated a higher-order grammar $\mathscr{H}(\pi)$ whose non-terminals and production rules exhibit an algorithm for constructing witness terms for the weak quantifiers in the end-sequent of π. If the end-sequent of π is Σ_1, then the language of the grammar $\mathscr{H}(\pi)$ presents a set of terms, which represents a Herbrand set for the end-sequent.

In a Herbrand scheme, the cut rule is interpreted directly as composition of the two cut premises, preserving the natural interpretation of cut from intuitionistic sequent calculus. This is possible due to the interpretation preserving the symmetry of the two sides of the sequent: the grammar ascribes the same computational meaning to a cut on a formula A as to the cut on $\neg A$ (obtained by first affecting the two cut formulas by $\neg$-rules). As a result, the Herbrand scheme for a proof ending with a cut is, in essence, the union of the Herbrand schemes assigned to the immediate subproofs. The other notable structural rule, contraction, is given an interpretation via Pierce's law $\big((A \to B) \to A\big) \to A$. Pierce's law, equivalent to excluded middle over intuitionistic logic, has played an important part in call-by-continuation programming as well as computational interpretations of classical proofs. Its appearance as a computation resolver of contraction, although not surprising, reinforces its value in this area of research.

The approach of [1], out of which the present work stems, also operates directly on sequent calculus but fails to fully respect the symmetry of the two sides in a Gentzen-style sequent (or, following the presentation of [1], the duality between formula and negation). The deficiency was most apparent in the asymmetric interpretation of the cut rule. Depending on which premise of a cut the witness extraction process would continue, the computation would either arrive at a simple composition of the two cut premises – in the way described above – or a more complex nested 'composition' involving one premise and two copies of the other. That design choice led to two shortcomings that the present work overcomes. First, an indirect result of the interpretation of cut is that the Herbrand schemes of [1] were restricted to proofs in which all cut formulas are in prenex normal form. Thus, to compute a Herbrand set from an arbitrary proof some initial processing of the proof is necessary before constructing the Herbrand scheme. Second, the witnesses extracted by the Herbrand scheme were related to cut-elimination strategies which always favoured a particular cut-premise for local cut reductions, i.e., strategies which treated cut asymmetrically.

Our main result shows that the language of a Herbrand scheme associated with a proof of a Σ_1 end sequent always corresponds to a valid Herbrand expansion of the end sequent. This can be interpreted as a soundness result for the analysis, supporting the notion that Herbrand schemes correctly represent the implicit witnesses of a proof. More precisely, the language contains all the Herbrand sets that can be extracted from cut-free normal forms of the proof. It should be remarked that we do not currently know whether the language of a Herbrand scheme is always finite; it may be that the language of a Herbrand scheme in some cases correspond to an infinite Herband disjunction. A similar situation occurs with the game-theoretic interpretation of Herbrand's theorem in [12], which derives infinitary Herbrand disjunctions from winning strategies associated with proofs and relies on a compactness argument to obtain a finite disjunction. Indeed, it seems to be currently unknown whether the set of witnesses that can be extracted from a given proof by the method of cut elimination is finite in general, since cut elimination is neither confluent nor strongly normalizing. In [1], finiteness of the language was made possible by the aforementioned restrictions to cut elimination strategies obeying certain constraints. Here our main focus is not first and foremost to extract finite Herband disjunctions, but rather to give an analysis of the constructive content implicit in classical sequent calculus proofs.

2 Sequent calculus for FOL

For the present work fix a countable first order signature $\mathscr{L}$ comprising predicates and function symbols each associated a finite natural number representing their arity. We assume $\mathscr{L}$ contains a sufficient stock of function symbols of each arity which includes a distinguished constant c. Terms in $\mathscr{L}$, henceforth *$\mathscr{L}$-terms*, are defined as usual starting from *two* denumerable sets of variables: *free* variables FV and *bound* variables BV. Symbols α, β, …, and x, y, z, … range over the two variable classes respectively. If t is an $\mathscr{L}$-term, $\mathsf{FV}(t)$ (resp. $\mathsf{BV}(t)$) denotes the set of free (resp. bound) variables occurring in t. If $\mathsf{BV}(t) = \emptyset$ we call t *open*. Symbols s, t, etc. range over $\mathscr{L}$-terms. For a bound variable x, $s[t/x]$ denotes the result of replacing all occurrences of x in s by t. In particular, $\mathsf{FV}(s[t/x]) = \mathsf{FV}(s) \cup \mathsf{FV}(t)$ if $x \in \mathsf{BV}(s)$, and $\mathsf{FV}(s[t/x]) = \mathsf{FV}(s)$ otherwise.

Formulas are constructed via the grammar:

$$A ::= P\vec{t} \mid (A \vee A) \mid \exists x A \mid \neg A$$

where P ranges over predicate symbols and $\vec{t}$ is a tuple of $\mathscr{L}$-terms whose length is the arity of P. We say the quantifier $\exists x$ *binds* x in $\exists x A$. Note that quantifiers can bind only bound variables.

Identity is omitted as an explicit predicate symbol as it can be expressed as a binary predicate with appropriate axioms. The remaining logical connectives are defined by the usual abbreviations: $\forall x A := \neg\exists x \neg A$, $A \wedge B := \neg(\neg A \vee \neg B)$ and $A \to B := \neg A \vee B$. The subformula relation is the standard one.

A formula is *open* if every occurrence of a bound variable is within the scope of a quantifier binding that variable. The *free* variables of a formula A, denoted $\mathsf{FV}(A)$, is the set of free variables occurring in A. Open formulas are identified up to α-equivalence, i.e., renaming of bound variables. Given a formula A, open term t and bound variable x we denote by $A[t/x]$ the result of replacing all free occurrences of x in A by t. Due to the separation of variables into free and bound classes, no condition is required on t and x beyond that t is open.

A formula which does not contain quantifiers is said to be *quantifier-free*. A Σ_1-*formula* is any formula with only weak quantifier occurrences where an occurrence of a quantifier $\exists x$ in A is *weak* (*strong*) if the number of negations on the path from A to the quantifier occurrence is even *(odd)*. As an example, the formula $A = \forall x \exists y((\exists z\, Pxyz) \to Qyx)$ contains one weak quantifier occurrence, marked by $\exists y$. The other two quantifier occurrences ($\forall x$ and $\exists z$) are strong.

An important operation on formulas is Skolemisation which, for present purposes, we take as the validity preserving transformation of arbitrary formulas into Σ_1 formulas by replacing strong quantifiers by function symbols. Let A be a formula in which all quantifier occurrences bind unique variables. Let S_A be the set of variables bound by strong quantifier occurrences in A. By assumption, each $x \in S_A$ uniquely identifies a strong quantifier occurrence $\exists x B$ in A. For each $x \in S_A$, let W_x be the set of variables bound by *weak* quantifiers on the path from $\exists x B$ to A. To each $x \in S_A$ is now associated a fresh function symbol f_x of arity $|W_x|$. The *skolemisation* of an open formula A, denoted A_{sk}, is the (open) Σ_1 formula obtained by removing all strong quantifiers in A and replacing every occurrence of a variable $x \in V$ by the term $f_x y_1 \cdots y_k$ where $(y_i)_i$ is some fixed enumeration of W_x.

Following the example $A = \forall x \exists y((\exists z\, Pxyz) \to Qyx)$ the set of strongly quantified variables is $S_A = \{x, z\}$. The associated weak quantifiers are $W_x = \emptyset$ and $W_z = \{y\}$, inducing function symbols f_x of arity 0 and f_z of arity 1 respectively. The skolemisation of A is then the Σ_1 formula

$$A_{\mathrm{sk}} = \exists y(P f_x y (f_z y) \to Q y f_x).$$

The following is immediate.

Proposition 1. *An open formula A is valid iff its skolemisation is.*

We employ a sequent calculus for classical predicate logic with explicit structural rules, namely a version of the calculus $\mathsf{G1c} + \mathsf{cut}$ of [13] for our chosen syntax. Thus, a *sequent* is a pair of finite sequences of open formulas, written $\Gamma \Rightarrow \Delta$. For $\Gamma = A_0, \ldots, A_k$ we set $\Gamma^\neg = \neg A_0, \ldots, \neg A_k$. We refer to $\Gamma \Rightarrow \Delta$ as a Σ_1-sequent if $\Gamma^\neg, \Delta \subseteq \Sigma_1$. The length of a sequent $\Gamma \Rightarrow \Delta$ is $|\Gamma \Rightarrow \Delta| \coloneqq |\Gamma| + |\Delta|$.

For $S = \Gamma \Rightarrow \Delta$ and $i < |S|$ we employ the shorthand $\Gamma \Rightarrow_i \Delta$ for the result of shifting the i-th formula in the sequent to right-hand side and negating the formula in the case $i < |\Gamma|$. That is,

$$\Gamma \Rightarrow_i \Delta \coloneqq \begin{cases} \Gamma \Rightarrow \Delta', \Delta'', A, & \text{if } i \geq |\Gamma|,\ \Delta = \Delta' A \Delta'' \text{ and } |\Gamma\Delta'| = i, \\ \Gamma', \Gamma'' \Rightarrow \Delta, \neg A, & \text{if } i < |\Gamma|,\ \Gamma = \Gamma' A \Gamma'' \text{ and } |\Gamma'| = i. \end{cases}$$

The rules of the calculus are listed in Figure 1 with the standard variable condition applying to $\exists\mathsf{L}$: The variable α, referred to as the *eigenvariable* of the inference, does not occur in the conclusion sequent $\exists x A, \Gamma \Rightarrow \Delta$. Note the non-standard form of the *exchange* rules eL and eR which move a formula occurrence in the premise sequent to the extreme. These inferences will often be omitted for brevity.

The first and final formula in a sequent are said to be *active* in the sequent, i.e., A and B are the active formulas in $A, \Gamma \Rightarrow \Delta, B$. All instances of the logical rules alter only a single active formula in each premise. The distinguished active formula A in the two premises of the cut rule is called the *cut formula*. For every rule instance one formula in the concluding sequent is denoted as the *principal* formula with the exception of cut which has no principal formula. In the case of the logical rules, this is the distinguished active formula(s), namely, $\neg A$ in $\neg\mathsf{L}$ and $\neg\mathsf{R}$, and both active formulas in id; the principal formula of the non-cut structural rules is the distinguished formula A. The formula in the premise corresponding to the principal formula (i.e., the distinguished active formula in logical rules and distinguished occurrences of A in structural rules) is called the *minor* formula.

Definition 1. A *proof* is a finite tree π of sequents and inferences locally correct with respect to the inferences in Figure 1 such that leaves are labelled by the zero-premise inference id. The sequent at the root of π is called the *endsequent*. Notation $\pi \vdash \Gamma \Rightarrow \Delta$ expresses that π is a proof with endsequent $\Gamma \Rightarrow \Delta$. The *size* of a proof π, in symbols $|\pi|$, is the number of vertices in the underlying tree. The *subproofs* of π are the proofs that occur as subtrees of π. A subproof of π is *immediate* if its root is a successor of the root of π.

A proof in which all cut formulas are quantifier-free is called *essentially cut-free*. If there are no applications of cut whatsoever, the proof is *cut-free*.

Definition 2. A proof π is *regular* if in every instance of a rule in π except $\exists\mathsf{L}$, all free variables of any premise sequent(s) are free variables of the conclusion.

Every proof can be transformed into a regular proof by simply substitution closed terms for the offending free variables.

Skolemisation can be lifted to proofs in a straightforward way. For each sequent $S = \Gamma \Rightarrow \Delta$, we fix a Σ_1 sequent $S_{\text{sk}} = \Gamma' \Rightarrow \Delta'$ such that $\Gamma'^\neg, \Delta'$ is a sequence

Structural rules

$$\frac{A, A, \Gamma \Rightarrow \Delta}{A, \Gamma \Rightarrow \Delta}\,\mathsf{cL} \qquad \frac{\Gamma, A, \Pi \Rightarrow \Delta}{A, \Gamma, \Pi \Rightarrow \Delta}\,\mathsf{eL} \qquad \frac{\Gamma \Rightarrow \Delta}{A, \Gamma \Rightarrow \Delta}\,\mathsf{wL}$$

$$\frac{\Gamma \Rightarrow \Delta, A, A}{\Gamma \Rightarrow \Delta, A}\,\mathsf{cR} \qquad \frac{\Gamma \Rightarrow \Delta, A, \Pi}{\Gamma \Rightarrow \Delta, \Pi, A}\,\mathsf{eR} \qquad \frac{\Gamma \Rightarrow \Delta}{\Gamma \Rightarrow \Delta, A}\,\mathsf{wR}$$

$$\frac{\Gamma \Rightarrow \Delta, A \qquad A, \Sigma \Rightarrow \Pi}{\Gamma, \Sigma \Rightarrow \Delta, \Pi}\,\mathsf{cut}$$

Logical rules

$$\frac{}{P\vec{s} \Rightarrow P\vec{s}}\,\mathsf{id}$$

$$\frac{\Gamma \Rightarrow \Delta, A}{\neg A, \Gamma \Rightarrow \Delta}\,\neg\mathsf{L} \qquad \frac{\Gamma \Rightarrow \Delta, A_i}{\Gamma \Rightarrow \Delta, A_0 \lor A_1}\,\lor\mathsf{R} \qquad \frac{A[\alpha/x], \Gamma \Rightarrow \Delta}{\exists x A, \Gamma \Rightarrow \Delta}\,\exists\mathsf{L}$$

$$\frac{A, \Gamma \Rightarrow \Delta}{\Gamma \Rightarrow \Delta, \neg A}\,\neg\mathsf{R} \qquad \frac{A_0, \Gamma \Rightarrow \Delta \qquad A_1, \Gamma \Rightarrow \Delta}{A_0 \lor A_1, \Gamma \Rightarrow \Delta}\,\lor\mathsf{L} \qquad \frac{\Gamma \Rightarrow \Delta, A[t/x]}{\Gamma \Rightarrow \Delta, \exists x A}\,\exists\mathsf{R}$$

Fig. 1 Inference rules for $\mathsf{G1c} + \mathsf{cut}$.

of skolemisations of the formulas in $\Gamma^{\neg}, \Delta$. From a proof $\pi \vdash S$ a proof $\pi_{\mathrm{sk}} \vdash S_{\mathrm{sk}}$ can be given by simply eliminating strong quantifier rule instances ($\exists\mathsf{L}$) contributing toward the endsequent and replacing all occurrences of these eigenvariables by terms determined by the skolemisation. In particular, $|\pi_{\mathrm{sk}}| \leq |\pi|$ and every cut formula of π_{sk} is a substitution instance of a cut formula of π.

2.1 A remark on prenex normal form

One restriction on the framework introduced in [1] was that all formulas in a proof were assumed to be in prenex normal form. Here we shall lift this restriction. It is worth commenting on why this matters.

From the standpoint of *provability*, the assumption of prenex normal form is without loss of generality. But from a computational point of view the assumption is not innocent: given a formula A, the proofs of a prenex normal form of A may be considerably different from, and be significantly larger than, proofs of A. The following example illustrates this point.

Let n be a positive integer and associate binary predicates R_i and constants a_i, b_i with each $i < n$. As an abbreviation we set $C_i x \coloneqq (R_i x b_i \lor \neg R_i a_i x)$. Consider the formula:

$$F_n \coloneqq \exists x\, C_0 x \land \cdots \land \exists x\, C_{n-1} x$$

Taking a prenex normal form of this formula according to a standard procedure yields:

$$G_n \coloneqq \exists x_0 \ldots \exists x_{n-1} (C_0 x_0 \land \ldots \land C_{n-1} x_{n-1})$$

Consider the following cut-free proof of F_2:

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{R_0a_0b_0 \Rightarrow R_0a_0b_0}{\Rightarrow R_0a_0b_0, \neg R_0a_0b_0}}{\Rightarrow C_0a_0, C_0b_0}}{\Rightarrow \exists x C_0x, \exists x C_0x}}{\Rightarrow \exists x C_0x} \qquad \dfrac{\dfrac{\dfrac{\dfrac{Ra_1b_1 \Rightarrow Ra_1b_1}{\Rightarrow R_1a_1b_1, \neg Ra_1b_1}}{\Rightarrow C_1a_1, C_1b_1}}{\Rightarrow \exists x C_1x, \exists x C_1x}}{\Rightarrow \exists x C_1x}}{\Rightarrow \exists x C_0x \wedge \exists x C_1x}$$

Here we used a derived right rule for $\wedge$ according to its definition in terms of $\neg, \vee$. Now compare the previous proof with a cut-free proof of G_2, where Dxy abbreviates $C_0x \wedge C_1y$ and the sequent arrow has been omitted for brevity:

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\vdots}{R_0a_0b_0, Da_0b_1, Db_0a_1, Db_0b_1}}{C_0a_0, Da_0b_1, Db_0a_1, Db_0b_1} \qquad \dfrac{\dfrac{\vdots}{R_1a_1b_1, Da_0b_1, Db_0a_1, Db_0b_1}}{C_1a_1, Da_0b_1, Db_0a_1, Db_0b_1}}{Da_0a_1, Da_0b_1, Db_0a_1, Db_0b_1}}{\exists y(C_0a_0 \wedge C_1y), \exists y(C_0a_0 \wedge C_1y), \exists y(C_0b_0 \wedge C_1y), \exists y(C_0b_0 \wedge C_1y)}}{\exists y(C_0a_0 \wedge C_1y), \exists y(C_0b_0 \wedge C_1y)}}{\exists x\exists y(C_0x \wedge C_1y), \exists x\exists y(C_0x \wedge C_1y)}}{\exists x\exists y(C_0x \wedge C_1y)}$$

The increase in size of the proof is not accidental. The number of contractions needed on the root formula G_n will grow exponentially with n, and therefore also the size of its cut-free proofs. By contrast, the size of the cut-free proofs of F_n will be of linear size in n (n subproofs of constant size containing a single contraction, followed by n applications of the right $\wedge$-rule). If we wish to view proofs as programs and cut-elimination as computation, this suggests that restricting to prenex normal form can lead to significantly less efficient programs.

Furthermore, since both formulas F_n and G_n are Σ_1, we obtain Herbrand expansions from their cut-free proofs. From our proof of F_2 we can read off the Herbrand expansion:

$$(R_0a_0b_0 \vee \neg R_0a_0b_0) \wedge (R_1a_1b_1 \vee \neg R_1a_1b_1)$$

From the corresponding proof of G_2 we get:

$$(R_0a_0b_0 \wedge R_1a_1b_1) \vee (R_0a_0b_0 \wedge \neg R_1a_1b_1) \vee (\neg R_0a_0b_0 \wedge R_1a_1b_1) \vee (\neg R_0a_0b_0 \wedge \neg R_1a_1b_1)$$

The latter is the disjunction normal form of the former Herbrand expansion, and it is not hard to see that the correspondence generalises to F_n and G_n. Therefore, prenex normal form can also lead to exponential blow-up in the size of Herbrand expansions.

$$\frac{}{\bot_\Sigma : \Sigma} \qquad \frac{\sigma : \Sigma \quad s : \iota \quad \alpha \in \mathsf{FV}}{[s/\alpha]\sigma : \Sigma} \qquad \frac{t \text{ an } \mathscr{L}\text{-term} \quad \sigma : \Sigma}{t \cdot \sigma : \iota}$$

Fig. 2 Typing rules for substitution stacks and individuals.

3 Types and terms

In the next section we will define the recursion schemes associated with sequent calculus proofs together with their rewrite rules. First, we set up the basic type theory as a derivation system of typing judgements, namely expressions of the form $t : U$. Derivability of the judgement $t : U$ is expressed by the phrase 't has type U'. When context is clear this assertion is expressed simply as $t : U$. Throughout the section we operate over the fixed first-order vocabulary $\mathscr{L}$.

The type system comprises three ground types and several type constructors split into three categories: *extended terms*, *formula types* and *general types*. The first group provides a representation of $\mathscr{L}$-terms and substitutions thereof; formula types denote the types of evidence of formulas; the general types close the above types under functions spaces.

3.1 Extended terms and substitutions

The first category of types comprises just two *ground* types:

- the type of *individuals*, ι.
- the type of *substitution stacks*, Σ.

Formation rules for terms of the above ground types are displayed in Figure 2. We will shortly introduce further term constructors amongst which one, *application*, provides an additional way to construct terms of either type above. However, the specific terms derivable from the three rules of Figure 2 represent canonical terms of type ι and Σ.

Definition 3. A *substitution stack* is any term σ for which the judgement $\sigma : \Sigma$ is derivable using *only* the typing judgements in Figure 2. Likewise, an *extended $\mathscr{L}$-term* is any term t for which the judgement $t : \iota$ is derivable using *only* the typing judgements in Figure 2.

Thus, substitution stacks are finite lists of the form $[s_n/\alpha_n] \cdots [s_1/\alpha_1]\bot_\Sigma$ where α_i is an eigenvariable and s_i is an extended $\mathscr{L}$-term for each $1 \le i \le n$, and extended $\mathscr{L}$-terms are $\mathscr{L}$-terms composed with a substitution stack. The $\mathscr{L}$-terms can thus be identified with extended $\mathscr{L}$-terms $t \cdot \bot_\Sigma$. Given substitution stacks σ and τ, we write $\sigma\tau$ as shorthand for the substitution stack

$$\sigma\tau \coloneqq [s_n/\alpha_n] \cdots [s_1/\alpha_1]\tau \qquad \text{where } \sigma = [s_n/\alpha_n] \cdots [s_1/\alpha_1]\bot_\Sigma.$$

Substitution stacks and extended $\mathscr{L}$-terms denote, respectively, substitutions and $\mathscr{L}$-terms. We now define this interpretation, starting with generalising the notion of substitution.

Definition 4. A *substitution* is a partial function $\theta \colon \mathsf{FV} \to \mathsf{Terms}(\mathscr{L})$ from free variables to $\mathscr{L}$-terms. The unique substitution with empty domain is denoted $\emptyset$. Application of a substitution θ to an $\mathscr{L}$-term t is denoted $t[\theta]$ and defined in the expected way:

$$\alpha[\theta] = \begin{cases} \theta(\alpha), & \text{if } \alpha \in \mathsf{dom}\,\theta, \\ \alpha, & \text{otherwise,} \end{cases} \qquad \begin{aligned} x[\theta] &= x \qquad x \in \mathsf{BV} \\ f(t_1,\dots,t_n)[\theta] &= f(t_1[\theta],\dots,t_n[\theta]). \end{aligned}$$

Two substitutions θ and θ' can be combined to form a substitution $\theta+\theta'$ with domain $\mathsf{dom}\,\theta \cup \mathsf{dom}\,\theta'$ which applies the substitution θ if on variables in its domain and θ' otherwise:

$$(\theta+\theta')(\alpha) = \begin{cases} \theta(\alpha), & \text{if } \alpha \in \mathsf{dom}\,\theta, \\ \theta'(\alpha), & \text{otherwise.} \end{cases}$$

The operation $\theta, \theta' \mapsto \theta + \theta'$ need not be commutative if $\mathsf{dom}\,\theta \cap \mathsf{dom}\,\theta' \neq \emptyset$. However, if θ and θ' have disjoint domains, then $\theta + \theta' = \theta' + \theta$. When there is no cause for confusion, we write $[t/\alpha]$ for the unique substitution with singleton domain $\{\alpha\}$ and codomain $\{t\}$. Thus, $[t/\alpha] + \theta$ denotes the substitution mapping α to t and $\beta(\neq \alpha)$ to $\theta(\beta)$.

We can now define the *value* of an extended individual term and of a substitution stack, which is to be an $\mathscr{L}$-term and a substitution respectively.

Definition 5. The *value* of an extended individual term v and stack σ is the $\mathscr{L}$-term $\mathsf{Val}(v)$ and substitution $\mathsf{Val}(\sigma)$ respectively defined by mutual recursion:

$$\begin{aligned} \mathsf{Val}(\bot_\Sigma) &= \emptyset \\ \mathsf{Val}([u/\alpha]\sigma) &= [\mathsf{Val}(u)/\alpha] + \mathsf{Val}(\sigma) \\ \mathsf{Val}(t \cdot \sigma) &= t[\mathsf{Val}(\sigma)] \end{aligned}$$

Proposition 2. *For all $\mathscr{L}$-terms t and substitution stacks σ, τ we have* $\mathsf{Val}(\sigma\tau) = \mathsf{Val}(\sigma) + \mathsf{Val}(\tau)$.

Proof. By induction on the length of the stack σ. □

3.2 Formula types and general types

From a third ground type □/, the *null* type, the *formula* types are generated via sum, product and (non-dependent) quantifiers. Closure of the basic and formula types under arrow forms the *general* types:

$$\begin{aligned} \mathrm{FmlTyp} \ni F, G &::= \Box \mid F + G \mid F \times G \mid \exists F \mid \forall F \\ \mathrm{Typ} \ni U, V &::= \iota \mid \Sigma \mid F \mid U \to V \end{aligned}$$

These types induce further term constructors listed in Figure 3. Every formula type F has an inhabitant $\bot_F$ and we omit the subscript when this can be determined from context. The operation $F \mapsto F^\perp$ maps each formula type to its type-theoretic *dual*: $\Box^\perp = \Box$, $(F \times G)^\perp = F^\perp + G^\perp$, $(F + G)^\perp = F^\perp \times G^\perp$, $(\forall F)^\perp = \exists F^\perp$ and $(\exists F)^\perp = \forall F^\perp$. A type $U \to V$ is called a *function* type with *domain* U and *codomain* V. Terms of function type are referred to as *functions*. As the only formation rule for terms of function type is application, there are no functions in the term calculus induced by the rules in Fig. 2 and 3 alone. Instead, functions will be instantiated by specific additional constants, called *nonterminals*, discussed in the next section.

The above constructors are not associated reduction rules and should be viewed as constructors of a programming language that are 'interpreted' by the Herbrand scheme. One of the constructors in Figure 3 deserve special mention, namely the *Peirce constant* $\mathtt{p}$. This is used in the interpretation of contraction, and supplies evidence for a formula A from a function mapping counter-evidence for A to evidence for A.

The Peirce constructor is essentially a witness of Peirce's axiom:

$$((A \to B) \to A) \to A,$$

or rather its particular instance $(\neg A \to A) \to A$, where F and $F^\perp$ represent A and $\neg A$ respectively. In this system the behaviour of the Peirce constant will be similar to the 'call-with-current-continuation' operator which has featured in various computational interpretations of classical theories, for example Krivine's classical realizability [9]. Its specific purpose in the present work is to provide computational reading of the contraction rule. The correlation to Peirce's axiom is made apparent by the 'formulas as types' encoding below.

Definition 6. Each $\mathscr{L}$-formula A is associated two formula types, denoted $[A]$ and $\langle A\rangle$, called respectively the type of *evidence* for A and *counter-evidence* for A:

$$\begin{array}{llll} [P\vec{s}] = \Box & [A \vee B] = [A] + [B] & [\exists x A] = \exists [A] & [\neg A] = \langle A \rangle \\ \langle P\vec{s}\rangle = \Box & \langle A \vee B\rangle = \langle A\rangle \times \langle B\rangle & \langle \exists x A\rangle = \forall \langle A\rangle & \langle \neg A\rangle = [A] \end{array}$$

A term s is said to have Σ_1-type if $s : [A]$ for some Σ_1 formula A.

Of note is that atomic formulas are uniformly mapped to the null type independently of the predicate and term structure. As such there is no type-theoretic distinction between evidence and counter-evidence of such formulas. More generally, evidence and counter-evidence are type-theoretic duals:

Proposition 3. *For all formulas A, $[A] = \langle A\rangle^\perp$.*

Proposition 4. *The evidence type of every Σ_1 formula is inhabited by a term not involving $\bot_F$ for F a formula type.*

Proof. By induction, noting that given an inhabitant $i : [A]$ we have $\mathtt{e}(\mathtt{c} \cdot \bot_\Sigma)i : [\exists x A]$. $\Box$

The 'formulas as types' encoding described above is extended to sequents as functions between evidence types.

$$\frac{}{\varepsilon : \Box} \qquad \frac{f_0 : F_0 \quad f_1 : F_1}{\mathtt{j} f_0 f_1 : F_0 \times F_1} \qquad \frac{f : F_i \quad i \in \{0,1\}}{\mathtt{i}_i f : F_0 + F_1} \qquad \frac{u : \iota \quad f : F}{\mathtt{e} u f : \exists F} \qquad \frac{u : \iota \to F}{\mathtt{a} u : \forall F}$$

$$\frac{}{\bot_F : F} \qquad \frac{u : F^\perp \to F}{\mathtt{p} u : F} \qquad \frac{u : U \quad v : U \to V}{v u : V}$$

Fig. 3 Typing rules for formula and general types. F, G range over formula types and U, V over general types.

Definition 7. The *type interpretation* of a sequent $\Gamma \Rightarrow \Delta$, in symbols $[\Gamma \Rightarrow \Delta]$, is defined as follows where $\Gamma = A_1, \ldots, A_m$ and $\Delta = B_1, \ldots, B_n$:

$$[\Gamma \Rightarrow \Delta] \coloneqq [A_1] \to \cdots \to [A_m] \to \langle B_1 \rangle \to \cdots \to \langle B_{n-1} \rangle \to [B_n].$$

Thus, an inhabitant of $[\Gamma \Rightarrow \Delta, B]$ is a function mapping evidence for each formula in Γ and count-evidence for each formula in Δ to evidence for B. This interpretation is a natural extension of the usual type-theoretic interpretation of an intuitionistic sequent $\Gamma \Rightarrow B$ as a function mapping justification for formulas in Γ to justification for B. The distinctive feature of classical sequent calculus is the permission of contraction to the right side of the sequent arrow. The type-theoretic interpretation presents right contraction as a transformation from $[\Gamma \Rightarrow \Delta, B, B] = \cdots \to \langle B \rangle \to [B]$ to $[\Gamma \Rightarrow \Delta, B] = \cdots \to [B]$, made possible by the Peirce constant.

Note, $[\Gamma \Rightarrow \Delta]$ yields the same type as $[\emptyset \Rightarrow \Gamma^{\neg}, \Delta]$. Applying the interpretation to the abbreviation $\Gamma \Rightarrow_i \Delta$, with Γ and Δ as in the definition above, we obtain for $i < m$ and $j < |\Delta|$:

$$[\Gamma \Rightarrow_i \Delta] = \begin{cases} \cdots \to [A_{i-1}] \to [A_{i+1}] \to \cdots \to \langle B_n \rangle \to \langle A_i \rangle, & i < m, \\ \cdots \to \langle B_{j-1} \rangle \to \langle B_{j+1} \rangle \to \cdots \to \langle B_n \rangle \to [B_j], & j = i - m. \end{cases}$$

4 Herbrand schemes

In this section we describe the recursion schemes associated with sequent calculus proofs. We assume some familiarity with higher-order recursion schemes. For background on recursion schemes see [14]. We will use a variant of recursion schemes involving *pattern matching* as in [1].

In the present setting, we define a higher-order recursion scheme generally to be a structure $\mathscr{S} = (\mathscr{N}, \mathscr{T}, \mathscr{V}, \mathscr{R}, \mathsf{S})$ consisting of the following data:

- A set $\mathscr{N}$ of typed *non-terminal symbols*,
- A set $\mathscr{T}$ of typed *terminal symbols*,
- A set $\mathscr{V}$ of typed *variables*,
- A set $\mathscr{R}$ of *rewrite*, or *production*, *rules*,
- A distinguished *start symbol* $\mathsf{S} \in \mathscr{N}$.

Let T be the set of well-typed terms in $\mathscr{N} \cup \mathscr{T} \cup \mathscr{V}$. Formally, a rewrite rule in $\mathscr{R}$ is a pair $(t, s) \in T \times T$ were s and t are of the same type. An *instance* of the rule (t_0, t_1) is a pair $(t_0[\sigma], t_1[\sigma])$ where σ is a type-preserving substitution from variables to T. We say that t_0 *one-step rewrites* to t_1, denoted $t_0 \longrightarrow^1 t_1$, if (t_0, t_1) is an instance of a rewrite rule in $\mathscr{R}$. We write $t_0 \longrightarrow t_1$ and say that t_0 rewrites to t_1 if the pair (t_0, t_1) is in the reflexive, transitive closure of the one-step rewrite relation. The *language* $\mathscr{L}(\mathscr{S})$ of a recursion scheme $\mathscr{S}$ with start symbol S is the set of terms t containing no variables or non-terminals such that $\mathsf{S} \longrightarrow t$. We will often write $t \longrightarrow u \parallel v$ ($t \longrightarrow^1 u \parallel v$) to say that $t \longrightarrow u$ ($t \longrightarrow^1 u$) and $t \longrightarrow v$ ($t \longrightarrow^1 v$).

Of course, not every recursion scheme in the sense above corresponds to a reasonable model of computation. At the very least one should require the set of rewrite rules to be recursive as well as certain constraints on the form of the first term in the pair $(t, s) \in \mathscr{R}$. Indeed, frequently the rewrites of a higher-order recursion scheme are required to be of the shape:

$$\mathsf{F} x_0 \cdots x_{n-1} \longrightarrow t$$

where F is a non-terminal, $x_0, \ldots, x_{n-1}$ are pairwise distinct variables of appropriate type and t is a term containing no variables other than the x_i.

The more general definition allows for context sensitive rewrites. While the recursion schemes that we associate with proofs require more general rewrites than immediately above, these will be constrained in comparison to the definition above. In particular, the only context sensitivity utilised in Herbrand schemes rewrites is *pattern matching*. Rewrites in the recursion schemes that follow will have the general form

$$\mathsf{F} t_0 \cdots t_{n-1} \longrightarrow t$$

where $t_i = \mathsf{f}_i x_{i,0} \cdots x_{i,k_i}$ is a term constructor with k_i the associated arity, and all the $x_{i,j}$ are pairwise distinct. Thus, some rewrite rules will depend not only on the outermost non-terminal but also the shape of its arguments.

Note that we also permit the sets of non-terminals and terminals to be infinite. This is just a technical convenience; rather than assigning separate sets of non-terminals and terminals to each proof, it will be simpler to have fixed sets of non-terminals and terminals with fixed rewrite rules. Essentially, there is a single infinite 'universal' recursion scheme $\mathscr{H}$ with no start symbol, and each individual scheme $\mathscr{H}(\pi)$ is specified by a start symbol and its rewrite rule. In practice, the recursion scheme associated with a proof will always be equivalent to one using only finitely many non-terminals, terminals and variables.

The remainder of this section defines the Herbrand scheme $\mathscr{H}(\hat{\pi})$ associated to a proof $\hat{\pi} \vdash \Gamma \Rightarrow A$ where A is Σ_1. The terminals of $\mathscr{H}(\hat{\pi})$ comprise all symbols in the type system introduced in Section 3.

The non-terminals of $\mathscr{H}(\hat{\pi})$ comprise the following symbols:

- A *start symbol* $\mathsf{S}_{\hat{\pi}} : [A]$.
- A *proof non-terminal* $\mathsf{F}^{\pi}_{i} : \Sigma \to [\Lambda \Rightarrow_i \Pi]$ for each subproof $\pi \vdash \Lambda \Rightarrow \Pi$ of $\hat{\pi}$ and $i < |\Lambda \Rightarrow \Pi|$. These non-terminals compute evidence for the i-th formula occurrence in $\Lambda \Rightarrow \Pi$ from counter-evidence for the remaining formula occurrences in the sequent.
- *Extractor* non-terminals $\mathsf{E}_B : [B] \to [B]$ for each B that is a subformula of A or the negation of a subformula of A. These non-terminals extract witnesses for the weak quantifiers in B from arbitrary evidence by recursively eliminating Peirce constants.
- *Helper* non-terminals used to express function composition and combinators and to simulate specific cases of λ-abstraction.

The remainder of this section presents the rewrite rules for the above non-terminals, starting with the helper functions.

4.1 Helper functions

Each of these non-terminals is assigned a single (deterministic) production rule that simulates a particular aspect of λ-abstraction over the underlying type system: composition, exchange, redundancy and substitution formation. By omitting abstraction as a formal constructor, we avoid the need to accommodate β-reduction alongside production rules and to reason about arbitrary λ-abstractions that cannot be simulated by the recursion scheme. Indeed, the particular λ-abstractions expressed by these non-terminals are all *sub-linear* in the sense that they express abstractions $\lambda x\, t$ with at most one occurrence of x in t.

For all types U, V, W and $\vec{V} = V_1, \ldots, V_n$, the following non-terminals are included in $\mathscr{H}(\pi)$ with associated production rule:

$$\circ : (V \to W) \to (U \to V) \to U \to W \qquad \mathsf{A}^n : (U \to \vec{V} \to W) \to \vec{V} \to U \to W$$
$$(x \circ y)z \coloneqq \circ xyz \longrightarrow x(yz) \qquad \mathsf{A}^n\, w\vec{x}z \longrightarrow wz\vec{x} \quad \text{where } |\vec{x}| = n$$

$$\mathsf{K} : U \to V \to U \qquad \mathsf{Sbs}_\alpha : \Sigma \to \iota \to \Sigma \quad \text{for each } \alpha \in \mathsf{FV}$$
$$\mathsf{K}\, xy \longrightarrow x \qquad \mathsf{Sbs}_\alpha\, xy \longrightarrow [y/\alpha]x$$

Despite foregoing λ-abstraction at the formal level, it nonetheless provides a convenient notation for expressing terms constructed from the helper non-terminals. Thus, in the sequel we will more often use the language of λ-calculus than the above non-terminals. Except where stated otherwise, such notation will be strictly confined to constructions that are expressible via the above symbols and other terms/non-terminals.

The three main examples of this notation are presented in Figures 2, 1 and 3 (cL, $\exists\mathsf{L}$ and Peirce reductions respectively), the expanded form of which are:

$$\begin{aligned}
\lambda v.\, \mathsf{F}\sigma v\vec{x} &\quad \text{expands to} \quad \mathsf{A}^{|\vec{x}|}(\mathsf{F}\sigma)\vec{x} \\
\lambda v.\, \mathsf{F}([v/\alpha]\sigma)\vec{x} &\quad \text{expands to} \quad (\mathsf{A}^{|\vec{x}|}\mathsf{F}\vec{x}) \circ (\mathsf{Sbs}_\alpha\sigma) \\
\lambda v.\, \mathsf{F}\sigma\vec{w}\vec{x}(z(\mathsf{G}\sigma\vec{w}v\vec{x}\vec{y}))\vec{y} &\quad \text{expands to} \quad (\mathsf{A}^{|\vec{y}|}(\mathsf{F}\sigma\vec{w}\vec{x})\vec{y}) \circ (z \circ (\mathsf{A}^{|\vec{x}\vec{y}|}(\mathsf{G}\sigma\vec{w})\vec{x}\vec{y}))
\end{aligned}$$

The above abbreviations do not utilise the non-terminal K. This non-terminal is useful in defining generic inhabitants of each formula type.

For every formula B we introduce a non-terminal $\mathsf{C}_B : [B]$, called the *generic evidence* for B with rewrite rules given as follows:

$$\begin{aligned}
&\mathsf{C}_B \longrightarrow \varepsilon,\ B \text{ atomic} \qquad \mathsf{C}_{B\vee C} \longrightarrow \mathtt{i}_0\mathsf{C}_B \mathbin{\|} \mathtt{i}_1\mathsf{C}_C \qquad \mathsf{C}_{\exists y B} \longrightarrow \mathtt{e}(\mathtt{c}\cdot\bot)\mathsf{C}_B \\
&\mathsf{C}_{\neg\neg B} \longrightarrow \mathsf{C}_B \qquad \mathsf{C}_{\neg B\vee C} \longrightarrow \mathtt{j}\mathsf{C}_{\neg B}\mathsf{C}_{\neg C} \qquad \mathsf{C}_{\neg\exists y B} \longrightarrow \mathtt{a}\big(\mathsf{K}\mathsf{C}_{\neg A[\mathtt{c}/y]}\big)
\end{aligned}$$

Intuitively, the term C_B describes the simplest strategy to provide evidence for the formula B.

4.2 Extractors

Each Σ_1-formula B has associated an *extractor* non-terminal $\mathsf{E}_B : [B] \to [B]$. These non-terminals follow a simple behaviour, recursively eliminating instances of the Peirce operator from their argument and commuting with all other term builders. When encountering a Peirce term the extractor simply evaluates the guarded function on generic evidence for the appropriate type. Their only role is with the start symbol of the Herbrand scheme for which they extract terms suitable for a Herbrand disjunction from arbitrary evidence (or counter-evidence) for a formula. These extractors have the following rewrite rules.

$$\begin{array}{rlrl} \mathsf{E}_B\varepsilon & \longrightarrow \varepsilon & \mathsf{E}_B\bot_{[B]} & \longrightarrow \bot_{[B]} \\ \mathsf{E}_{B_0\vee B_1}(\mathtt{i}_i x) & \longrightarrow \mathtt{i}_i(\mathsf{E}_{B_i}x) & \mathsf{E}_{\neg\neg B}x & \longrightarrow \mathsf{E}_B x \\ \mathsf{E}_{\neg B_0\vee B_1}(\mathtt{j}xy) & \longrightarrow \mathtt{j}(\mathsf{E}_{\neg B_0}x)(\mathsf{E}_{\neg B_1}y) & \mathsf{E}_B(\mathtt{p}z) & \longrightarrow \mathsf{E}_B(z\mathsf{C}_{\neg B}) \\ \mathsf{E}_{\exists xB}(\mathtt{e}xy) & \longrightarrow \mathtt{e}x(\mathsf{E}_B y) & & \end{array}$$

The above rewrites rely on pattern-matching to determine which rewrite rule is applicable.

Proposition 5. *For every Σ_1 formula B and every final term t, $\mathsf{E}_B\mathsf{C}_B \longrightarrow t$ iff $\mathsf{C}_B \longrightarrow t$.*

4.3 Start symbol

The start symbol $\mathsf{S}_{\hat\pi}$ of $\mathscr{H}(\hat\pi)$ is associated a single rewrite rule:

$$\mathsf{S}_{\hat\pi} \longrightarrow \mathsf{E}_A\big(\mathsf{F}^{\hat\pi}_n\bot\mathsf{C}_{A_1}\cdots\mathsf{C}_{A_n}\big)$$

where, recall, the endsequent of $\hat\pi$ is $\Gamma \Rightarrow A$ (with $A \in \Sigma_1$) and $\Gamma = A_1, \ldots, A_n$.

4.4 Proof non-terminals

Finally, and most importantly, for each proof π with end sequent $\Lambda \Rightarrow \Pi$ and each index $i < |\Lambda \Rightarrow \Pi|$ there is a non-terminal $\mathsf{F}^\pi_i : \Sigma \to [\Lambda \Rightarrow_i \Pi]$. These non-terminals are each associated re-write rules according to the final inference of π and particular structure of certain arguments. These are listed in Figures 1, 2 and 3 and have two essential forms:

Inference productions. A single production rule for F^π_i determined completely by the final inference in π and the immediate subproofs. These production rules are listed in Figures 1 and 2. Only one such inference production is associated to each proof non-terminal whose general shape is determined by whether the i-th formula of $\Lambda \Rightarrow \Pi$ is principal. The inference production for F^π_i maps this non-terminal to a term built from terminals and non-terminals F^π_j where π_0 is an immediate subproof of π. In some cases, this term will include constant and helper non-terminals. The production rule for F^π_i depends on pattern-matching only if the final inference is a logical inference and the i is not the index of the principal formula.

Internal productions. These are production rules associated every proof non-terminal F^π_i and applicable whenever at least one argument is $\bot$ or guarded by the Peirce or choice constant $\mathtt{p}x$. These productions are listed in Figure 3. Each rewrite 'consumes' the matched term constructor and 'reduces' the term F^π_i to one involving the same non-terminal F^π_i and, in the case of the Peirce reduction, a non-terminal F^π_j for $j \neq i$.

In the remainder of the section we discuss the production rules in more detail with some examples. Further examples are given in Section 6 where two Herbrand schemes associated to proofs are examined in detail.

4.4.1 Undefined input

Each such non-terminal proof has the following distinguished rewrite rules, in order to handle 'undefined' inputs.

$$\mathsf{F}^\pi_i \vec{x} \bot_C \vec{y} \longrightarrow \bot_{\langle A \rangle} \qquad \mathsf{F}^\pi_j \vec{x} \bot_C \vec{y} \longrightarrow \bot_{[B]}$$

In the above rewrites i is an index corresponding to a left formula occurrence A, j is an index corresponding to a right formula occurrence B, and the formula corresponding to the position of the argument $\bot_C$ is either a left or right occurrence of C. The notation $\bot_C$ is to be understood as $\bot_{[C]}$ or $\bot_{\langle C \rangle}$ as appropriate. Since the distinction between $\bot_{[A]}$ and $\bot_{\langle A \rangle}$ will not affect the evaluation of terms, we may abuse notation

Inference	Principal rewrite	Non-principal rewrite
$\dfrac{}{P\vec{s} \Rightarrow P\vec{s}}\,\mathsf{id}$	$\mathsf{F}^\pi_* \sigma x \longrightarrow \varepsilon$	
$\dfrac{\overset{\pi_0}{A, \Gamma \Rightarrow \Delta}}{\Gamma \Rightarrow \Delta, \neg A}\,\neg\mathsf{R}$	$\mathsf{F}^\pi_* \sigma \vec{x} \longrightarrow \mathsf{F}^{\pi_0}_* \sigma \vec{x}$	$\mathsf{F}^\pi_i \sigma \vec{x} z \longrightarrow \mathsf{F}^{\pi_0}_{i+1} \sigma z \vec{x}$
$\dfrac{\overset{\pi_0}{\Gamma \Rightarrow \Delta, A}}{\neg A, \Gamma \Rightarrow \Delta}\,\neg\mathsf{L}$	$\mathsf{F}^\pi_* \sigma \vec{x} \longrightarrow \mathsf{F}^{\pi_0}_* \sigma \vec{x}$	$\mathsf{F}^\pi_i \sigma z \vec{x} \longrightarrow \mathsf{F}^{\pi_0}_{i-1} \sigma \vec{x} z$
$\dfrac{\overset{\pi_0}{\Gamma \Rightarrow \Delta, A_j}}{\Gamma \Rightarrow \Delta, A_0 \vee A_1}\,\vee\mathsf{R}$	$\mathsf{F}^\pi_* \sigma \vec{x} \longrightarrow \mathtt{i}_j \big(\mathsf{F}^{\pi_0}_* \sigma \vec{x}\big)$	$\mathsf{F}^\pi_i \sigma \vec{x} \big(\mathtt{j} z_0 z_1\big) \longrightarrow \mathsf{F}^{\pi_0}_i \sigma \vec{x} z_j$
$\dfrac{\overset{\pi_0}{A_0, \Gamma \Rightarrow \Delta} \quad \overset{\pi_1}{A_1, \Gamma \Rightarrow \Delta}}{A_0 \vee A_1, \Gamma \Rightarrow \Delta}\,\vee\mathsf{L}$	$\mathsf{F}^\pi_* \sigma \vec{x} \longrightarrow \mathtt{j}(\mathsf{F}^{\pi_0}_* \sigma \vec{x})(\mathsf{F}^{\pi_1}_* \sigma \vec{x})$	$\mathsf{F}^\pi_i \sigma \vec{x} \big(\mathtt{i}_j z\big) \longrightarrow \mathsf{F}^{\pi_j}_i \sigma \vec{x} z$
$\dfrac{\overset{\pi_0}{\Gamma \Rightarrow \Delta, A[t/x]}}{\Gamma \Rightarrow \Delta, \exists x A}\,\exists\mathsf{R}$	$\mathsf{F}^\pi_* \sigma \vec{x} \longrightarrow \mathtt{e}(t \cdot \sigma)(\mathsf{F}^{\pi_0}_* \sigma \vec{x})$	$\mathsf{F}^\pi_i \sigma \vec{x} (\mathtt{a} z) \longrightarrow \mathsf{F}^{\pi_0}_i \sigma \vec{x} (z (t \cdot \sigma))$
$\dfrac{\overset{\pi_0}{A[\alpha/x], \Gamma \Rightarrow \Delta}}{\exists x A, \Gamma \Rightarrow \Delta}\,\exists\mathsf{L}$	$\mathsf{F}^\pi_* \sigma \vec{x} \longrightarrow \mathtt{a}\big(\lambda w. \mathsf{F}^{\pi_0}_* \big([w/\alpha]\sigma\big) \vec{x}\big)$	$\mathsf{F}^\pi_i \sigma (\mathtt{e} z_0 z_1) \vec{x} \longrightarrow \mathsf{F}^{\pi_0}_i ([z_0/\alpha]\sigma) z_1 \vec{x}$

Table 1 Production rules for F^π_i derived from logical inferences. In all cases $\vec{x}$ is a sequence of variables of length $|\Gamma\Delta|$ typed appropriately. Subscript $*$ denotes the index(es) of the principal formula(s) of π (in the context of F^π_*) and the index of the corresponding minor formula (in $\mathsf{F}^{\pi_0}_*$ etc.).

writing the more informal $\bot_A$. This term should always be read as a term of type $[A]$ or $\langle A \rangle$, which should be clear from context.

4.4.2 Axiom

Consider an instance of a weakened axiom:

$$\dfrac{\dfrac{}{P\vec{s} \Rightarrow P\vec{s}}\,\mathsf{id}}{A, P\vec{s} \Rightarrow P\vec{s}, B}\,\mathsf{wL}+\mathsf{wR}$$

Inference	Principal rewrite	Non-principal rewrite
$\dfrac{\overset{\pi_0}{\Gamma \Rightarrow \Delta, A, A}}{\Gamma \Rightarrow \Delta, A}\,\mathsf{cR}$	$\mathsf{F}^{\pi}_{*}\sigma\vec{x} \longrightarrow \mathtt{p}\big(\mathsf{F}^{\pi_0}_{*}\sigma\vec{x}\big) \parallel \mathtt{p}\big(\mathsf{F}^{\pi_0}_{*'}\sigma\vec{x}\big)$	$\mathsf{F}^{\pi}_{i}\sigma\vec{x}z \longrightarrow \mathsf{F}^{\pi_0}_{i}\sigma\vec{x}zz$
$\dfrac{\overset{\pi_0}{A, A, \Gamma \Rightarrow \Delta}}{A, \Gamma \Rightarrow \Delta}\,\mathsf{cL}$	$\mathsf{F}^{\pi}_{*}\sigma\vec{x} \longrightarrow \mathtt{p}\big(\hat{\mathsf{F}}^{\pi_0}_{*}\sigma\vec{x}\big) \parallel \mathtt{p}\big(\hat{\mathsf{F}}^{\pi_0}_{*'}\sigma\vec{x}\big)$	$\mathsf{F}^{\pi}_{i}\sigma z\vec{x} \longrightarrow \mathsf{F}^{\pi_0}_{i+1}\sigma zz\vec{x}$
$\dfrac{\overset{\pi_0}{\Gamma \Rightarrow \Delta, A, \Pi}}{\Gamma \Rightarrow \Delta, \Pi, A}\,\mathsf{eR}$	$\mathsf{F}^{\pi}_{*}\sigma\vec{x}\vec{y} \longrightarrow \mathsf{F}^{\pi_0}_{*}\sigma\vec{x}\vec{y}$	$\mathsf{F}^{\pi}_{i}\sigma\vec{x}z \longrightarrow \begin{cases} \mathsf{F}^{\pi_0}_{i}\sigma\vec{x}\vec{y}, & \text{if } i < \lvert\Gamma\Delta\rvert, \\ \mathsf{F}^{\pi_0}_{i+1}\sigma\vec{x}\vec{y}, & \text{if } i \geq \lvert\Gamma\Delta\rvert. \end{cases}$
$\dfrac{\overset{\pi_0}{\Pi, A, \Gamma \Rightarrow \Delta}}{A, \Pi, \Gamma \Rightarrow \Delta}\,\mathsf{eL}$	$\mathsf{F}^{\pi}_{*}\sigma\vec{y}\vec{x} \longrightarrow \mathsf{F}^{\pi_0}_{*}\sigma\vec{y}\vec{x}$	$\mathsf{F}^{\pi}_{i}\sigma z\vec{y}\vec{x} \longrightarrow \begin{cases} \mathsf{F}^{\pi_0}_{i-1}\sigma\vec{y}z\vec{z}, & \text{if } 0 < i \leq \lvert\Pi\rvert, \\ \mathsf{F}^{\pi_0}_{i}\sigma\vec{y}z\vec{x}, & \text{if } i > \lvert\Pi\rvert. \end{cases}$
$\dfrac{\overset{\pi_0}{\Gamma \Rightarrow \Delta}}{\Gamma \Rightarrow \Delta, A}\,\mathsf{wR}$	$\mathsf{F}^{\pi}_{*}\sigma\vec{x} \longrightarrow \bot_{[A]}$	$\mathsf{F}^{\pi}_{i}\sigma\vec{x}y \longrightarrow \mathsf{F}^{\pi_0}_{i}\sigma\vec{x}$
$\dfrac{\overset{\pi_0}{\Gamma \Rightarrow \Delta}}{A, \Gamma \Rightarrow \Delta}\,\mathsf{wL}$	$\mathsf{F}^{\pi}_{*}\sigma\vec{x} \longrightarrow \bot_{\langle A \rangle}$	$\mathsf{F}^{\pi}_{i}\sigma y\vec{x} \longrightarrow \mathsf{F}^{\pi_0}_{i-1}\sigma\vec{x}$
$\dfrac{\overset{\pi_0}{\Gamma \Rightarrow \Delta, C} \quad \overset{\pi_1}{C, \Lambda \Rightarrow \Pi}}{\Gamma, \Lambda \Rightarrow \Delta, \Pi}\,\mathsf{cut}$		$\mathsf{F}^{\pi}_{i}\sigma\vec{x}_0\vec{y}\vec{x}_1 \longrightarrow \begin{cases} \mathsf{F}^{\pi_0}_{i}\sigma\vec{x}(\mathsf{F}^{\pi_1}_{*}\sigma\vec{y}), & \text{if } i < \lvert\Gamma\rvert, \\ \mathsf{F}^{\pi_0}_{j}\sigma\vec{x}(\mathsf{F}^{\pi_1}_{*}\sigma\vec{y}), & \text{if } i = \lvert\Gamma\Lambda\rvert + j < \lvert\Gamma\Lambda\Delta\rvert, \\ \mathsf{F}^{\pi_1}_{j+1}\sigma(\mathsf{F}^{\pi_0}_{*}\sigma\vec{x})\vec{y}, & \text{if } i = \lvert\Gamma\rvert + j < \lvert\Gamma\Lambda\rvert, \\ \mathsf{F}^{\pi_1}_{j+1}\sigma(\mathsf{F}^{\pi_0}_{*}\sigma\vec{x})\vec{y}, & \text{if } i = \lvert\Gamma\Lambda\Delta\rvert + j. \end{cases}$

Table 2 Production rules for F^{π}_{i} derived from structural inferences. Subscript $*$ denotes the index of the principal formula of π and the corresponding minor formula/cut formula in the premise(s). In the case of the contraction rules, indices of the two minor formulas are denoted $*$ and $*'$ respectively. In the principal rewrite for cL, $\hat{\mathsf{F}}^{\pi_0}_{i}\sigma\vec{x}$ abbreviates $\lambda v.\, \mathsf{F}^{\pi_0}_{i}\sigma v\vec{x} := \mathsf{A}^{\lvert\vec{x}\rvert}(\mathsf{F}^{\pi_0}_{i}\sigma)\vec{x}$.

$\mathsf{F}^{\pi}_{i}\sigma\vec{w}\vec{x}(\mathtt{p}z)\vec{y} \longrightarrow \mathtt{p}\big(\lambda v.\, \mathsf{F}^{\pi}_{i}\sigma\vec{w}\vec{x}(z(\mathsf{F}^{\pi}_{i+k}\sigma\vec{w}v\vec{x}\vec{y}))\vec{y}\big)$	for $i = \lvert\vec{w}\rvert$ and $k = \lvert\vec{x}\rvert$
$\mathsf{F}^{\pi}_{i+k}\sigma\vec{w}(\mathtt{p}z)\vec{x}\vec{y} \longrightarrow \mathtt{p}\big(\lambda v.\, \mathsf{F}^{\pi}_{i+k}\sigma\vec{w}(z(\mathsf{F}^{\pi}_{i}\sigma\vec{w}\vec{x}v\vec{y}))\vec{x}\vec{y}\big)$	for $i = \lvert\vec{w}\rvert$ and $k = \lvert\vec{x}\rvert > 0$
$\mathsf{F}^{\pi}_{i}\sigma\vec{x}\bot\vec{y} \longrightarrow \bot$	

Table 3 Production rules associated to Peirce constructors and ‘undefined’ inputs.

Let π denote the proof above and π_0 the trivial subproof comprising the axiom only. For $i \in \{0, 1\}$ there is a proof non-terminal $\mathsf{F}_i^{\pi_0}$ of type $\Sigma \to \Box \to \Box$ (since $\langle P\vec{s}\rangle = [P\vec{s}] = \Box$). The rewrite in each case is simply $\mathsf{F}_i^{\pi_0}\sigma x \longrightarrow \varepsilon$.

Turning to the weakened axiom π, let l be the length of the root sequent. For each $i < l$ there is a proof non-terminal F_i^{π} which is a function from substitution stacks to the type interpretation of $A, P\vec{s} \Rightarrow_i P\vec{s}, B$. This takes one of three forms depending on the value of i:

$$\mathsf{F}_i^{\pi} : \Sigma \to \begin{cases} \Box \to \Box \to \langle B\rangle \to \langle A\rangle, & \text{if } i = 0, \\ [A] \to \Box \to \langle B\rangle \to \Box, & \text{if } i = 1, 2, \\ [A] \to \Box \to \Box \to [B], & \text{if } i = 3. \end{cases}$$

The rewrite for F_i^{π} is constant in each case:

$$\mathsf{F}_i^{\pi}\sigma\vec{x} \longrightarrow \begin{cases} \bot_{\langle A\rangle}, & \text{if } i = 0, \\ \varepsilon, & \text{if } i = 1, 2, \\ \bot_{[B]}, & \text{if } i = 3. \end{cases}$$

In all the above rewrites σ is a variable of type Σ and $\vec{x} = x_0x_1x_2$ is a sequence of variables of appropriate type dependent on the value of i.

4.4.3 Negation

The first non-trivial logical inference we consider is the right $\neg$-rule. (The rules for the left rule are entirely analogous.) Consider a proof π:

$$\dfrac{\begin{array}{c}\pi_0 \\ A, \Gamma \Rightarrow \Delta\end{array}}{\Gamma \Rightarrow \Delta, \neg A}\,\neg\mathsf{R}$$

Let the length of Γ and Δ be m and n respectively. The principal rewrite for this inference is:

$$\mathsf{F}_{m+n}^{\pi}\sigma\vec{x}\vec{y} \longrightarrow \mathsf{F}_0^{\pi_0}\sigma\vec{x}\vec{y}$$

where the i-th variable in $\vec{x}$ has the type of evidence for the i-th formula in the sequence $\Gamma\Delta^{\neg}$. In Figure 1 the two indices $m + n$ and 0 are abbreviated as $*$ as they refer to the principal/minor formula of the inference.

The non-principal rewrites associated to this inference are the following where $z : \langle\neg A\rangle$ and $\vec{x}$ is typed appropriately for the formulas in $\Gamma\Delta^{\neg}$ minus the i-th formula:

$$\mathsf{F}_i^{\pi_0}\sigma\vec{x}z \longrightarrow \mathsf{F}_{i+1}^{\pi_1}\sigma z\vec{x}.$$

4.4.4 Disjunction

Suppose π is a proof of the form:

$$\dfrac{\begin{array}{c}\pi_0\\ \Gamma \Rightarrow \Delta, A_j\end{array}}{\Gamma \Rightarrow \Delta, A_0 \vee A_1}\ \vee\mathsf{R}$$

where $j \in \{0,1\}$. Henceforth we adopt the convention in Fig. 1 of using $*$ to refer to the index of the principal and minor formula (in this example, $* = |\Gamma\Delta|$ in both cases). As principal rewrite we have:

$$\mathsf{F}^{\pi}_{*}\sigma\vec{x} \longrightarrow \mathtt{i}_j(\mathsf{F}^{\pi_0}_{*}\sigma\vec{x})$$

As above, $\vec{x}$ is a sequence of variables of length $|\Gamma\Delta^{\neg}|$ with corresponding type. The principal rule is required to produce evidence for the formula in focus, the disjunction $A_0 \vee A_1$. From the immediate subproof we extract evidence for the disjunct A_j via the term $\mathsf{F}^{\pi_0}_{*}\sigma\vec{x}$ which is combined with the constructor $\mathtt{i}_j$ to obtain evidence for the disjunction.

As non-principal rewrites we have the following.

$$\mathsf{F}^{\pi}_{i}\sigma\vec{x}(\mathtt{j}z_0z_1) \longrightarrow \mathsf{F}^{\pi_0}_{i}\sigma\vec{x}z_j$$

This rewrite is producing evidence for the i-th formula in the sequent (with $i \neq *$). As input, the non-terminal has counter-evidence for the other formulas, including counter-evidence for the principal disjunction. From the immediate subproof evidence for the required formula can be obtained if we can provide counter-evidence for the specific disjunct A_j. This will be available provided the counter-evidence for $A_0 \vee A_1$ is in 'canonical' form $\mathtt{j}t_0t_1$. Thus, the non-principal rewrite for $\vee\mathsf{R}$ is a example of a *pattern matching* rule as it can only be applied when the final argument of F^{π}_{i} is of the specific form $\mathtt{j}t_0t_1$ (for any terms t_0 and t_1) whence the rule discards the constructor and one of t_0 or t_1. In particular, the rewrite above cannot be applied if the final argument (corresponding to the principal disjunction) is an instance of choice $t_0 \parallel t_1$, the Peirce guard $\mathtt{p}t$ or a non-terminal such as $\mathsf{F}^{\hat{\pi}}_{j}\tau\vec{s}$ for some proof $\hat{\pi}$.

Via disjunction and negation the familiar rule for conjunction can be derived:

$$\dfrac{\begin{array}{c}\pi_0\\ A_j, \Gamma \Rightarrow \Delta\end{array}}{\pi : A_0 \wedge A_1, \Gamma \Rightarrow \Delta}\ \wedge\mathsf{L} \quad \coloneqq \quad \dfrac{\dfrac{\dfrac{\begin{array}{c}\pi_0\\ A_j, \Gamma \Rightarrow \Delta\end{array}}{\pi_1 : \Gamma \Rightarrow \Delta, \neg A_j}\ \neg\mathsf{R}}{\pi_2 : \Gamma \Rightarrow \Delta, \neg A_0 \vee \neg A_1}\ \vee\mathsf{R}}{\pi : A_0 \wedge A_1, \Gamma \Rightarrow \Delta}\ \neg\mathsf{L}$$

Following the rewrite rules for the three inferences give rise to a derived rewrite rules for conjunction-left identical to the rewrites for disjunction:

$$\begin{aligned} \mathsf{F}^{\pi}_{*}\sigma\vec{x} &\longrightarrow \mathtt{i}_j\left(\mathsf{F}^{\pi_0}_{*}\sigma\vec{x}\right) \\ \mathsf{F}^{\pi}_{i}\sigma(\mathtt{j}z_0z_1)\vec{x} &\longrightarrow \mathsf{F}^{\pi_0}_{i}\sigma z_j\vec{x} \end{aligned}$$

for principal and non-principal rewrites respectively.

4.4.5 Existential

We begin with the left rule:

$$\dfrac{\overset{\pi_0}{A[\alpha/x], \Gamma \Rightarrow \Delta}}{\exists x A, \Gamma \Rightarrow \Delta}\,\exists\mathsf{L}$$

Utilising the abbreviations, the principal rewrite associated with this rule is:

$$\mathsf{F}^{\pi}_{*}\sigma\vec{x} \longrightarrow \mathsf{a}\big(\lambda w.\, \mathsf{F}^{\pi_0}_{*}\big([w/\alpha]\sigma\big)\vec{x}\big)$$

Observe that the type of $\lambda w.\, \mathsf{F}^{\pi_0}_{*}\big([w/\alpha]\sigma\big)\vec{x}$ is $\iota \to \langle A\rangle$, so the right-hand term above is well-typed.

The non-principal rewrite for $\exists\mathsf{L}$ utilise pattern-matching, separating evidence for $\exists x A$ into a concrete witness for the quantifier (passed to the substitution stack) and evidence for the immediate subformula (used to continue the computation):

$$\mathsf{F}^{\pi}_{i}\sigma(\mathsf{e}tz)\vec{x} \longrightarrow \mathsf{F}^{\pi_0}_{i}([t/\alpha]\sigma)z\vec{x}$$

The right rule for existential quantification provides the only rewrites which interact with the substitution stack. Suppose π is given as:

$$\dfrac{\overset{\pi_0}{\Gamma \Rightarrow \Delta, A[t/x]}}{\Gamma \Rightarrow \Delta, \exists x A}\,\exists\mathsf{R} \quad \text{with rewrites} \quad \begin{cases} \mathsf{F}^{\pi}_{*}\sigma\vec{x} \longrightarrow \mathsf{e}(t\cdot\sigma)(\mathsf{F}^{\pi_0}_{*}\sigma\vec{x}) \\ \mathsf{F}^{\pi}_{i}\sigma\vec{x}(\mathsf{a}z) \longrightarrow \mathsf{F}^{\pi_0}_{i}\sigma\vec{x}(z(t\cdot\sigma)) \end{cases}$$

This is again a pattern matching rule where we require that $z : \iota \to \langle A\rangle$. Note that the rule is type-preserving as $\mathsf{a}z : \langle \exists x A\rangle$. In the principal case, the rewrite returns partial evidence for the existential quantifier in the form of the literal witness t guarded by the maintained substitution stack, and a further computation $(\mathsf{F}^{\pi_0}_{*}\cdots)$ for evidence for the immediate subformula. The point of the substitution stack is to possibly instantiate values for the free variables in t. If t contains a free variable, say α, which is associated to a strong quantifier in a cut formula (further down the proof than π), then the rewrite for the corresponding instance of $\exists\mathsf{L}$ will have inserted into the stack σ a substitution of the form $[u/\alpha]$ which, by the principal rewrite above, is recorded alongside the term t. In the non-principal rewrite, the function for which z is a placeholder is expecting input of type ι and is correspondingly fed the witness (t) an the current substitution (σ). The rewrites for this pair of inference rules will be revisited in the examination of the cut rule, below.

Using $\exists\mathsf{L}$ and negation inferences, derived rules for universal quantification are available in the obvious way. The reader can check that the derived principal and non-principal production rules for these inferences are identical to those above.

4.4.6 Cut

Suppose we are given a proof ending with an application of cut:

$$\dfrac{\begin{matrix}\pi_0 \\ \Gamma \Rightarrow \Delta, C\end{matrix} \qquad \begin{matrix}\pi_0 \\ C, \Lambda \Rightarrow \Pi\end{matrix}}{\pi : \Gamma\Lambda \Rightarrow \Delta, \Pi}\,\mathsf{cut}$$

To describe the rewrite rules for the non-terminals corresponding to this proof it will be easiest to adopt some notational conventions. For each index i corresponding to a formula occurrence among $\Gamma, \Lambda, \Delta, \Pi$ in the conclusion let i' be the index of the corresponding formula occurrence in the left or right premise depending on whether the formula occurrence is in Γ, Δ or Λ, Π. In the former case we say that i is a *left index*, in the latter case we say it is a *right index*. Let $*$ denote the index of the cut formula occurrence in the context of either premise. The rewrite rule can now be written as:

$$\mathsf{F}^{\pi}_{i}\sigma\vec{w}\vec{x}\vec{y}\vec{z} \longrightarrow \begin{cases} \mathsf{F}^{\pi_0}_{i'}\sigma\vec{w}\vec{y}(\mathsf{F}^{\pi_1}_{*}\sigma\vec{x}\vec{z}) & i \text{ a left index} \\ \mathsf{F}^{\pi_1}_{i'}\sigma(\mathsf{F}^{\pi_0}_{*}\sigma\vec{w}\vec{y})\vec{x}\vec{z} & i \text{ a right index} \end{cases}$$

where $\vec{w}$, $\vec{x}$, $\vec{y}$, $\vec{z}$ are sequences of variables typed in correspondence with Γ, Λ, Δ, Π, excluding the i-th formula in the sequence.

It is most appropriate to examine the rewrite for cut in case of reductive cut elimination. Thus, consider the case above in which the cut formula is existentially quantified, $C = \exists x D$, and is principal in each of the two subproofs:

$$\pi_0 = \left\{ \begin{matrix} \dfrac{\begin{matrix}\hat{\pi}_0 \\ \Gamma \Rightarrow \Delta, D[t/x]\end{matrix}}{\Gamma \Rightarrow \Delta, \exists x D}\,\exists\mathsf{R} \end{matrix} \right. \quad \left. \begin{matrix} \dfrac{\begin{matrix}\hat{\pi}_1 \\ D[\alpha/x], \Gamma \Rightarrow \Delta\end{matrix}}{\exists x D, \Gamma \Rightarrow \Delta}\,\exists\mathsf{L} \end{matrix} \right\} = \pi_1$$
$$\dfrac{\qquad\qquad\qquad\qquad\qquad\qquad}{\Gamma, \Lambda \Rightarrow \Delta, \Pi}\,\mathsf{cut}$$

The rewrites for the three inferences combine in the following way. For i a left index:

$$\begin{aligned} \mathsf{F}^{\pi}_{i}\sigma\vec{w}\vec{x}\vec{y}\vec{z} &\longrightarrow \mathsf{F}^{\pi_0}_{i'}\sigma\vec{w}\vec{y}(\mathsf{F}^{\pi_1}_{*}\sigma\vec{x}\vec{z}) & \\ &\longrightarrow \mathsf{F}^{\pi_0}_{i'}\sigma\vec{w}\vec{y}\big(\mathsf{a}\big(\lambda w.\, \mathsf{F}^{\hat{\pi}_1}_{*}\big([w/\alpha]\sigma\big)\vec{x}\vec{z}\big)\big) & \text{reducing } \mathsf{F}^{\pi_1}_{0} \\ &\longrightarrow \mathsf{F}^{\hat{\pi}_0}_{i'}\sigma\vec{w}\vec{y}\big(\big(\lambda w.\, \mathsf{F}^{\hat{\pi}_1}_{*}\big([w/\alpha]\sigma\big)\vec{x}\vec{z}\big)(t\cdot\sigma)\big) & \text{reducing } \mathsf{F}^{\pi_0}_{i'} \\ &\longrightarrow \mathsf{F}^{\hat{\pi}_0}_{i'}\sigma\vec{w}\vec{y}\big(\mathsf{F}^{\hat{\pi}_1}_{*}\big([t\cdot\sigma/\alpha]\sigma\big)\vec{x}\vec{z}\big) & \text{reducing } \lambda w \end{aligned}$$

For i a right index we derive instead:

$$\begin{aligned} \mathsf{F}^{\pi}_{i}\sigma\vec{w}\vec{x}\vec{y}\vec{z} &\longrightarrow \mathsf{F}^{\pi_1}_{i'}\sigma(\mathsf{F}^{\pi_0}_{*}\sigma\vec{w}\vec{y})\vec{x}\vec{z} & \\ &\longrightarrow \mathsf{F}^{\pi_1}_{i'}\sigma\big(\mathsf{e}(t\cdot\sigma)\big(\mathsf{F}^{\hat{\pi}_0}_{*}\sigma\vec{w}\vec{y}\big)\big)\vec{x}\vec{z} & \text{reducing } \mathsf{F}^{\pi_0}_{*} \\ &\longrightarrow \mathsf{F}^{\hat{\pi}_1}_{i'}([t\cdot\sigma/\alpha]\sigma)\big(\mathsf{F}^{\hat{\pi}_0}_{*}\sigma\vec{w}\vec{y}\big)\vec{x}\vec{z} & \text{reducing } \mathsf{F}^{\pi_1}_{i'} \end{aligned}$$

In the next section we observe that the final terms of the two derivations above induce equivalent languages as the ‘reduced’ cut:

$$\dfrac{\begin{matrix}\pi_0 \\ \Gamma \Rightarrow \Delta, D[t/x]\end{matrix} \quad \begin{matrix}\hat{\pi}_1[t/\alpha] \\ D[t/x], \Gamma \Rightarrow \Delta\end{matrix}}{\Gamma, \Lambda \Rightarrow \Delta, \Pi}\,\mathsf{cut}$$

4.4.7 Contraction

We give an example of the contraction inference combined with a cut:

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\vdots}{5 : A \Rightarrow C_0, C_1}\,\vee\mathsf{R}}{4 : A \Rightarrow C_0, C_0 \vee C_1}\,\mathsf{eR}}{3 : A \Rightarrow C_0 \vee C_1, C_0}\,\vee\mathsf{R}}{2 : A \Rightarrow C_0 \vee C_1, C_0 \vee C_1}\,\mathsf{cR}}{1 : A \Rightarrow C_0 \vee C_1} \quad \dfrac{\begin{matrix}\vdots \\ 7 : C_0 \Rightarrow B\end{matrix} \quad \begin{matrix}\vdots \\ 8 : C_1 \Rightarrow B\end{matrix}}{6 : C_0 \vee C_1 \Rightarrow B}\,\vee\mathsf{L}}{0 : A \Rightarrow B}\,\mathsf{cut}$$

Subproofs of the above proof are referred to via the numeral labels, so F^4_i refers to F^π_i where π is the subproof above with endsequent labelled 4. Let $i = 0$ be the single left index for the concluding cut and $y : \langle B \rangle$. The following derivation obtains.

$$\begin{aligned}
\mathsf{F}^0_i \sigma y &\longrightarrow \mathsf{F}^1_i \sigma\big(\mathsf{F}^6_0 \sigma y\big) && \\
&\longrightarrow \mathsf{F}^2_i \sigma\big(\mathsf{F}^6_0 \sigma y\big)\big(\mathsf{F}^6_0 \sigma y\big) && \text{reducing } \mathsf{F}^1_i \\
&\longrightarrow \mathsf{F}^2_i \sigma\big(\mathsf{F}^6_0 \sigma y\big)\big(\mathtt{j}(\mathsf{F}^7_0 \sigma y)(\mathsf{F}^8_0 \sigma y)\big) && \text{reducing } \mathsf{F}^6_0 \\
&\longrightarrow \mathsf{F}^3_i \sigma\big(\mathsf{F}^6_0 \sigma y\big)\big(\mathsf{F}^7_0 \sigma y\big) && \text{reducing } \mathsf{F}^2_i \\
&\longrightarrow \mathsf{F}^4_i \sigma\big(\mathsf{F}^7_0 \sigma y\big)\big(\mathtt{j}(\mathsf{F}^7_0 \sigma y)(\mathsf{F}^8_0 \sigma y)\big) && \text{reducing } \mathsf{F}^3_i \text{ and } \mathsf{F}^6_0 \\
&\longrightarrow \mathsf{F}^5_i \sigma\big(\mathsf{F}^7_0 \sigma y\big)\big(\mathsf{F}^8_0 \sigma y\big) && \text{reducing } \mathsf{F}^4_i
\end{aligned}$$

For $i = 1$ being the right index and $x : [A]$:

$$\begin{aligned}
\mathsf{F}^0_i \sigma x &\longrightarrow \mathsf{F}^6_1 \sigma\big(\mathsf{F}^1_1 \sigma x\big) && \\
&\longrightarrow \mathsf{F}^6_1 \sigma\big(\mathtt{p}\big(\mathsf{F}^2_1 \sigma x\big)\big) \parallel \mathsf{F}^6_1 \sigma\big(\mathtt{p}\big(\mathsf{F}^2_2 \sigma x\big)\big) && \text{reducing } \mathsf{F}^1_m \\
&\longrightarrow \mathtt{p}\big(\lambda v.\, \mathsf{F}^6_1 \sigma\big(\mathsf{F}^2_1 \sigma x\big(\mathsf{F}^6_0 \sigma v\big)\big)\big) \parallel \mathtt{p}\big(\lambda v.\, \mathsf{F}^6_1 \sigma\big(\mathsf{F}^2_2 \sigma x\big(\mathsf{F}^6_0 \sigma v\big)\big)\big) && \text{Peirce reduction}
\end{aligned}$$

Formally, the above terms cannot be further reduced as the above abstractions are expressed via composition. Viewing the λv as a formal construct it is possible simulate a continuation of the reduction. Note, the non-terminals F^6_1, F^2_1 and F^2_2 rely on pattern-matching, though F^6_0 can be reduced:

$$\begin{aligned}
&\longrightarrow \mathtt{p}\big(\lambda v.\, \mathsf{F}^6_1 \sigma\big(\mathsf{F}^2_1 \sigma x\big(\mathtt{j}(\mathsf{F}^7_0 \sigma v)(\mathsf{F}^8_0 \sigma v)\big)\big)\big) \parallel \cdots && \text{reducing } \mathsf{F}^6_0 \\
&\longrightarrow \mathtt{p}\big(\lambda v.\, \mathsf{F}^6_1 \sigma\big(\mathsf{F}^4_2 \sigma x\big(\mathsf{F}^7_0 \sigma v\big)\big)\big) \parallel \cdots && \text{reducing } \mathsf{F}^2_1 \text{ \& } \mathsf{F}^2_2
\end{aligned}$$

$$\longrightarrow \mathtt{p}\big(\lambda v.\, \mathsf{F}_1^6\sigma\big(\mathtt{i}_1\big(\mathsf{F}_2^5\sigma x\big(\mathsf{F}_0^7\sigma v\big)\big)\big)\big) \parallel \cdots \qquad \text{reducing } \mathsf{F}_2^4 \ \& \ \ldots$$
$$\longrightarrow \mathtt{p}\big(\lambda v.\, \mathsf{F}_1^8\sigma\big(\mathsf{F}_2^5\sigma x\big(\mathsf{F}_0^7\sigma v\big)\big)\big) \parallel \mathtt{p}\big(\lambda v.\, \mathsf{F}_1^7\sigma\big(\mathsf{F}_1^5\sigma x\big(\mathsf{F}_0^8\sigma v\big)\big)\big) \quad \text{reducing } \mathsf{F}_1^6 \ \& \ \ldots$$

The final term is what is computed from nested cuts:

$$\dfrac{\dfrac{\dfrac{\dfrac{\vdots}{A \Rightarrow C_0, C_1} \quad \dfrac{\vdots}{C_0 \Rightarrow B}}{A \Rightarrow B, C_1}\,\mathsf{cut} \qquad \dfrac{\vdots}{C_1 \Rightarrow B}}{A \Rightarrow B, B}\,\mathsf{cut}}{A \Rightarrow B}\,\mathsf{cR}$$

The connection between Herbrand schemes and cut elimination is not accidental and is the main focus of the second part of this article.

5 Languages

We can now show how to read off Herbrand expansions for Σ_1-formulas from terms in the language of a Herbrand scheme.

5.1 Final terms and expansions

Definition 8. The *final terms* are the terms of (any) Σ_1 formula type that contain neither the Peirce constructor $\mathtt{p}$ nor any non-terminals.

Notice that the final terms are those generated by the grammar where F ranges over formula types:

$$\mu ::= \varepsilon \mid \bot_F \mid \mathtt{i}_i\mu \mid \mathtt{j}\mu\mu \mid \mathtt{e}t\mu$$

where t ranges over extended $\mathscr{L}$-terms. For every final term μ there is a Σ_1 formula A with $\mu : [A]$. We employ symbols μ and ν (with adornments) as meta-variables for final terms. Inspection of the production rules for the extractor non-terminals yields

Proposition 6. *If π is a proof of a Σ_1 sequent then $\mathscr{L}(\pi)$ comprises final terms only.*

Definition 9. Let $A \in \Sigma_1$ and $\mu : [A]$ a final term. The *μ-expansion* of A is the quantifier-free formula $A\{\mu\}$ defined by

$$\begin{aligned} P\vec{s}\{\varepsilon\} &= P\vec{s} & (A_0 \vee A_1)\{\mathtt{i}_i\mu\} &= A_i\{\mu\} \\ \neg P\vec{s}\{\varepsilon\} &= \neg P\vec{s} & \neg(A_0 \vee A_1)\{\mathtt{j}\mu_0\mu_1\} &= \neg A_0\{\mu_0\} \wedge \neg A_1\{\mu_1\} \\ A\{\bot\} &= \bot & (\exists x A)\{\mathtt{e}s\mu\} &= (A\{\mu\})[\mathsf{Val}(s)/x] \\ (\neg\neg A)\{\mu\} &= A\{\mu\} \end{aligned}$$

If U is a set of final terms all of type $[A]$ we write $A\{U\}$ for $\bigvee_{u \in U} A\{u\}$; note that this may in general be an infinite disjunction. An *expansion* of a Σ_1 sequent $\Gamma \Rightarrow A$ is any set of final terms $\mu_i : [A], i \in I$. If U is an expansion and the (possibly infinite) disjunction $\bigvee \Gamma^{\neg} \vee A\{U\}$ is valid, we call $\{\mu_i\}_{i \in I}$ a *valid* expansion of the sequent.

Theorem 1. *Let π be an (essentially) cut-free proof of a Σ_1 sequent $\Gamma \Rightarrow A$. Then $\mathscr{L}(\pi)$ is a non-empty and finite, and valid expansion of $\Gamma \Rightarrow A$.*

Proof. By induction on π. □

5.2 Language subsumption and equivalence

We will make frequent use of the concepts of language *subsumption* and *equivalence*, which we introduce now. We begin with final terms.

Definition 10. For extended terms $s, t : \iota$ we define $s \sim t$ iff $\mathsf{Val}(s) = \mathsf{Val}(t)$. The subsumption relation $\sqsubseteq$ on final terms is the smallest reflexive and transitive relation satisfying the following clauses for all final terms $\mu, \nu, \xi : F$.

1. $\bot_F \sqsubseteq \mu$.
2. If $s \sim t : \iota$ and $\mu \sqsubseteq \nu$, then $\mathtt{e}s\mu \sqsubseteq \mathtt{e}t\nu$.
3. If $\mu \sqsubseteq \nu$ then $\mathtt{i}_i\mu \sqsubseteq \mathtt{i}_i\nu$, $\mathtt{j}\mu\xi \sqsubseteq \mathtt{j}\nu\xi$ and $\mathtt{j}\xi\mu \sqsubseteq \mathtt{j}\xi\nu$.

Given sets U, V of final terms we write $U \sqsubseteq V$ if, for all $u \in U$ there is $v \in V$ with $u \sqsubseteq v$. We write $u \sqsubseteq V$ rather than $\{u\} \sqsubseteq V$ and $U \sqsubseteq v$ rather than $U \sqsupseteq \{v\}$.

A simple induction confirms that $\sqsubseteq$ preserves type: if $\mu \sqsubseteq \nu$ then μ and ν are terms of the same type. The following observation is immediate.

Proposition 7. *If $U \sqsubseteq V$ are sets of final terms of type $[A]$ then $A\{U\} \to A\{V\}$ is a tautology. In particular, if U is a valid expansion for a sequent $\Gamma \Rightarrow A$ then so is V.*

The next definition extends the subsumption relation to arbitrary terms. In the following, by a *context* $c(\cdot)$ we mean some term c of Σ_1-type with a distinguished variable z of arbitrary type U, and we write $c(t)$ for the substitution $c[t/z]$ under where $t : U$.

Definition 11. Given arbitrary terms $s, t : U$ of some type, we say that t *subsumes* s if, for every context $c(\cdot)$ and every final term μ such that $c(s) \longrightarrow \mu$, there is a final term ν such that $c(t) \longrightarrow \nu$ and $\mu \sqsubseteq \nu$. We write $s \sqsubseteq t$ if t *subsumes* s. We write $s \equiv t$, and say that s, t are *equivalent*, if $t \sqsubseteq s$ and $s \sqsubseteq t$.

The next proposition confirms that the two definitions above agree on final terms.

Proposition 8. *For final terms μ, ν, $\mu \sqsubseteq \nu$ (in the sense of defn 10) iff ν subsumes μ (in the sense of defn 11).*

Proof. The right to left direction is immediate by choosing the trivial context. For the other direction we suppose $\mu \sqsubseteq \nu$ (per definition 10) and show, by induction on the length of reduction sequences, that if $c(\mu) \longrightarrow \mu'$, then $c(\nu) \longrightarrow \nu'$ for some $\mu' \sqsubseteq \nu'$. □

Equivalence and subsumption provide a means to compare the language of Herbrand schemes. By showing, for instance, that S_π is subsumed by S_ρ it follows that the Herbrand expansion computed from π is implicit in the expansion computed from ρ. In the following we describe a reduction on proofs via a form of cut elimination and establish that if π can be reduced to ρ then $\mathsf{S}_\rho \sqsubseteq \mathsf{S}_\pi$. Applying Theorem 1 will lead to the observation that if π can be reduced to an essentially cut-free proof then $\mathscr{L}(\pi)$ is a valid expansion of the endsequent.

The following observations will become useful:

Proposition 9. *$t \cdot \sigma \equiv t[\mathsf{Val}(\sigma)] \cdot \bot$ for every extended term $t \cdot \sigma$.*

Proof. Proposition 8 and the definition of value. □

Proposition 10. *If* $t \longrightarrow s$ *then* $s \sqsubseteq t$.

Proof. Immediate. ☐

Proposition 11. *If* $t \longrightarrow s$ *via a production in which a proof non-terminal is* not *principal then* $s \equiv t$.

Proof. It suffices to show $t \sqsubseteq s$. Let $c(\cdot)$ be a context and suppose $c(t) \longrightarrow \mu$. By induction on the length of this reduction it can be shown that $c(s) \longrightarrow \mu$. ☐

Proposition 12. *Let* A *be a formula and* σ *a substitution for the free variables of* A. *Then* $\mathsf{C}_A \equiv \mathsf{C}_{A[\sigma]}$.

Proof. Obvious by inspection of the rewrite rules. ☐

We will use Proposition 12 without further mention, in practice identifying the non-terminals C_A and $\mathsf{C}_{A[\sigma]}$.

5.3 Substitutions, regular terms and regular stacks

Definition 12. The *bound variables* of a substitution stack σ and extended individual term $t \cdot \sigma$ is the set $\mathsf{BV}(\sigma)$ and $\mathsf{BV}(t \cdot \sigma)$ respectively, given by

$$\mathsf{BV}(\bot) = \emptyset \qquad \mathsf{BV}([u/\alpha]\sigma) = \mathsf{BV}(\sigma) \cup \{\alpha\} \qquad \mathsf{BV}(t \cdot \sigma) = \mathsf{BV}(\sigma)$$

The *free variables* of an individual term t is the set $\mathsf{FV}(t)$ of free variables occurring in t. For a substitution stack σ and extended individual term $t \cdot \sigma$, the free variables are defined by mutual recursion:

$$\mathsf{FV}(\bot) = \emptyset \quad \mathsf{FV}([u/\alpha]\sigma) = \mathsf{FV}(u) \cup \mathsf{FV}(\sigma) \quad \mathsf{FV}(t \cdot \sigma) = (\mathsf{FV}(t) \setminus \mathsf{BV}(\sigma)) \cup \mathsf{FV}(\sigma).$$

Definition 13. An extended term t is *regular* if every extended individual subterm s of t is such that $\mathsf{FV}(s) = \emptyset$. A substitution stack σ is regular if every extended individual subterm of σ is regular.

The next two observations about regular substitution stacks hold by inspection.

Proposition 13. *If* σ *is a regular substitution stack then* $\mathsf{FV}(\sigma) = \emptyset$.

Proposition 14. *Let* π *be a regular proof and suppose* $\mathsf{S}_\pi \longrightarrow t$. *Then every substitution stack appearing in* t *is regular.*

Our main task in this subsection is to establish a connection between regular substitution stacks and actual substitutions performed on regular derivations:

Proposition 15. *Let* π *be any regular proof and* σ *a regular substitution stack. Then*

$$\mathsf{F}_i^\pi \sigma \equiv \mathsf{F}_i^{\pi[\mathsf{Val}(\sigma)]} \bot.$$

Proof. See Appendix A.1. ☐

6 Extracting computational content from proofs: examples

In this section we use Herbrand schemes to analyse the computational content of three concrete proofs. The first is a proof of the "Drinker paradox". The second is a proof of the infinite pigeonhole principle from suitable Π_1-assumptions, and the third is a proof of the finite pigeonhole principle proved by a cut on the infinite pigeonhole principle. The key feature of our Herbrand schemes that we want to illustrate with these examples is *compositionality*. While Herbrand schemes are designed to extract Herbrand disjunctions from proofs of Σ_1-sequents, they do more than that. We shall see that, using the rewrite rules for proof non-terminals we can extract meaningful programs or encapsulated "modules" from arbitrary proofs, regardless of the shape of the end sequent. Once the computational content of a proof π has been made explicit, it can then be used and re-used for Herbrand extraction from proofs obtained by cuts on π.

Intuitively, and borrowing some ideas from game semantics, the programs that we extract from Herbrand schemes can be thought of as representing strategies in a game played between two players, *Prover* and *Refuter*. Prover tries to provide evidence for formulas on the right-hand side of the end sequent, and counter-evidence to formulas on the left. Refuter acts dually, trying to provide evidence for formulas on the left-hand side, and counter-evidence to formulas on the right. More precisely, a proof of a sequent $\Gamma \Rightarrow \Delta$ translate, through the Herbrand scheme, to a procedure generating evidence for each $A \in \Delta$ from evidence for Γ and counter-evidence for $\Delta \setminus \{A\}$, and counter-evidence to each $B \in \Delta$ from evidence for $\Gamma \setminus \{B\}$ and counter-evidence for Δ.

An application of the cut rule with premises $\Gamma_0 \Rightarrow \Delta_0, A$ and $A, \Gamma_1 \Rightarrow \Delta_1$ can be interpreted as composing strategies from the two games. In the game for these premises, Prover has a strategy to evidence the formula A (the left premise) and a strategy to counter-evidence A (the right premise), given (counter-)evidence for the formulas in Γ_0 and Γ_1 (Δ_0 and Δ_1) respectively. Dually, Refuter has a strategy to evidence A in the second in the first game, and counter-evidence A in the first. Prover's strategy for the conclusion to the cut – the sequent $\Gamma_0\Gamma_1 \Rightarrow \Delta_0\Delta_1$ – is a composition of the two strategies: To evidence a $D \in \Delta_1$, we assume Prover is provided evidence for each formula in $\Gamma_0 \cup \Gamma_1$ and counter-evidence for each formula in $\Delta_0 \cup (\Delta_1 \setminus \{D\})$. Prover utilises her strategy for the left premise of the cut to obtain evidence for the formula A and feeds this into the other strategy to extract evidence for D. Evidencing $D \in \Delta_0$ is symmetric: Prover obtains counter-evidence for A from the strategy for the right premise of the cut and combines this with her strategy for the left premise.

To simplify notation a bit, in the following we will write $\tilde{\varepsilon}$ as a common abbreviation for terms of quantifier-free formula type. For example, consider the following proof:

$$\dfrac{\dfrac{\dfrac{\dfrac{Pc \Rightarrow Pc}{\Rightarrow \neg Pc, Pc}\,\neg\mathsf{R}}{\Rightarrow Pc \vee \neg Pc, Pc \vee \neg Pc}\,\vee\mathsf{R}^*}{\Rightarrow Pc \vee \neg Pc}\,\mathsf{cR}}{\Rightarrow \exists x(Pc \vee \neg Px)}\,\exists\mathsf{R}$$

$$\cfrac{}{7.\ D\alpha \Rightarrow D\alpha}\ \mathsf{id}$$

$$\vdots\ \lor\mathsf{R}, \lor\mathsf{R}, \neg\mathsf{R}$$

$$\cfrac{\cfrac{\cfrac{\cfrac{\cfrac{6.\ \Rightarrow \neg Da \lor D\alpha, \neg D\alpha \lor D\beta}{5.\ \Rightarrow \neg Da \lor D\alpha, \forall y(\neg D\alpha \lor Dy)}\ \forall\mathsf{R}(\beta)}{4.\ \Rightarrow \neg Da \lor D\alpha, \exists x \forall y(\neg Dx \lor Dy)}\ \exists\mathsf{R}}{3.\ \Rightarrow \forall y(\neg Da \lor Dy), \exists x \forall y(\neg Dx \lor Dy)}\ \forall\mathsf{R}(\alpha)}{2.\ \Rightarrow \exists x \forall y(\neg Dx \lor Dy), \exists x \forall y(\neg Dx \lor Dy)}\ \exists\mathsf{R}}{1.\ \Rightarrow \exists x \forall y(\neg Dx \lor Dy)}\ \mathsf{cR}$$

$$\cfrac{}{D(\mathsf{f}a) \Rightarrow D(\mathsf{f}a)}\ \mathsf{id}$$

$$\vdots$$

$$\cfrac{\cfrac{\cfrac{\Rightarrow \neg Da \lor D(\mathsf{f}a), \neg D(\mathsf{f}a) \lor D(\mathsf{f}(\mathsf{f}a))}{\Rightarrow \neg Da \lor D(\mathsf{f}a), \exists x(\neg Dx \lor D(\mathsf{f}x))}\ \exists\mathsf{R}}{\Rightarrow \exists x(\neg Dx \lor D(\mathsf{f}x)), \exists x(\neg Dx \lor D(\mathsf{f}x))}\ \exists\mathsf{R}}{\Rightarrow \exists x(\neg Dx \lor D(\mathsf{f}x))}\ \mathsf{cR}$$

Fig. 4 Sequent calculus proof of the Drinker paradox (left) and its skolemised form (right).

From its Herbrand scheme we can extract the terms $\mathtt{e}c(\mathtt{i}_0\varepsilon)$ and $\mathtt{e}c(\mathtt{i}_1\varepsilon)$, which we simplify to $\mathtt{e}c\tilde{\varepsilon}$.

6.1 The drinker paradox

The following classical validity provides a simple yet instructive example:

$$\exists x \forall y(\neg Dx \lor Dy) \tag{1}$$

This formula is often referred to as the Drinker paradox from the interpretation 'there is some person such that, if they drink, then everybody drinks'. The skolemised form is $\exists x(\neg Dx \lor D(\mathsf{f}x))$. Figure 4 presents cut-free sequent calculus proofs of (1) and its skolemisation. The language of skolemised proof is the single final term $\mathtt{e}(a \cdot \bot)(\mathtt{i}_1\epsilon) \parallel \mathtt{e}((\mathsf{f}a) \cdot \bot)(\mathtt{i}_0\epsilon)$. We leave the verification to reader and instead analyse the behaviour of the Herbrand scheme for the unskolemised formula.

Abbreviate $\exists x \forall y(\neg Dx \lor Dy)$ as A, $\forall y(\neg D\gamma \lor Dy)$ as B, $\exists y \neg(\neg D\gamma \lor Dy)$ as $B^\perp$ and $\neg(\neg D\gamma \lor Dy)$ as C, and let π denote the proof of A in Figure 4. The associated Herbrand scheme rewrites as:

$$\mathsf{F}^0_0\sigma \longrightarrow \mathtt{p}(\mathsf{F}^1_0\sigma) \parallel \mathtt{p}(\mathsf{F}^1_1\sigma) \tag{2}$$

The computation stops as no further production rule is applicable. To continue the computation the Peirce constructor $\mathtt{p}$ must be removed and the non-terminals F^1_0 and F^1_1 be supplied an argument corresponding to the copy of the formula A not in focus. We can, however, analyse the continuing computation by providing generic, archetypal inputs of the appropriate type. We get the following rewrites:

$$\mathsf{F}^1_0\sigma(\mathtt{a}(\lambda w.\mathtt{e}tv)) \sqsupseteq \mathtt{e}a(\mathtt{a}(\lambda u.\tilde{\varepsilon})) \tag{3}$$

$$\mathsf{F}^1_1\sigma(\mathtt{a}(\lambda w.\mathtt{e}tv)) \sqsupseteq \mathtt{e}t(a)(\mathtt{a}(\lambda u.\tilde{\varepsilon})) \tag{4}$$

Here we let $t = t(w)$ be a metavariable for a term with a displayed occurrence of the variable w. The 'equations' (2)–(4) can be seen together as describing the behaviour of the program F^0_0, or with the game semantic view, as a strategy for Prover to provide evidence for the end sequent. First, Prover splits the game into two parallel copies played concurrently with the aim of winning in at least one copy ((2)). In the 'left'

$$
\begin{array}{c}
\vdots \\
\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{8.\ \Gamma \Rightarrow I_0\alpha_0(\mathsf{m}\alpha_0\alpha_1), I_1\alpha_1(\mathsf{m}\alpha_0\alpha_1)}{7.\ \Gamma \Rightarrow I_0\alpha_0(\mathsf{m}\alpha_0\alpha_1), \exists y\, I_1\alpha_1 y}\ \exists\mathsf{R}}{6.\ \Gamma \Rightarrow \exists y\, I_0\alpha_0 y, \exists y\, I_1\alpha_1 y}\ \exists\mathsf{R}}{5.\ \Gamma \Rightarrow \exists y\, I_0\alpha_0 y, \forall x\exists y\, I_1 xy}\ \forall\mathsf{R}}{4.\ \Gamma \Rightarrow \forall x\exists y\, I_0 xy, \forall x\exists y\, I_1 xy}\ \forall\mathsf{R}}{3.\ \Gamma \Rightarrow \forall x\exists y\, I_0 xy, \forall x\exists y\, I_0 xy \vee \forall x\exists y\, I_1 xy}\ \vee\mathsf{R}}{2.\ \Gamma \Rightarrow \forall x\exists y\, I_0 xy \vee \forall x\exists y\, I_1 xy, \forall x\exists y\, I_0 xy \vee \forall x\exists y\, I_1 xy}\ \vee\mathsf{R}}{1.\ \Gamma \Rightarrow \forall x\exists y\, I_0 xy \vee \forall x\exists y\, I_1 xy}\ \mathsf{cR}
\end{array}
$$

Fig. 5 Sequent proof π_{iph} of the infinite pigeonhole principle.

copy of the game ((3)), Prover simply provides a candidate witness a. In the 'right' copy ((4)), the proposed witness is instead $t(a)$, using as input the proposed counter-evidence of Refuter which is essentially the function $w \mapsto t(w)$. Clearly the strategy is sound, since a loss for Prover in the 'left' copy of the game means that the literal $Dt(a)$ is false, which means $\forall y(\neg Dt(a) \vee Dy)$ is true and so the witness provided in the 'right' copy of the game succeeds. The point is that we can extract meaningful computational content from arbitrary proofs, which is re-usable and can be applied via applications (cuts). In the special case of proofs for Σ_1 end sequents, we obtain Herbrand disjunctions.

6.2 Infinite pigeonhole principle

The second example is a version of the infinite pigeonhole principle: Given an infinite binary string, two indices are labelled by the same bit.

The formal language $\mathscr{L}$ comprises two binary relation symbols, $=$ and $\leq$, a constant symbol $\mathsf{0}$, two unary function symbols s and f and one binary function symbol m. In the sequel, we employ the following abbreviations of formulas

- $I_0xy \coloneqq x \leq y \wedge \mathsf{f}y = \mathsf{0}$, expressing that y is an index not earlier than x labelled 0.
- $I_1xy \coloneqq x \leq y \wedge \mathsf{f}y = \mathsf{s0}$, expressing that y is an index not earlier than x labelled 1.

Let Γ consists of sufficient Π_1-formulas to enable the leaf sequent in the derivation in Figure 5 is provable. In particular, we assume that Γ contains the formula $\forall x(\mathsf{f}x = \mathsf{0} \vee \mathsf{f}x = \mathsf{s0})$. The intension of m is as the binary max-function on natural numbers and Γ contains Π_1-axioms matching that interpretation. To explore the computational content of this proof we consider how its proof non-terminal rewrites:

$$(1) \quad \mathsf{F}^1_0\sigma \longrightarrow \mathtt{p}(\mathsf{F}^2_0\sigma) \parallel \mathtt{p}(\mathsf{F}^2_1\sigma)$$

The computation stops after this first rewrite since the rewrite rules for F^2_0 and F^2_1 need to pattern match on the input. To understand the computational behaviour of these non-terminals we feed them generic, archetypal input terms built out of variables, constructors and destructors. For F^2_0 we have:

$$(2) \quad \mathsf{F}^2_0\sigma(\mathtt{j}(\mathtt{e}x_0w_0)(\mathtt{e}x_1w_1)) \sqsupseteq \mathtt{i}_0\mathtt{a}(\lambda z.\mathtt{e}(\mathsf{m}(z,x_1))\hat{\varepsilon})$$

$$
\begin{array}{c}
\vdots \\
\hline
\Gamma \Rightarrow I_1\mathsf{c}_1(\mathsf{c}_0)(\mathsf{m}\mathsf{c}_0(\mathsf{c}_1(\mathsf{c}_0)))) \,, I_0\mathsf{c}_0(\mathsf{m}\mathsf{c}_0(\mathsf{c}_1(\mathsf{c}_0))) \\
\hline
\Gamma \Rightarrow I_1\mathsf{c}_1(\mathsf{c}_0)(\mathsf{m}\mathsf{c}_0(\mathsf{c}_1(\mathsf{c}_0)))) \,, I_0\mathsf{c}_0(\mathsf{m}\mathsf{c}_0(\mathsf{c}_1(\mathsf{c}_0))) \vee I_1\mathsf{c}_1(\mathsf{m}\mathsf{c}_0(\mathsf{c}_1(\mathsf{c}_0)))\mathsf{c}_0) \quad \vee\mathsf{R} \\
\hline
\Gamma \Rightarrow I_1\mathsf{c}_1(\mathsf{c}_0)(\mathsf{m}\mathsf{c}_0(\mathsf{c}_1(\mathsf{c}_0)))) \,, \exists y_1(I_0\mathsf{c}_0(\mathsf{m}\mathsf{c}_0(\mathsf{c}_1(\mathsf{c}_0))) \vee I_1\mathsf{c}_1(\mathsf{m}\mathsf{c}_0(\mathsf{c}_1(\mathsf{c}_0)))y_1) \quad \exists\mathsf{R} \\
\hline
\Gamma \Rightarrow I_1\mathsf{c}_1(\mathsf{c}_0)(\mathsf{m}\mathsf{c}_0(\mathsf{c}_1(\mathsf{c}_0)))) \,, \exists y_0\exists y_1(I_0\mathsf{c}_0 y_0 \vee I_1\mathsf{c}_1(y_0)y_1) \quad \exists\mathsf{R} \\
\hline
\Gamma \Rightarrow I_0\mathsf{c}_0\mathsf{c}_0 \vee I_1\mathsf{c}_1(\mathsf{c}_0)(\mathsf{m}\mathsf{c}_0(\mathsf{c}_1(\mathsf{c}_0)))) \,, \exists y_0\exists y_1(I_0\mathsf{c}_0 y_0 \vee I_1\mathsf{c}_1(y_0)y_1) \quad \vee\mathsf{R} \\
\hline
\Gamma \Rightarrow \exists y_1(I_0\mathsf{c}_0\mathsf{c}_0 \vee I_1\mathsf{c}_1(\mathsf{c}_0)y_1) \,, \exists y_0\exists y_1(I_0\mathsf{c}_0 y_0 \vee I_1\mathsf{c}_1(y_0)y_1) \quad \exists\mathsf{R} \\
\hline
\Gamma \Rightarrow \exists y_0\exists y_1(I_0\mathsf{c}_0 y_0 \vee I_1\mathsf{c}_1(y_0)y_1) \,, \exists y_0\exists y_1(I_0\mathsf{c}_0 y_0 \vee I_1\mathsf{c}_1(y_0)y_1) \quad \exists\mathsf{R} \\
\hline
\Gamma \Rightarrow \exists y_0\exists y_1(I_0\mathsf{c}_0 y_0 \vee I_1\mathsf{c}_1(y_0)y_1) \quad \mathsf{cR}
\end{array}
$$

Fig. 6 Skolemized infinite pigeonhole principle.

For F_1^2:

$$(3) \quad \mathsf{F}_1^2\sigma(\mathtt{j}(\mathtt{e}x_0w_0)(\mathtt{e}x_1w_1)) \sqsupseteq \mathtt{i}_1\mathtt{a}(\lambda z.\mathtt{e}(\mathtt{m}(x_0, z))\tilde{\varepsilon})$$

The equations (1) − (3) provide us with a computational analysis of the proof of the infinite pigeon-hole principle, which serves as a re-usable "module".

Viewed as a strategy for Prover, it runs as follows. The aim for Prover is to provide evidence for the right formula of the end sequent, which should consist of either a choice of the left disjunct or right disjunct together with a function from individuals to individuals. If the left disjunct is chosen then the function h provided should map each number n to a larger or equal one which f maps to 0. If the right disjunct is chosen, f should always map $h(n)$ to 1. Prover's strategy is to first ask Refuter for counter-evidence to the same formula, which is expected to consist of a counter-example x_0 to the left disjunct and a counter-example x_1 to the right disjunct. In the rewrites (2) and (3) the term $\mathtt{j}(\mathtt{e}x_0w_0)(\mathtt{e}x_1w_1)$ describes the generic form expected from Refuter's response. Now Prover plays two strategies in two parallel copies of the game: in one copy she chooses the left disjunct and offers the witnessing function $\lambda z.\mathsf{m}(z, x_1)$. In the other copy of the game, she chooses the right disjunct and witnessing function $\lambda z.\mathsf{m}(x_0, z)$. Since the assumptions in Γ entail that f is bounded by 1, if x_0 is a counter-example to the left disjunct then f maps every greater or equal number to 1, and hence the function $\lambda z.\mathsf{m}(x_0, z)$ witnesses the formula $\forall x\exists y I_1(x, y)$. Similarly, if x_1 is a counter-example to the right disjunct then $\lambda z.\mathsf{m}(z, x_1)$ witnesses $\forall x\exists y I_0(x, y)$.

It may be instructive to compare this analysis to the witness extraction for the Skolemised version of the infinite pigeonhole principle. The prenex normal form is:

$$\forall x_0\exists y_0\forall x_1\exists y_1(I_0x_0y_0 \vee I_1x_1y_1)$$

If we Skolemise the universal quantifiers we get the valid formula:

$$\exists y_0\exists y_1(I_0\mathsf{c}_0y_0 \vee I_1\mathsf{c}_1(y_0)y_1)$$

A proof is displayed in Figure 6. From this proof we obtain the substitutions $y_0 \mapsto \mathsf{c}_0, y_1 \mapsto \mathsf{m}\mathsf{c}_0(\mathsf{c}_1(\mathsf{c}_0))$ and $y_0\mathsf{m}\mathsf{c}_0(\mathsf{c}_1(\mathsf{c}_0)), y_1 \mapsto \mathsf{c}_0$. From the Herbrand scheme of the proof we can derive corresponding final terms with the values:

$$\mathsf{e}\mathsf{c}_0(\mathsf{e}(\mathsf{m}\mathsf{c}_0(\mathsf{c}_1(\mathsf{c}_0)))\tilde{\varepsilon}) \parallel \mathsf{e}(\mathsf{m}\mathsf{c}_0(\mathsf{c}_1(\mathsf{c}_0)))(\mathsf{e}\mathsf{c}_0\tilde{\varepsilon})$$

$$\vdots \qquad\qquad\qquad\qquad\qquad\qquad \vdots$$

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{17.\ \Delta, I_0 0\beta_0, I_0(\mathsf{s}\beta_0)\gamma_0 \Rightarrow D\beta_0\gamma_0}{16.\ \Delta, I_0 0\beta_0, I_0(\mathsf{s}\beta_0)\gamma_0 \Rightarrow \exists y\, D\beta_0 y}\exists\mathsf{R}}{15.\ \Delta, I_0 0\beta_0, I_0(\mathsf{s}\beta_0)\gamma_0 \Rightarrow \exists x\exists y\, Dxy}\exists\mathsf{R}}{14.\ \Delta, I_0 0\beta_0, \exists y\, I_0(\mathsf{s}\beta_0)y \Rightarrow \exists x\exists y\, Dxy}\exists\mathsf{L}}{13.\ \Delta, I_0 0\beta_0, \forall x\exists y\, I_0 xy \Rightarrow \exists x\exists y\, Dxy}\exists\mathsf{R}}{12.\ \Delta, \exists y\, I_0 0y, \forall x\exists y\, I_0 xy \Rightarrow \exists x\exists y\, Dxy}\exists\mathsf{L}}{11.\ \Delta, \forall x\exists y\, I_0 xy, \forall x\exists y\, I_0 xy \Rightarrow \exists x\exists y\, Dxy}\exists\mathsf{R}}{10.\ \Delta, \forall x\exists y\, I_0 xy \Rightarrow \exists x\exists y\, Dxy}\mathsf{cL}$$

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{25.\ \Delta, I_0 0\beta_1, I_1(\mathsf{s}\beta_1)\gamma_1 \Rightarrow D\beta_1\gamma_1}{24.\ \Delta, I_0 0\beta_1, I_1(\mathsf{s}\beta_1), \gamma_1 \Rightarrow \exists y\, D\beta_1 y}}{23.\ \Delta, I_1 0\beta_1, I_1(\mathsf{s}\beta_1)\gamma_1 \Rightarrow \exists x\exists y\, Dxy}}{22.\ \Delta, I_1 0\beta_1, \exists y\, I_1(\mathsf{s}\beta_1)y \Rightarrow \exists x\exists y\, Dxy}}{21.\ \Delta, I_1 0\beta_1, \forall x\exists y\, I_1 xy \Rightarrow \exists x\exists y\, Dxy}}{20.\ \Delta, \exists y\, I_1 0y, \forall x\exists y\, I_1 xy \Rightarrow \exists x\exists y\, Dxy}}{19.\ \Delta, \forall x\exists y\, I_1 xy, \forall x\exists y\, I_1 xy \Rightarrow \exists x\exists y Dxy}}{18.\ \Delta, \forall x\exists y\, I_1 xy \Rightarrow \exists x\exists y\, Dxy}$$

$$\dfrac{\pi_{\mathrm{iph}} \qquad \dfrac{10. \qquad 18.}{9.\ \Delta, \forall x\exists y\, I_0 xy \vee \forall x\exists y\, I_1 xy \Rightarrow \exists x\exists y Dxy}\vee\mathsf{L}}{0.\ \Gamma, \Delta \Rightarrow \exists x\exists y D(x,y)}\mathsf{cut}$$

Fig. 7 Sequent proof π_{ph} of the pigeonhole principle from the infinite pigeonhole principle. Dxy abbreviates $x < y \wedge \mathsf{f}x = \mathsf{f}y$, and Δ consists of sufficient Π_1-formulas for the leaves to be provable. Sequents in the right-hand subproof (lines 18–25) are subject to the identical rule applications as the left-hand (lines 10–17).

As a reduction of the infinite pigeonhole principle to a propositional validity, this is perfectly fine. As a computational interpretation however, it is not what we would want. It doesn't tell us anything about how the infinite pigeonhole principle can be used for reasoning in other contexts, i.e. how it can be applied. By contrast, we shall see in the following section how the analysis of π_{iph} can be imported into another example, where we extract a Herbrand disjunction for the finite pigeonhole principle via a cut on π_{iph}.

6.3 Finite pigeonhole principle

In the final example we consider a proof involving a cut on the infinite pigeonhole principle π_{iph}. This example demonstrates how the implicit strategy of the previous example plays out in practice.

The Herbrand scheme associated with this proof can be rewritten in different ways; the rewrite rules for Herbrand schemes are not confluent. We focus on a part of the language of the scheme which is large enough to give a valid Herbrand expansion:

$$\begin{aligned}
&\mathsf{e(m00)(e(m0(s(m00)))}\tilde{\varepsilon})\\
&\parallel \mathsf{e(m00)(e(m(s(m0(s(m00))))(s(m00)))}\tilde{\varepsilon}),\\
&\parallel \mathsf{e(m(s(m00))0)(e(m0(s(m(s(m00))0)))}\tilde{\varepsilon})\\
&\parallel \mathsf{e(m(s(m00))0)(e(m(s(m0(s(m(s(m00))0))))(s(m(s(m00))0)))}\tilde{\varepsilon})
\end{aligned}$$

These four terms represents four possible substitutions for the bound variables x and y in $\exists x\exists y\, Dxy$: $(0,1)$, $(0,2)$, $(1,2)$ and $(1,3)$. We omit the details of the computation. What is interesting is how the strategy extracted from the proof of the infinite pigeonhole principle occurs in the witness extraction. What happens can be explained informally as follows. Recall that the strategy told Prover to create two parallel games, expecting an input from Refuter of the form $\mathsf{j}(\mathsf{e}x_0w_0)(\mathsf{e}x_1w_1)$. In response to a given input of this shape, Prover plays the following moves in each copy of the game: the function $w \mapsto \mathsf{m}wx_1$ in the "left" game, and the function $w \mapsto \mathsf{m}x_0w$ in the "right" game (ignoring the "wrapping" of these functions in constructors and destructors to simplify).

Starting the computation from the Herbrand scheme associated with the proof π_{ph} above, and employing the same abbreviations of notation and ignoring the side-formulas Γ, Δ, we have:

$$\begin{aligned} \mathsf{S}_{\pi_{\mathrm{ph}}} &\longrightarrow \mathsf{E}(\mathsf{F}^0_0\bot) \\ &\longrightarrow \mathsf{E}(\mathsf{F}^9_1\bot(\mathsf{F}^1_0\bot)) \end{aligned}$$

So the strategy obtained from the previous proof of the infinite pigeonhole principle is now fed as an input via the cut. The strategy will now instead be played by Refuter, in response to moves made by Prover as she tries to provide witnesses for the existential quantifiers in $\exists x \exists y\, Dxy$.

Skipping ahead a bit, consider how we can extract witnesses for these quantifiers in the central subproof beginning on line 10. Prover is trying to find values for the eigenvariables β_0, γ_0 which are then given as witnesses to the existential quantifiers. She can find the value of β_0 as follows: First, she proposes 0 as a counterexample to A_0, and applies the function $\lambda w.\ \mathsf{m} w x_1)$ obtained from Refuter's response to 0 to arrive at a value for β_0. This function uses a counterexample x_1 to A_1 as a parameter. The rightmost subproof proposes two such counterexamples: One is simply $\mathsf{0}$, in which case we get the value $\mathsf{m00}$ (i.e., 0) for β_0. The other counter-example proposed is $\mathsf{s}\beta_1$, the successor of β_1, where the value of β_1 found by applying the function $\lambda w. \mathsf{m} x_0 w$ to 0. This function in turn uses a counterexample x_0 to A_0 as a parameter. But Prover has already proposed such a counterexample, namely 0. So in this case we obtain $\mathsf{s(m00)}$, namely 1, as a counterexample to A_1, and when Prover applies the function $\lambda w.\ \mathsf{m} w x_1)$ with this parameter the result is $\mathsf{m0(s(m00))}$, expressing 1. So in this way Prover arrives at the two possible values 0 and 1 for β_0. (In reality, if one goes through the details of the computation, this informally described "passing of values" will be controlled by the rewrite rules for the Peirce operator.)

Now to find the value of γ_0 Prover first proposes the term $\mathsf{s}\beta_0$ as another counterexample to A_0. It has just been established that β_0 denotes either 0 or 1, and so $\mathsf{s}\beta_0$ denotes 1 or 2. Again, Refuter responds with the function $\lambda w.\ \mathsf{m} w x_1$, and the parameter x_1 is either $\mathsf{0}$ or $\mathsf{s}\beta_1$ where the value of β_1 is found by applying the function $\lambda w.\ \mathsf{m} x_0 w$ to the counterexample to A_0 already proposed by Prover, namely $\mathsf{s}\beta_0$ (which therefore also becomes the current value of the parameter x_0). In the case where β_0 denotes 0 we thus get:

$$\begin{aligned}
\gamma_0 &\coloneqq (\lambda w.\, \mathsf{m} w x_1)(\mathsf{s}\beta_0) & \\
&\coloneqq \mathsf{m}(\mathsf{s}\beta_0)x_1 & \beta\text{-reduction} \\
&\coloneqq \mathsf{m}1x_1 & \text{using } \beta_0 \mapsto 0 \\
&\coloneqq \mathsf{m}10 \;\text{ or }\; \mathsf{m}1(\mathsf{s}\beta_1) & x_1 \text{ is 0 or } \mathsf{s}\beta_1 \\
&\coloneqq \mathsf{m}10 \;\text{ or }\; \mathsf{m}1(\mathsf{s}(\lambda w.\, \mathsf{m}x_0 w)(\mathsf{s}\beta_0)) & \text{using inferred value of } \beta_1 \\
&\coloneqq \mathsf{m}10 \;\text{ or }\; \mathsf{m}1(\mathsf{s}(\mathsf{m}x_0(\mathsf{s}\beta_0))) & \beta\text{-reduction} \\
&\coloneqq \mathsf{m}10 \;\text{ or }\; \mathsf{m}1(\mathsf{s}(\mathsf{m}(\mathsf{s}\beta_0)(\mathsf{s}\beta_0))) & x_0 \text{ is now } s\beta_0! \\
&\coloneqq \mathsf{m}10 \;\text{ or }\; \mathsf{m}1(\mathsf{s}(\mathsf{m}11)) & \text{using } \beta_0 \mapsto 0 \\
&\coloneqq 1 \;\text{ or }\; 2 &
\end{aligned}$$

Here, $t \coloneqq s$ means informally that t evaluates to s through the values chosen for variables according to Prover's strategy. This process of evaluation results in the pair of substitutions $(0,1)$ and $(0,2)$ for (x,y). Repeating the same argument under the assumption $\beta_0 \mapsto 1$ provides the two remaining substitutions $(1,2)$ and $(1,3)$. The same substitutions are derived from the rightmost subproof in a similar way.

Essentially the same proof was analysed in [1], yielding the set of substitutions:

$$\{(0,1),(0,2),(1,2),(1,3),(2,3)\}$$

Here, we arrived at the set:

$$\{(0,1),(0,2),(1,2),(1,3)\}$$

.

7 Multicut analysis

In this section we introduce the multicut rule, which will form the basis of the approach to cut-elimination we will follow.

7.1 Multicut instances

It will be convenient to regard a *formula occurrence* in a sequent formally as a pair (S,i) where S is a sequent and i is an index indicating the i-th formula occurrence in the sequent. We write $p(S,i)=0$ if (S,i) is a left formula occurrence and $p(S,i)=1$ if (S,i) is a right formula occurrence. Given a formula occurrence (S,i) we write $S[i]$ for the i-th formula in the sequent S. So we can think of formula occurrences as "addresses" of formulas appearing in sequents, and the operation $S[-]$ retrieves the actual formula from its address. The *length* $l(S)$ of a sequent $S = \Gamma \Rightarrow \Delta$ is $|\Gamma| + |\Delta|$.
Definition 14. A quasi-instance of multicut is a pair $(\mathbb{S}, \triangledown)$ where $\mathbb{S}$ is a non-empty finite multi-set of sequents and $\triangledown$ is a symmetric binary relation over the set of formula occurrences from sequents in $\mathbb{S}$. Given sequents S, U we write $S \triangledown U$ if there are i, j

such that $(S,i) \bigtriangledown (U,j)$. The (undirected) graph with vertices $\mathbb{S}$ and edges $\{S, S'\}$ such that $S \bigtriangledown S'$ is called the *cut graph* of the multicut quasi-instance, and denoted $G(\mathbb{S}, \bigtriangledown)$.

A multicut quasi-instance $(\mathbb{S}, \bigtriangledown)$ is called an *instance* of multicut if the following constraints are satisfied:

- The cut graph $G(\mathbb{S}, \bigtriangledown)$ is a tree, i.e. connected and acyclic.
- For each formula occurrence (S,i) there is at most one formula occurrence (U,j) such that $(S,i) \bigtriangledown (U,j)$.
- For each pair of sequents S, U, there is at most one formula occurrence (S,i) and at most one formula occurrence (U,j) such that $(S,i) \bigtriangledown (U,j)$.
- If $(S,i) \bigtriangledown (U,j)$ then:
  - $S[i] = U[j]$, and
  - $p(S,i) = 1 - p(U,j)$.

If there is some (U,j) with $(S,i) \bigtriangledown (U,j)$ then we call (U,j) the *cut companion* of (S,i). A formula occurrence (S,i) is said to be a *cut formula occurrence* of a multicut if it has a cut companion and a *side formula occurrence* otherwise.

If $(\mathbb{S}, \bigtriangledown)$ is a multicut instance, we say that the inference with sequents $\mathbb{S}$ as premises and the associated cut relation $\bigtriangledown$, written in short-hand as:

$$\frac{(\mathbb{S}, \bigtriangledown)}{U}$$

where U is some sequent, is *valid*, if there is a one-to-one map f from formula occurrences (U,i) for $i < l(U)$ to side formula occurrences in the multicut instance. We often speak rather informally of the formula occurrence corresponding to (U,i) in the multicut, by which we mean $f(U,i)$ for some map f which we usually leave implicit.

Definition 15. We say that a formula occurrence (S,i) *directly depends on* a formula occurrence (U,j), written $(U,j) \prec (S,i)$, if there is some $k < l(S)$ with $k \neq i$ and $(U,j) \bigtriangledown (S,k)$. We say that (S,i) depends on (U,j) if $(U,j) \prec^+ (S,i)$, where $\prec^+$ denotes the transitive closure of $\prec$.

Intuitively, (S,i) directly depends on (U,j) if an output for the latter is needed as input for the calculation of an output for the former. As a direct consequence of the definition of a multicut instance, the dependency relation is well-founded.

An alternative formulation of dependency may help to clarify it. Given a multicut instance $(\mathbb{S}, \bigtriangledown)$ and two distinct sequents $S, U \in \mathbb{S}$, there must be a unique sequence of formula occurrences:

$$(S, i_0) \bigtriangledown (V_1, j_n) \neq (V_1, i_n) \bigtriangledown \dots \bigtriangledown (V_{n-1}, i_{n-1}) \neq (V_{n-1}, j_{n-1}) \bigtriangledown (U, i_n)$$

We write $\mathsf{link}(S,U)$ for the index i_0 and, conversely, $\mathsf{link}(U,S)$ for the index i_n.

Proposition 16. *Given two formula occurrences (S,i) and (U,j) with $S \neq U$, we have $(S,i) \prec^+ (U,j)$ if and only if $i = \mathsf{link}(S,U)$ and $j \neq \mathsf{link}(U,S)$.*

Proof. See Appendix A.2. □

An *extended multicut instance* $(\mathbb{S}, \bigtriangledown, \pi)$ is an instance of multicut together with an assignment π of a multicut-free (but not necessarily cut-free) proof π_S to each sequent

$S \in \mathbb{S}$. Given an extended multicut instance we define the *canonical input term* $I(S,i)$ and the *canonical output term* $O(S,i)$ associated with a formula occurrence (S,i) by well-founded induction on $\prec$:

- If (S,i) is a left side formula occurrence of the multicut, then $I(S,i) = \mathsf{C}_{[S[i]]}$.
- If (S,i) is a right side formula occurrence of the multicut, then $I(S,i) = \mathsf{C}_{\langle S[i] \rangle}$.
- If (S,i) has cut companion (U,j) then $I(S,i) = O(U,j)$.
- $O(S,i) = F_i^{\pi_S} \bot I(S,0) \dots I(S,i-1) I(S,i+1) \dots I(S,l(S)-1)$

Note that this definition is indeed by well-founded induction, as each term $I(S,j)$ for $j \neq i$ is either a constant non-terminal or given recursively as $O(U,k)$ for $(U,k) \bigtriangledown (S,j)$ and hence $(U,k) \prec (S,i)$. We extend the notation π_S for $S \in \mathbb{S}$ by writing also π_{S_0} for the subtree of π_S generated by S_0, if S_0 is a premise of S in π_S.

7.2 Multicut guarded proofs

Definition 16. Let π be any regular proof. The *end piece* of π is the smallest sub-tree π_{ep} of π containing the end sequent and closed under premises of all rules except cut and multicut. Given a sequent S in π_{ep}, define $\pi_{ep}|S$ to be the part of π_{ep} containing all sequents from the end sequent of π up to and including S. We say that π is *multicut guarded* if each leaf of π_{ep} is either an axiom or the conclusion of a multicut rule.

Proposition 17. *For every constant non-terminal C_A where A is a Σ_1-formula, there is a final term t with $C_A \longrightarrow t$.*

Proof. Obvious. □

Definition 17. Let π be a multicut guarded regular proof of a Σ_1-sequent. For each sequent S belonging to π_{ep}, and each index $i < l(S)$, we define the *target term set* $T(S,i,\pi)$ by well-founded induction on the descendant relation in the tree π_{ep}. For the base case, where S is a leaf in π_{ep}, there are two possible cases: either S is an axiom, or S is the conclusion of a multicut. If S is an axiom then it consists of two occurrences of the same formula, which is quantifier free. We set $T(S,i,\pi) = \{\varepsilon\}$ for $i \in \{0,1\}$.

The remaining case is where S is the conclusion of a multicut. We can define an extended multicut instance $(\mathbb{S}, \bigtriangledown, \pi')$ by letting $\mathbb{S}$ consist of the premises of the multicut, with the relation $\bigtriangledown$ given as specified by the multicut, and letting π' assign to each premise U the corresponding generated sub-proof of π. We define $T(S,i,\pi) = \{O(S',i')\}$, where (S',i') is the side formula occurrence of the multicut corresponding to (S,i).

Now suppose S has at least one premise belonging to π_{ep}. Then S is the conclusion of a logical rule or a structural rule other than cut. Let S_0, S_1 be the left and right premise of S, with S_0 denoting the unique premise if the rule used was unary. We define $T(S,i,\pi)$ by a case distinction on the rule used:

Weakening: Put $T(S,*,\pi) = \{\bot_A\}$ for the principal index, where the principal formula is A. For other indices we put $T(S,i,\pi) = T(S',i',\pi)$ where (S',i') is the formula occurrence associated with (S,i) via the weakening inference.

Contraction: Put $T(S,*,\pi) = T(S_0,*,\pi) \cup T(S_0,*+1,\pi)$ for the principal index. For a non-principal index put $T(S,i,\pi) = T(S_0,i',\pi)$ where i' is the corresponding index in the premise.
Exchange: Assume the exchanged formulas have indices i,j, put $T(S,i,\pi) = T_(S_0,j,\pi)$, $T(S,j,\pi) = T(S_0,i,\pi)$ and $T(S,k,\pi) = T(S_0,k,\pi)$ for $k \notin \{i,j\}$.
Right ∃-rule: For the principal index put $T(S,*,\pi) = \{\mathtt{e}(t \cdot \bot)v \mid v \in T(S_0,*,\pi))\}$ where t is the witness used in the rule application. For a non-principal index put $T(S,i,\pi) = T(S_0,i,\pi)$.
Right ∨-rule: For the principal index put $T(S,*,\pi) = \{\mathtt{i}_j v \mid v \in (T(S_0,*,\pi))\}$ where j is 0 or 1 depending on whether the minor formula is the left or right disjunct. For a non-principal index put $T(S,*,\pi) = T(S_0,*,\pi)$.
Left ∨-rule: For the principal index put

$$T(S,*.\pi) = \{\mathtt{j}vw \mid v \in T(S_0,*,\pi) \ \&\ w \in T(S_1,*,\pi)\}.$$

For a non-principal index put $T(S,i,\pi) = T(S_0,i,\pi) \cup T(S_1,i,\pi)$.

For a proof π of a Σ_1-sequent S of the form $\Gamma \Rightarrow A$ with a single formula on the right we abbreviate $T(S,i,\pi)$ by $T(\pi)$, where i the index of the right-most formula. Note that target term sets always consist of final terms.
Definition 18. The *guarded version* of a regular proof π is the proof π' obtained by replacing every bottom-most cut in π by a multicut with the same premises. We say that the proof π reduces to the proof π^* via multicut reductions if its guarded version does.
Proposition 18. *Let π be a proof with Σ_1 end sequent S of the form $\Gamma \Rightarrow A$, and let π' be its guarded version. Then:*

$$\mathsf{S}_\pi \sqsupseteq T(\pi')$$

Definition 19. A proof is said to be *essential-cut-free* if every cut formula in a cut or multicut appearing in the proof is quantifier free.
Note that any essential-cut-free proof is multicut guarded.

7.3 Multicut reductions

In this section we introduce reduction rules for multicuts, and show how these can be simulated using rewrite rules for Herbrand schemes. In the following we assume an extended multicut instance $(\mathbb{S}, \bigtriangledown, \pi)$ which reduces to an e-m-i $(\mathbb{S}^*, \bigtriangledown^*, \pi^*)$ by one of several possible reductions which fall in one of two categories: *external* reductions in which a the multicut is permuted with the rule at the root of one of the premise derivations, and *internal* reductions in which the derivation below the multicut is unchanged, modulo structural rules, and one or more of the premises are reduced in size.

7.3.1 External reductions

If the principal formula(s) of a premise of a multicut are not cut-formulas, then a multicut reduction is applicable which permutes the multicut with the final rule of that premise.

Axiom

If a premise to the multicut is an axiom and neither occurrences of the principal formula is a cut formula, then the multicut contains exactly one sequent and may be discharged as an axiom:

$$\dfrac{\dfrac{}{A \Rightarrow A}\,\mathsf{id}}{A \Rightarrow A}\,\triangledown \quad \Rightarrow \quad \dfrac{}{A \Rightarrow A}\,\mathsf{id}$$

Proposition 19. *If π reduces to π^* under the above reduction then $\pi^* \sqsubseteq \pi$.*

Proof. Immediate. □

Disjunction (right)

This external reduction permutes the multicut with an application of $\vee\mathsf{R}$:

$$\dfrac{\cdots \quad \dfrac{\begin{array}{c}\vdots\\ \Sigma \Rightarrow \Pi, A_j\end{array}}{\Sigma \Rightarrow \Pi, A_0 \vee A_1}\,\vee\mathsf{R}}{\Gamma \Rightarrow \Delta, A_0 \vee A_1}\,\triangledown \quad \Rightarrow \quad \dfrac{\dfrac{\cdots \quad \begin{array}{c}\vdots\\ \Sigma \Rightarrow \Pi, A_j\end{array}}{\Gamma \Rightarrow \Delta, A_j}\,\triangledown^*}{\Gamma \Rightarrow \Delta, A_0 \vee A_1}\,\vee\mathsf{R}$$

We first properly define the above reduction as transformation on e-m-is $(\mathbb{S}, \triangledown, \pi) \mapsto (\mathbb{S}^*, \triangledown^*, \pi^*)$. Suppose the permuted side formula occurrence is (P_0, p) where $P_0[p] = A_0 \vee A_1$ (and the formula occurrence is on the right) and the premise of P_0 in π_{P_0} is P_1 with $P_1[p] = A_j$. We define the permuted multicut instance $(\mathbb{S}^*, \triangledown^*, \pi^*)$ as follows:

- Replace P_0 by P_1, and keep all other sequents from $\mathbb{S}$.
- Define $\pi^*_{P_1}$ to be the subproof of π_{P_0} generated by P_1.
- Define $\triangledown^*$ to be the smallest symmetric relation satisfying:
  - If $(S, i) \triangledown (P_0, j)$ then $(S, i) \triangledown^* (P_1, j)$.
  - If $(S, i) \triangledown (U, j)$, $S \neq P_0$ and $U \neq P_0$, then $(S, i) \triangledown^* (U, j)$.

This definition corresponds to the reduction rule on a proof ending with a multicut:

$$\dfrac{(\mathbb{S}, \triangledown, \pi)}{\Gamma \Rightarrow \Delta, A_0 \vee A_1} \quad \Rightarrow \quad \dfrac{\dfrac{(\mathbb{S}^*, \triangledown^*, \pi^*)}{\Gamma \Rightarrow \Delta, A_j}}{\Gamma \Rightarrow \Delta, A_0 \vee A_1}\,\vee\mathsf{R}$$

where the displayed formula occurrence $A_0 \vee A_1$ corresponds to the formula occurrence (P_0, p) in the multicut.

Proposition 20. *For any formula occurrence (S, i), we have:*

1. *If $S = P_0$ and $i \neq p$ then $O(P_0, i) \sqsupseteq O^*(P_1, i)$.*
2. *$O(P_0, p) \sqsupseteq \mathtt{i}_j(O^*(P_1, p))$.*
3. *If $S \neq P_0$ then $O(S, i) \sqsupseteq O^*(S, i)$.*

Proof. Left to the reader. □

Disjunction (left)

As $\vee\mathsf{L}$ is a branching rule, the external reduction for this inference duplicates the multicut over the two premises:

$$\dfrac{\cdots \quad \dfrac{\dfrac{\vdots}{\Sigma, A_0 \Rightarrow \Pi} \quad \dfrac{\vdots}{\Sigma, A_1 \Rightarrow \Pi}}{\Sigma, A_0 \vee A_1 \Rightarrow \Pi}\vee\mathsf{L}}{\Gamma, A_0 \vee A_1 \Rightarrow \Delta}\triangledown \quad \Rightarrow \quad \dfrac{\dfrac{\cdots \quad \dfrac{\vdots}{\Sigma, A_0 \Rightarrow \Pi}}{\Gamma, A_0 \Rightarrow \Delta}\triangledown^* \quad \dfrac{\cdots \quad \dfrac{\vdots}{\Sigma, A_1 \Rightarrow \Pi}}{\Gamma, A_1 \Rightarrow \Delta}\triangledown^*}{\Gamma, A_0 \vee A_1 \Rightarrow \Delta}\vee\mathsf{L}$$

Suppose the permuted side formula occurrence is (P_0, p) where $P_0[p] = A_0 \vee A_1$ (and the formula occurrence is on the left) and the premises of P_0 in π_{P_0} is P_1 with $P_1[p] = A_0$ and P_2 with $P_2[p] = A_1$. We define the *left* permuted multicut instance $(\mathbb{S}^*_L, \triangledown^*_L, \pi^*_L)$ as follows:

- Replace P_0 by P_1, and keep all other sequents from $\mathbb{S}$.
- Define $(\pi^*_L)_{P_1}$ to be the subproof of π_{P_0} generated by P_1.
- Define $\triangledown^*_L$ to be the smallest symmetric relation satisfying:
  - If $(S, i) \triangledown (P_0, j)$ then $(S, i) \triangledown^*_L (P_1, j)$.
  - If $(S, i) \triangledown (U, j)$, $S \neq P_0$ and $U \neq P_0$, then $(S, i) \triangledown^*_L (U, j)$.

Similarly we define the *right* permuted multicut instance $(\mathbb{S}^*_R, \triangledown^*_R, \pi^*_R)$ as follows:

- Replace P_0 by P_2, and keep all other sequents from $\mathbb{S}$.
- Define $(\pi^*_R)_{P_2}$ to be the subproof of π_{P_0} generated by P_2.
- Define $\triangledown^*_R$ to be the smallest symmetric relation satisfying:
  - If $(S, i) \triangledown (P_0, j)$ then $(S, i) \triangledown^*_R (P_1, j)$.
  - If $(S, i) \triangledown (U, j)$, $S \neq P_0$ and $U \neq P_0$, then $(S, i) \triangledown^*_R (U, j)$.

This definition corresponds to a reduction rule for a proof ending with a multicut:

$$\dfrac{(\mathbb{S}, \triangledown, \pi)}{\Gamma, A_0 \vee A_1 \Rightarrow \Delta} \quad \Rightarrow \quad \dfrac{\dfrac{(\mathbb{S}^*_L, \triangledown^*_L, \pi^*_L)}{\Gamma, A_0 \Rightarrow \Delta} \quad \dfrac{(\mathbb{S}^*_R, \triangledown^*_R, \pi^*_R)}{\Gamma, A_1 \Rightarrow \Delta}}{\Gamma, A_0 \vee A_1 \Rightarrow \Delta}\vee\mathsf{L}$$

where the displayed formula occurrence $A_0 \vee A_1$ corresponds to the formula occurrence (P_0, p) in the multicut.

Proposition 21. *For any formula occurrence (S, i) of the multicut, we have:*

1. $O(P_0, p) \sqsupseteq \mathsf{j}^j(O^*_L(P_1, p), O^*_R(P_2, p))$.
2. *If $i \neq p$ then:* $O(P_0, i) \sqsupseteq \{O^*_L(P_1, i), O^*_R(P_2, i)\}$
3. *If $S \neq P_0$ and (S, i) is a side formula occurrence then* $O(S, i) \sqsupseteq \{O^*_L(S, i), O^*_R(S, i)\}$.

Here, the terms $O^*_L(S, i)$ and $O^*_R(S, i)$ refer to the output terms in the left and right reduced multicut respectively.

Proof. See Appendix A.4.1. □

Existential (right)

The case of $\exists\mathsf{R}$ is similar to $\vee\mathsf{R}$. The external reduction for this rule is:

$$\dfrac{\cdots \quad \dfrac{\overset{\vdots}{\Sigma \Rightarrow \Pi, A[t/x]}}{\Sigma \Rightarrow \Pi, \exists x A}\exists\mathsf{R}}{\Gamma \Rightarrow \Delta, \exists x A}\bigtriangledown \quad \Rightarrow \quad \dfrac{\dfrac{\cdots \quad \overset{\vdots}{\Sigma \Rightarrow \Pi, A[t/x]}}{\Gamma \Rightarrow \Delta, A[t/x]}\bigtriangledown^*}{\Gamma \Rightarrow \Delta, \exists x A}\exists\mathsf{R}$$

We first define the permutation of an extended multicut instance $(\mathbb{S}, \bigtriangledown, \pi)$. Suppose the permuted side formula occurrence is (P_0, p) where $P_0[p] = \exists x A$ (and the formula occurrence is on the right) and the premise of P_0 in π_{P_0} is P_1 with $P_1[p] = A[t/x]$. We define the permuted multicut instance $(\mathbb{S}^*, \bigtriangledown^*, \pi^*)$ as follows:

- Replace P_0 by P_1, and keep all other sequents from $\mathbb{S}$.
- Define $\pi^*_{P_1}$ to be the subproof of π_{P_0} generated by P_1.
- Define $\bigtriangledown^*$ to be the smallest symmetric relation satisfying:
  - If $(S, i) \bigtriangledown (P_0, j)$ then $(S, i) \bigtriangledown^* (P_1, j)$.
  - If $(S, i) \bigtriangledown (U, j)$, $S \neq P_0$ and $U \neq P_0$, then $(S, i) \bigtriangledown^* (U, j)$.

This definition corresponds to a reduction rule for a proof ending with a multicut:

$$\dfrac{(\mathbb{S}, \bigtriangledown, \pi)}{\Gamma \Rightarrow \exists x A, \Delta} \quad \Rightarrow \quad \dfrac{\dfrac{(\mathbb{S}^*, \bigtriangledown^*, \pi^*)}{\Gamma \Rightarrow A[t/x], \Delta}}{\Gamma \Rightarrow \exists x A, \Delta}\exists\mathsf{R}$$

where the displayed formula occurrence $\exists x A$ corresponds to the formula occurrence (P_0, p) in the multicut.

Proposition 22. *For any formula occurrence (S, i), we have:*

1. *If $S = P_0$ and $i \neq p$ then $O(P_0, i) \sqsupseteq O^*(P_1, i)$.*
2. *$O(P_0, p) \sqsupseteq \mathsf{e}(t \cdot \bot, O^*(P_1, p))$.*
3. *If $S \neq P_0$ then $O(S, i) \sqsupseteq O^*(S, i)$.*

Proof. See Appendix A.4.2. □

Existential (left)

If a premise of the multicut ends in $\exists\mathsf{L}$ then the principal formula is necessarily a cut formula because the conclusion is a Σ_1-sequent. Hence, there is no external reduction for the case of $\exists\mathsf{L}$.

Negation

The case of negation is straightforward and left to the reader.

Contraction

The final rule of a premise is contraction and the contracted formula is a side formula of the multicut. In the case of right-contraction, this induces the following multicut reduction:

$$\dfrac{\cdots \quad \dfrac{\dfrac{\vdots}{\Sigma \Rightarrow \Pi, A, A}}{\Sigma \Rightarrow \Pi, A}\,\mathsf{cR}}{\Gamma \Rightarrow \Delta, A}\,\triangledown \quad \Rightarrow \quad \dfrac{\dfrac{\cdots \quad \dfrac{\vdots}{\Sigma \Rightarrow \Pi, A, A}}{\Gamma \Rightarrow \Delta, A, A}\,\triangledown^*}{\Gamma \Rightarrow \Delta, A}\,\mathsf{cR}$$

The left contraction rule is handled in the same manner. Suppose the permuted side formula occurrence is (P_0, p) where $P_0[p] = A$ and the premise of P_0 in π_{P_0} is P_1 with $P_1[p] = P_1[p+1] = A$. We define the permuted multicut instance $(\mathbb{S}^*, \triangledown^*, \pi^*)$ as follows:

- Replace P_0 by P_1, and keep all other sequents from $\mathbb{S}$.
- Define $\pi^*_{P_1}$ to be the subproof of π_{P_0} generated by P_1.
- Define $\triangledown^*$ to be the smallest symmetric relation satisfying:
  - If $(S, i) \triangledown (P_0, j)$ and $j < p$ then $(S, i) \triangledown^* (P_1, j)$.
  - If $(S, i) \triangledown (P_0, j)$ and $j > p$ then $(S, i) \triangledown^* (P_1, j+1)$.
  - If $(S, i) \triangledown (U, j)$, $S \neq P_0$ and $U \neq P_0$, then $(S, i) \triangledown^* (U, j)$.

This definition corresponds to a reduction rule for a proof ending with a multicut:

$$\dfrac{(\mathbb{S}, \triangledown, \pi)}{\Gamma \Rightarrow \Delta, A} \quad \Rightarrow \quad \dfrac{\dfrac{(\mathbb{S}^*, \triangledown^*, \pi^*)}{\Gamma \Rightarrow \Delta, A, A}}{\Gamma \Rightarrow \Delta, A}\,\mathsf{cR}$$

Where the displayed formula occurrence A corresponds to the formula occurrence (P_0, p) in the multicut.

Proposition 23. *For any formula occurrence (S, i) we have:*

1. *If $S = P_0$ and $i < p$ then $O(P_0, j) \sqsupseteq O^*(P_1, j)$.*
2. *If $S = P_0$ and $i > p$ then $O(P_0, j) \sqsupseteq O^*(P_1, j+1)$.*
3. *$\mathsf{E}(O(P_0, p)) \sqsupseteq \{\mathsf{E}(O^*(P_1, p)), \mathsf{E}(O^*(P_1, p+1))\}$.*
4. *If $S \neq P_0$ then $O(S, i) \sqsupseteq O^*(S, i)$.*

Proof. See Appendix A.4.3. □

Weakening and exchange

The multicut reductions for the two remaining inferences are much simpler, so we will leave the precise definition of the reductions as easy exercises and merely sketch the reductions. In each case we can state and prove suitable propositions describing how the reductions can be simulated by rewrite rules for the associated Herbrand schemes. The arguments follow the same structure as the cases we've already seen, but are much simpler.

The final rule of a premise is weakening and the introduced formula is a side formula of the multicut. In the case of right-weakening, this induces the following multicut reduction:

$$\dfrac{\cdots \quad \dfrac{\dfrac{\vdots}{\Sigma \Rightarrow \Pi}}{\Sigma \Rightarrow \Pi, A}\,\mathsf{wR}}{\Gamma \Rightarrow \Delta, A}\,\triangledown \quad \Rightarrow \quad \dfrac{\dfrac{\cdots \quad \dfrac{\vdots}{\Sigma \Rightarrow \Pi}}{\Gamma \Rightarrow \Delta}\,\triangledown^*}{\Gamma \Rightarrow \Delta, A}\,\mathsf{wR}$$

The case of weakening of the left is handled in the same manner.

Exchange is handled in the same way. For the right exchange rule, where both exchanged formula occurrences are side formulas, the multicut reduction takes the following form:

$$\dfrac{\cdots \quad \dfrac{\dfrac{\vdots}{\Sigma \Rightarrow \Pi, B, A}}{\Sigma \Rightarrow \Pi, A, B}\,\mathsf{eR}}{\Gamma \Rightarrow \Delta, A, B}\,\triangledown \quad \Rightarrow \quad \dfrac{\cdots \quad \dfrac{\vdots}{\Sigma \Rightarrow \Pi, B, A}}{\Gamma \Rightarrow \Delta, A, B}\,\triangledown$$

Note that there is no need to introduce an instance of the exchange rule below the multicut, as we can simply re-arrange the map f connecting side formulas of the conclusion to side formulas in the multicut to make the multicut application in the reduced proof valid. The cases where one or both exchanged formulas are cut formulas is similar, but there we also need to suitable re-arrange the cut relation.

7.3.2 Internal reductions

In the external reductions, a rule of inference from one of the premise derivations was permuted with the multicut. This occurred because the principal formula(s) of the rule were not cut formulas. The internal reductions, which we now present, treat the case that the principal formula of the root inference of one or more premises is a cut formula. The result of the reduction is that the rule in the premise(s) is removed without inducing any change to the generated cut-free proof.

Axiom

If a premise of the multicut is an axiom and both principal formulas are cut formulas, then the multicut can be transformed by the following reduction:

$$\dfrac{\cdots \quad \dfrac{}{A \Rightarrow A}\,\mathsf{id} \quad \dfrac{\vdots}{A, \Sigma \Rightarrow \Pi} \quad \cdots}{\Gamma \Rightarrow \Delta}\,\triangledown \quad \Rightarrow \quad \dfrac{\cdots \quad \dfrac{\vdots}{A, \Sigma \Rightarrow \Pi} \quad \cdots}{\Gamma \Rightarrow \Delta}\,\triangledown^*$$

Suppose the extended multicut instance to be reduced is $(\mathbb{S}, \triangledown, \pi)$, and suppose $R \in \mathbb{S}$ is an axiom, where $(R, 0)$ and $(R, 1)$ are the two occurrences of the principal formula of the axiom. Suppose at least one of these two occurrences is a cut formula occurrence, say 1, and let Q be the unique sequent such that $(Q, q) \triangledown (R, 1)$ for some $q < l(Q)$. We define the reduced extended multicut instance $(\mathbb{S}^*, \triangledown^*, \pi^*)$ as follows:

- Remove R from $\mathbb{S}$ to form $\mathbb{S}^*$.
- Let π^* be the restriction of π to $\mathbb{S}^*$.
- Define $\triangledown^*$ to be the smallest symmetric relation such that:

– If $(S,i) \bigtriangledown (U,j)$ and $\{S,U\} \cap \{R,Q\} = \emptyset$, then $(S,i) \bigtriangledown^* (U,j)$.
– If $(S,i) \bigtriangledown (Q,j)$ where $S \neq R$ then $(S,i) \bigtriangledown^* (Q,j)$.
– If $(S,i) \bigtriangledown (R,0)$ where $S \neq Q$ then $(S,i) \bigtriangledown^* (Q,q)$.

Proposition 24. *Let (S,i) be any side formula occurrence with $S \in \mathbb{S}$.*

1. If $S = Q$ then $O(Q,i) \sqsupseteq O^*(Q,i)$.
2. If $S \notin \{Q,R\}$ then $O(S,i) \sqsupseteq O^*(S,i)$.

Proof. Left to the reader. □

Weakening

If a cut formula has been introduced by weakening, the associated multicut reduction simply removes the application of weakening and all parts of the multicut that depended on the weakened cut formula. In the case of wL, this reduction can be visualised as:

$$\dfrac{\cdots \quad \dfrac{\vdots}{\Sigma \Rightarrow \Pi, A} \quad \dfrac{\dfrac{\vdots}{\Sigma' \Rightarrow \Pi'}}{A, \Sigma' \Rightarrow \Pi'}\,\mathsf{wL}}{\Gamma \Rightarrow \Delta}\,\bigtriangledown \quad \Rightarrow \quad \dfrac{\dfrac{\cdots \quad \dfrac{\vdots}{\Sigma' \Rightarrow \Pi'}}{\Gamma' \Rightarrow \Delta'}\,\bigtriangledown^*}{\Gamma \Rightarrow \Delta}\,(\mathsf{wL}+\mathsf{wR})^*$$

where $\Gamma' \Rightarrow \Delta'$ is the subsequent of $\Gamma \Rightarrow \Delta$ obtained by removing the side formula occurrences emanating from the premise $\Sigma \Rightarrow \Pi, A$ and other premises connected to the cut formula A.

More precisely, given an extended multicut instance $(\mathbb{S}, \bigtriangledown, \pi)$ in which the sequent P is the conclusion of a (left or right) weakening on a cut formula occurrence (P,p). Say that a premise S is *affected* if $S \neq P$ and the unique $\bigtriangledown$-path from S to P ends with the cut on (P,p), and *unaffected* otherwise. We define the reduced multicut $(\mathbb{S}^*, \bigtriangledown^*, \pi^*)$ as follows:

- Let $\mathbb{S}^*$ consist of the unaffected premises of $\mathbb{S}$ together with the premise P' of P in π_P.
- Let π^* assign the generated subproof of π_P to P', and π_S to each unaffected premise S.
- Let $\bigtriangledown^*$ be the smallest symmetric relation such that:
  – If U,V are both unaffected and $(U,i) \bigtriangledown (V,j)$ then $(U,i) \bigtriangledown^* (V,j)$.
  – If U is unaffected and $(U,i) \bigtriangledown (P,j)$ then $(U,i) \bigtriangledown^* (P',j')$, where (P',j') is the formula occurrence corresponding to (P,j) via the weakening inference.

This definition corresponds to a cut reduction:

$$\dfrac{(\mathbb{S}, \bigtriangledown, \pi)}{\Gamma \Rightarrow \Delta} \quad \Rightarrow \quad \dfrac{\dfrac{(\mathbb{S}^*, \bigtriangledown^*, \pi^*)}{\Gamma' \Rightarrow \Delta'}}{\Gamma \Rightarrow \Delta}\,(\mathsf{wL}+\mathsf{wR})^*$$

where Γ', Δ' are those side formula occurrences belonging to unaffected premises of the original multicut, together with the side formulas occurring in the conclusion of the weakening inference.

Proposition 25. *Let* $(\mathbb{S}, \triangledown, \pi)$ *be a multicut as above and let* (S, i) *be a side formula occurrence in one of the premises.*

1. *If* S *is unaffected then* $O(S, i) \sqsupseteq O^*(S, i)$*.*
2. *If* $S = P$ *then* $O(P, i) \sqsupseteq O^*(P', i')$*, where* i' *is the index associated with* i *via the weakening inference.*
3. *If* S *is affected then* $O(S, i) \sqsupseteq \bot_{S[i]}$*.*

Proof. See Appendix A.5.1. □

Contraction

If a cut formula has been contracted, the associated multicut reduction removes the application of contraction and duplicates all necessary premises of the multicut. In the case of cR, this reduction can be visualised as:

$$\dfrac{\cdots \quad \dfrac{\dfrac{\vdots}{\Sigma \Rightarrow \Pi, A, A}}{\Sigma \Rightarrow \Pi, A}\,\mathsf{cR} \qquad \dfrac{\vdots}{A, \Sigma' \Rightarrow \Pi'}}{\Gamma, \Gamma' \Rightarrow \Delta, \Delta'}\,\triangledown \quad \Rightarrow \quad \dfrac{\dfrac{\cdots \quad \dfrac{\vdots}{\Sigma \Rightarrow \Pi, A, A} \quad \dfrac{\vdots}{A, \Sigma' \Rightarrow \Pi'} \quad \dfrac{\vdots}{A, \Sigma' \Rightarrow \Pi'}}{\Gamma, \Gamma', \Gamma' \Rightarrow \Delta, \Delta', \Delta'}\,\triangledown^*}{\Gamma, \Delta' \Rightarrow \Delta, \Delta'}\,\mathsf{c}^*$$

Suppose the extended multicut instance to be reduced is $(\mathbb{S}, \triangledown, \pi)$, and suppose $R_0 \in \mathbb{S}$ is the conclusion of a contraction in π_R with with premise R_1, where (R_0, p) is the principal formula occurrence of the contraction. Let Q be the unique sequent such that $(Q, q) \triangledown (R_0, p)$ for some $q < l(Q)$. Say that a sequent $S \in \mathbb{S}$, $S \neq R_0$, is *affected* by the reduction if the unique path from S to R_0 in $G(\mathbb{S}, \triangledown)$ contains Q, and *unaffected* otherwise. We define the reduced extended multicut instance $(\mathbb{S}^*, \triangledown^*, \pi^*)$ as follows:

- We let the premise R_1 replace R_0. Set $\pi^*_{R_1}$ to be the subproof of π_{R_0} generated by R_1.
- For each affected $S \in \mathbb{S}$ we include two sequents $S_L, S_R \in \mathbb{S}^*$ with $l(S_L) = l(S_R) = l(S)$ and $S_L[i] = S_R[i] = S[i]$ for each $i < l(S)$. We put $\pi^*_{S_L} = \pi_S$ and let $\pi^*_{S_R}$ a copy of π_S with suitably renamed eigenvariables.
- For each unaffected $S \in \mathbb{S}$, we put $S \in \mathbb{S}^*$ and $\pi^*_S = \pi_S$.
- Define $\triangledown^*$ to be the smallest symmetric relation for which the following conditions hold:
  - We have $(Q_L, q) \triangledown^* (R_1, p)$ and $(Q_R, q) \triangledown^* (R_1, p+1)$.
  - If $(S, i) \triangledown (R_0, j)$ for $j \neq p$ then $(S, i) \triangledown^* (R_1, j)$ if $j < p$ and $(S, i) \triangledown^* (R_1, j+1)$ if $j > p$. (Note we must have $S \neq Q$ and that S is unaffected.)
  - If $(S, i) \triangledown (U, j)$ and both S and U are unaffected then $(S, i) \triangledown^* (U, j)$.
  - If $(S, i) \triangledown (U, j)$ and both S and U are affected then $(S_L, i) \triangledown^* (U_L, j)$ and $(S_R, i) \triangledown^* (U_R, j)$.

This definition corresponds to a reduction rule for a proof ending with a multicut:

$$\frac{(\mathbb{S}, \triangledown, \pi)}{\Gamma, \Gamma' \Rightarrow \Delta, \Delta'} \quad \Rightarrow \quad \dfrac{\dfrac{(\mathbb{S}^*, \triangledown^*, \pi^*)}{\Gamma, \Gamma', \Gamma' \Rightarrow \Delta, \Delta', \Delta'}}{\Gamma, \Gamma' \Rightarrow \Delta, \Delta'}\,\mathsf{c}^*$$

where Γ', Δ' are the formula occurrences of the conclusion corresponding to formula occurrences in affected premises of the multicut. Given $S \in \mathbb{S}$ we write $O(S, i)$ and $I(S, i)$ for the canonical output and input terms for (S, i) with respect to $(\mathbb{S}, \triangledown, \pi)$, and call these the *prior* canonical output/input terms. Similarly for $(S, i) \in \mathbb{S}^*$ we write $O^*(S, i)$ and $I^*(S, i)$ for the canonical output/input terms for (S, i) with respect to $(\mathbb{S}^*, \triangledown^*, \pi^*)$ and call these the *posterior* canonical output/input terms.

Proposition 26. *Let $(\mathbb{S}, \triangledown, \pi)$ and $(\mathbb{S}^*, \triangledown^*, \pi^*)$ be as above, let (S, i) be a side formula occurrence with $S \in \mathbb{S}$.*

- *If $S \neq R_0$ is affected by the reduction then*

$$\mathsf{E}(O(S, i)) \sqsupseteq \{E(O^*(S_L, i)), E(O^*(S_R, i))\}$$

- *If $S = R_0$ we have $O(R_0, i) \sqsupseteq O^*(R_1, i)$ if $i < p$, and $O(R_0, i) \sqsupseteq O^*(R_1, i+1)$ if $i > p$.*
- *In all other cases we have $\mathsf{E}(O(S, i)) \sqsupseteq \mathsf{E}(O^*(S, i))$.*

Proof. See Appendix A.5.2. □

Existential

An existentially quantified cut-formula induces the following multicut reduction:

$$\dfrac{\cdots \quad \dfrac{\dfrac{\vdots}{\Sigma \Rightarrow \Pi, A[t/x]}}{\Sigma \Rightarrow \Pi, \exists x A}\,\exists\mathsf{R} \quad \dfrac{\dfrac{\vdots}{A[\alpha/x], \Sigma' \Rightarrow \Pi'}}{\exists x A, \Sigma' \Rightarrow \Pi'}\,\exists\mathsf{L}}{\Gamma \Rightarrow \Delta}\,\triangledown \quad \Rightarrow \quad \dfrac{\cdots \quad \dfrac{\vdots}{\Sigma \Rightarrow \Pi, A[t/x]} \quad \dfrac{\vdots\,[t/x]}{A[t/x], \Sigma' \Rightarrow \Pi'}}{\Gamma \Rightarrow \Delta}\,\triangledown^*$$

We first define the reduction. Suppose the extended multicut instance to be reduced is $(\mathbb{S}, \triangledown, \pi)$, and suppose the principal formula occurrences are (L_0, l) and (R_0, r), where (L_0, l) is a left formula occurrence, (R_0, r) is a right formula occurrence, and $L[l] = R[r] = \exists x A$. Suppose the premise of R_0 is R_1 with $R_1[r] = A[t/x]$ and suppose the premise of L_0 is L_1 with $L_1[l] = A[\alpha/x]$. We define the reduced extended multicut instance $(\mathbb{S}^*, \triangledown^*, \pi^*)$ as follows:

- Replace R_0 by R_1 and L_0 by L_1. We set $\pi^*_{R_1}$ to be the subproof of π_{R_0} generated by R_1, and $\pi^*_{L_1}$ is obtained by uniformly substituting t for α in the subproof of L_0 generated by L_1.
- All other sequences in $\mathbb{S}$ are carried over as they are to $\mathbb{S}^*$.
- The relation $\triangledown^*$ is defined as the smallest symmetric relation satisfying the following conditions:
  - $(L_1, l) \triangledown (R_1, r)$.

- $(S,i) \triangledown^* (U,j)$ for all formula occurrences $(S,i) \triangledown (U,j)$ with $S,U \notin \{L_0, R_0\}$.
- If $(S,i) \triangledown (L_0,j)$ and $S \neq R_0$ then $(S,i) \triangledown^* (L_1,j)$.
- If $(S,i) \triangledown (R_0,j)$ and $S \neq L_0$ then $(S,i) \triangledown^* (R_1,j)$.

This definition corresponds to a reduction rule for a proof ending with a multicut:

$$\frac{(\mathbb{S}, \triangledown, \pi)}{\Gamma \Rightarrow \Delta} \quad \Rightarrow \quad \frac{(\mathbb{S}^*, \triangledown^*, \pi^*)}{\Gamma \Rightarrow \Delta}$$

Proposition 27. *Let (S,i) be any side formula occurrence with $S \in \mathbb{S}$.*

1. *If $S = L_0$ then $O(L_0,i) \sqsupseteq O^*(L_1,i)$.*
2. *If $S = R_0$ then $O(R_0,i) \sqsupseteq O^* R_1, i)$.*
3. *If $S \notin \{L_0, R_0\}$ then $O(S,i) \sqsupseteq O^*(S,i)$.*

Proof. See Appendix A.5.3. □

Disjunction

If a disjunction is cut-formula and principal in both it's occurring premises, the following reduction is applicable:

$$\dfrac{\cdots \quad \dfrac{\dfrac{\vdots}{\Sigma \Rightarrow \Pi, A}}{\Sigma \Rightarrow \Pi, A \vee B} \vee\mathsf{R} \quad \dfrac{\dfrac{\vdots}{A, \Sigma' \Rightarrow \Pi'} \quad \dfrac{\vdots}{B, \Sigma' \Rightarrow \Pi'}}{A \vee B, \Sigma' \Rightarrow \Pi'} \vee\mathsf{L}}{\Gamma \Rightarrow \Delta} \triangledown \quad \Rightarrow \quad \dfrac{\cdots \quad \dfrac{\vdots}{\Sigma \Rightarrow \Pi, A} \quad \dfrac{\vdots}{A, \Sigma' \Rightarrow \Pi'}}{\Gamma \Rightarrow \Delta} \triangledown^*$$

We first define the reduction of an extended multicut instance $(\mathbb{S}, \triangledown, \pi)$. Suppose (L_0, l) is the reduced principal cut formula occurrence in which l is a left index, where $L_0[l] = A \vee B$, and suppose the premises of L_0 in π_{L_0} are L_1 and L_2 with $L_1[l] = A$ and $L_2[l] = B$. Suppose $(R_0, r) \triangledown (L_0, l)$ is the reduced principal cut formula occurrence in which r is a right index, and suppose the premise of R_0 in π_{R_0} is R_1 with $R_1[r] = A$ (the other case is similar). We define the reduced extended multicut instance $(\mathbb{S}^*, \triangledown^*, \pi^*)$ as follows:

- Replace L_0 by L_1 and R_0 by R_1, and keep all other sequents from $\mathbb{S}$. Set $\pi^*_{L_1}$ to be the subproof of π_{L_0} generated by L_1 and set $\pi^*_{R_1}$ to be the subproof generated by π_{R_0}.
- Define $\triangledown^*$ to be the smallest symmetric relation satisfying:
  - $(L_1, l) \triangledown^* (R_1, r)$.
  - If $(S,i) \triangledown (L_0,j)$ and $S \notin \{L_0, R_0\}$ then $(S,i) \triangledown^* (L_1,j)$.
  - If $(S,i) \triangledown (R_0,j)$ and $S \notin \{L_0, R_0\}$ then $(S,i) \triangledown^* (R_1,j)$.
  - If $(S,i) \triangledown (U,k)$ and $S,U \notin \{L_0, R_0\}$ then $(S,i) \triangledown^* (U,k)$.

This definition corresponds to the following reduction rule for a proof ending with a multicut:

$$\frac{(\mathbb{S}, \triangledown, \pi)}{\Gamma \Rightarrow \Delta} \quad \Rightarrow \quad \frac{(\mathbb{S}^*, \triangledown^*, \pi^*)}{\Gamma \Rightarrow \Delta}$$

Proposition 28. *For each side formula occurrence (S,i), we have:*

1. *If $S = L_0$ then $O(L_0, i) \sqsupseteq O^*(L_1, i)$.*
2. *If $S = R_0$ then $O(R_0, i) \sqsupseteq O^*(R_1, i)$.*
3. *If $S \notin \{L_0, R_0\}$ then $O(S, i) \sqsupseteq O^*(S, i)$.*

Proof. See Appendix A.5.4. □

Negation

The case of negation induces the following multicut reduction.

$$\dfrac{\cdots \quad \dfrac{\dfrac{\vdots}{\Sigma, A \Rightarrow \Pi}}{\Sigma \Rightarrow \Pi, \neg A}\neg\mathsf{R} \quad \dfrac{\dfrac{\vdots}{\Sigma' \Rightarrow \Pi', A}}{\neg A, \Sigma' \Rightarrow \Pi'}\neg\mathsf{L}}{\Gamma \Rightarrow \Delta}\bigtriangledown \quad \Rightarrow \quad \dfrac{\cdots \quad \dfrac{\vdots}{\Sigma' \Rightarrow \Pi', A} \quad \dfrac{\vdots}{\Sigma, A \Rightarrow \Pi}}{\Gamma \Rightarrow \Delta}\bigtriangledown^*$$

We leave the formal presentation and verification to the reader.

Cut

The case that a premise of the multicut ends with application of (binary) cut induces a reduction which absorbs the cut.

$$\dfrac{\cdots \quad \dfrac{\dfrac{\vdots}{\Sigma_0 \Rightarrow \Pi_0, C} \quad \dfrac{\vdots}{C, \Sigma_1 \Rightarrow \Pi_1}}{\Sigma_0\Sigma_1 \Rightarrow \Pi_0\Pi_1}\mathsf{cut}}{\Gamma \Rightarrow \Delta}\bigtriangledown \quad \Rightarrow \quad \dfrac{\cdots \quad \dfrac{\vdots}{\Sigma_0 \Rightarrow \Pi_1, C} \quad \dfrac{\vdots}{C, \Sigma_1 \Rightarrow \Pi_1}}{\Gamma \Rightarrow \Delta}\bigtriangledown^*$$

Suppose $(\mathbb{S}, \bigtriangledown, \pi)$ is an extended multicut instance, and P_0 is a sequent which is the conclusion of a binary cut rule application in π_{P_0}, with premises P_1, P_2, where (P_1, p_1) is the right-hand occurrence of the cut formula and (P_2, p_2) is the left-hand occurrence of the cut formula. Let f denote the map sending each given formula occurrence (P_0, i) to the corresponding formula occurrence in either the premise P_1 or P_2.

We define the reduced extended multicut instance $(\mathbb{S}^*, \bigtriangledown^*, \pi^*)$ as follows:

- Replace P_0 by P_1 and P_2, keeping all other sequents in $\mathbb{S}$.
- Define $\pi^*_{P_1}$ and $\pi^*_{P_2}$ to be the sub-proofs of π_{P_0} generated by P_1, P_2 respectively.
- Define $\bigtriangledown^*$ to be the smallest symmetric relation such that:
  - $(P_1, p_1) \bigtriangledown (P_2, p_2)$,
  - If $(S, i) \bigtriangledown (U, j)$ and $S \neq P_0$, $U \neq P_0$ then $(S, i) \bigtriangledown^* (U, j)$.
  - If $(S, i) \bigtriangledown (P_0, j)$, then $(S, i) \bigtriangledown^* f(P_0, j)$.

Proposition 29. *1. For all $i < l(P_0)$, $O(P_0, i) \sqsupseteq O^*(f(P_0, i))$.*
2. For every formula occurrence (S, i) with $S \neq P_0$, we have $O(S, i) \sqsupseteq O^*(S, i)$.

Proof. Left to the reader. □

8 Extracting Herbrand expansions

The main result of this article relates Herbrand schemes to cut elimination. Specifically, we prove the language of a Herbrand scheme covers all Herbrand expansions that can be extracted from a proof of a Σ_1-sequent by the cut reductions rules of the previous section:

Theorem 2. *Let π be a regular proof of a Σ_1-sequent of the form $\Gamma \Rightarrow A$ and suppose π reduces to an essential-cut-free proof π' (via multicut reduction). Then $S_\pi \longrightarrow t$ for some term $t \equiv T(\pi')$.*

Theorem 2 is an immediate consequence of proposition 18 and

Proposition 30. *Suppose π is a multicut guarded regular proof of a Σ_1-sequent. Suppose π^* is obtained from π by applying any of the multicut reduction or permutation rules. Then π^* is also a multicut guarded regular proof. Furthermore, for any sequent S appearing in π_{ep}, S is in π^*_{ep} also, and for all $i < |S|$ we have $T(S, i, \pi) \sqsupseteq T(S, i, \pi^*)$.*

Proof. It can be seen by inspection of the multicut reduction and permutation rules that they all preserve the property of being a regular multicut guarded proof. We omit the details.

The second item of the proposition is proved by induction on the difference $h(\pi_{ep}) - h(\pi_{ep}|S)$. The only non-trivial part of the induction is for the base case, where S is a leaf in π_{ep} and is therefore either an axiom or the conclusion of a multicut. In the former case, $T(S, i, \pi) = T(S, i, \pi^*)$ for each $i < l(S)$. In the latter case, we have an extended multicut instance $(\mathbb{S}, \bigtriangledown, \pi)$ where $\mathbb{S}$ consists of the premises of S, and we overload the notation to let π denote the assignment to each premise of the multicut the corresponding generated subproof of the proof π. We make a case distinction on what sort of cut reduction was applied to the multicut to obtain π^*, ignoring the trivial case in which some other multicut was reduced.

Internal existential, disjunction or negation reduction

Suppose the side formula occurrence corresponding to (S, i) in the multicut is (S', i'). Then Proposition 27 and Proposition 28 immediately give us:

$$T(S, i, \pi) = \{O(S', i')\} \sqsupseteq \{O^*(S', i')\} = T(S, i, \pi^*)$$

as required.

External existential reduction

As the end-sequent is Σ_1, this case only concerns the $\exists$R-rule. Let the permuted premise of the multicut be U_0, its premise in π be U_1 and the principal formula of the $\exists$-rule is $U_0[p] = \exists x A$, minor formula $U_1[p] = A[t/x]$, then let (S, i) be the formula occurrence corresponding to the formula occurrence (U_0, p) in the multicut, and let S' be the new premise of S in π^* Then we have:

$$\begin{aligned} T(S,i) &= \{\mathsf{E}(O(U_0,p))\} \\ &\sqsupseteq \{\mathsf{E}(\mathsf{e}(t\cdot\bot, O^*(U_1,p)))\} && \text{Proposition 22} \\ &\sqsupseteq \{\mathsf{e}(t\cdot\bot, \mathsf{E}(O^*(U_1,p))\} \\ &= \{\mathsf{e}(t\cdot\bot, v) \mid v \in T(S',i,\pi^*))\} && T(S',i,\pi^*) = \{\mathsf{E}(O^*(U_1,p))\} \\ &= T(S,i,\pi^*) \end{aligned}$$

as required. For $j \neq i$, if the formula occurrence corresponding to (S,j) in the multicut is (U_0,k) then $k \neq p$ and we get:

$$\begin{aligned} T(S,j) &= \{\mathsf{E}(O(U_0,k))\} \\ &\sqsupseteq \{\mathsf{E}(O^*(U_1,k))\} && \text{Proposition 22} \\ &= T(S',j,\pi^*) \\ &= T(S,j,\pi^*) \end{aligned}$$

If the formula occurrence corresponding to (S,j) in the multicut is (V,k) for $V \neq U_0$ then we have:

$$\begin{aligned} T(S,j) &= \{\mathsf{E}(O(V,k))\} \\ &\sqsupseteq \{\mathsf{E}(O^*(V,k))\} && \text{Proposition 22} \\ &= T(S',j,\pi^*) \\ &= T(S,j,\pi^*) \end{aligned}$$

as required.

Internal weakening reduction

By inspection of the reduction rule for weakening of a cut formula, it easy to verify that for the conclusion S of the multicut in π, if π reduces to π^* by such a reduction then $T(S,i,\pi^*) = \{\bot_A\}$ if (S,i) is a formula occurrence associated with an occurrence in an affected premise (where $S[i] = A$), and otherwise $T(S,i,\pi^*) = \{O^*(S',i')\}$ where (S',i') is the formula occurrence associated with (S,i) via the reduced multicut. In both cases we get $T(S,i,\pi) = \{O(S'',i'')\} \sqsupseteq T(S,i,\pi^*)$ by Proposition 25, where (S'',i'') is the premise associated with (S,i) via the multicut inference in π.

Internal contraction reduction

Let (R_0,p) be the principal formula of the reduced contraction and R_1 the premise of R_0 in π, let S' be the new conclusion of the reduced multicut. For each index $i < l(S)$, let i' be the index of the corresponding formula occurrence in S' if (S,i) is not copied by a contraction leading up to S', and let i'_0, i'_1 be the two indices corresponding to the two copies of (S,i) otherwise. Then, if (S,i) corresponds to an unaffected side formula occurrence (U, j) of the multicut with $U \neq R_0$ we get:

$$\begin{aligned} T(S,i,\pi) &= \{\mathsf{E}(O(U,j))\} \\ &\sqsupseteq \{\mathsf{E}(O^*(U,j))\} \quad \text{Proposition 26} \\ &= T(S',i',\pi^*) \\ &= T(S,i,\pi^*) \end{aligned}$$

If (S, i) corresponds to a side formula occurrence (R_0, j) with $j > p$ then:

$$\begin{aligned} T(S,i,\pi) &= \{\mathsf{E}(O(R_0,j))\} \\ &\sqsupseteq \{\mathsf{E}(O^*(R_1,j+1))\} \quad \text{Proposition 26} \\ &= T(S',i',\pi^*) \\ &= T(S,i,\pi^*) \end{aligned}$$

The case with $j < p$ is similar. Finally if (S, i) corresponds to an affected side formula occurrence (U, j) then:

$$\begin{aligned} T(S,i,\pi) &= \{\mathsf{E}(O(U,j))\} \\ &\sqsupseteq \{\mathsf{E}(O^*(U_L,j)), \mathsf{E}(O^*(U_R,j))\} \quad \text{Proposition 26} \\ &= T(S',i'_0,\pi^*) \cup T(S',i'_1,\pi^*) \\ &= T(S,i,\pi^*) \end{aligned}$$

as required.

External disjunction reduction

Suppose π^* was obtained from π by a left $\vee$-permutation. Let (P, p) be the principal formula occurrence in π, $P[p] = A \vee B$, where P is a premise of the multicut, let P_L, P_R be the premises of P and let (S, i) be the formula occurrence corresponding to (P, p) in the conclusion so $S[i] = A \vee B$ also. Let S'_L, S'_R be the new premises of S in π^*_{ep}. We have:

$$\begin{aligned} T(S,i,\pi) &= \{\mathsf{E}(O(P,p))\} \\ &\sqsupseteq \{\mathsf{E}(\mathsf{j}(O^*(P_L,p), O^*(P_R,p)))\} \qquad \text{Proposition 21} \\ &\sqsupseteq \{\mathsf{j}(\mathsf{E}(O^*(P_L,p)), \mathsf{E}(O^*(P_R,p)))\} \\ &= \{\mathsf{j}(T(S'_L,i,\pi^*), T(S'_R,i,\pi^*))\} \\ &= T(S',i,\pi^*) \end{aligned}$$

For $j \neq i$, let (U, k) be the formula occurrence corresponding to (S, j) in the multicut. If $U = P$ then we have:

$$\begin{aligned} T(S, j, \pi) &= \{\mathsf{E}(O(P, k))\} \\ &\sqsupseteq \{\mathsf{E}(O^*_L(P_L, k)), \mathsf{E}(O^*_R(P_R, k))\} \quad \text{Proposition 21} \\ &= T(S'_L, j, \pi^*) \cup T(S'_R, j, \pi^*) \\ &= T(S, j, \pi^*) \end{aligned}$$

If $U \neq P$ then:

$$\begin{aligned} T(S, j, \pi) &= \{\mathsf{E}(O(U, k))\} \\ &\sqsupseteq \mathsf{E}(O^*_L(U, k)) \parallel \mathsf{E}(O^*_R(U, k)) \quad \text{Proposition 21} \\ &= T(S'_L, j, \pi^*) \cup T(S'_R, j, \pi^*) \\ &= T(S, j, \pi^*) \end{aligned}$$

as required.

The remaining cases are trivial. □

To obtain Herbrand's theorem, it remains to establish weak normalisation of the multicut reductions.

Theorem 3. *Let π be a proof of a Σ_1-end-sequent. Then π reduces to an essential cut-free proof.*

We show that there is a terminating multicut reduction strategy by reducing to the case of propositional linear logic and appealing to a result on limits of fair multicut strategies in that context by Saurin [15]:

Theorem ([15, Theorem 44])**.** *For a proof in classical propositional logic, every fair multicut reduction sequence terminates.*

Saurin's cut-elimination argument interprets a propositional proof as a non-wellfounded proof in linear logic and shows that such proofs are closed under fair multicut elimination. The resultant proof will be well-founded (i.e., finite) by the assumption on the vocabulary of the end-sequent (i.e., that it is sequent in propositional logic).

Proof of Theorem 3. Let $\pi \vdash \Gamma \Rightarrow \Delta$ be a proof of a Σ_1 sequent. Stripping π of all first-order content we envisage π as a proof π_P in classical propositional logic of a sequent $\Gamma_P \Rightarrow \Delta_P$ via a translation that removes all terms for formulas:

$$(P\vec{t})_P = P \quad (A \vee B)_P = A_P \vee B_P \quad (\exists x A)_P = A_P \vee \bot \quad (\neg A)_P = \neg A_P$$

That is, π_P is the proof in classical propositional logic in which every formula occurrence A in π has been replaced by A_P and instances of $\exists\mathsf{L}$ and $\exists\mathsf{R}$ are changed appropriately. The special interpretation of $\exists x A$ is to ensure that applications of $\exists\mathsf{L}$ and $\exists\mathsf{R}$ are recorded in π_P with a corresponding principal formula.

It is clear that every multicut reduction $\pi \Rightarrow \pi^*$ corresponds to a multicut reduction $\pi_P \Rightarrow^* \pi^*_P$ where $\pi^*_P = (\pi^*)_P$ is the result of applying the above translation to π^*.

Indeed, for most reductions, we have $\pi_P \Rightarrow \pi_P^*$, but the special treatment of quantifiers induces two multicut reduction steps.

All that remains is to observe that if $\pi \Rightarrow \pi_0 \Rightarrow \cdots$ is a fair multicut reduction sequence, then the simulating reduction $\pi_P \Rightarrow^* (\pi_0)_P \Rightarrow^* \cdots$ can be chosen to be fair. Saurin's theorem implies that the reduction sequence necessarily terminates in a cut-free proof. □

Corollary (Soundness)**.** *Let π be a proof of a Σ_1-sequent $\Gamma \Rightarrow A$. Then $\mathscr{L}(\pi)$ is a valid Herbrand expansion of this sequent.*

Proof. From Theorems 1, 2 and 3. □

9 Conclusion

We have associated a higher-order recursion scheme to any sequent calculus proof for classical first-order logic that can compute Herbrand expansions for Σ_1-formulas. The framework builds on [1] offering several conceptual advantages. Primarily, cut is given a symmetric interpretation in contrast the to 'nested' composition in [1]. The introduction of explicit constructors and destructors for the logical connectives, and the 'call-with-current-continuation'-like Peirce operator for handling contraction, lifts the Herbrand scheme representation to arbitrary sequent calculus proofs. From a game-theoretic perspective, Herbrand schemes are extracting strategies that evidence formulas in the end sequent. The relationship between Herbrand schemes and the explicitly game-theoretic proof semantics in [12] seems to merit further investigation.

The main merit of the model is perhaps in giving a refined representation of Herbrand schemes amenable for treating extensions of FOL, in particular, with Martin-Löf style inductive definitions. With the restriction of prenex formulas lifted, one can now investigate the computational content ascribed by Herbrand schemes to inductive proofs in a classical setting. The latter is the main motivation behind the present work and is currently under investigation. The game theoretic interpretation also extends naturally to this setting where terms/programs represent strategies in an infinite game.

Acknowledgments

This work was supported by the Knut and Alice Wallenberg Foundation [2020.0199], Swedish Research Council [2020-01873, 2017-05111], and Dutch Research Council [OCENW.M20.048].

Appendix A Omitted proofs

A.1 Language equivalence is preserved through substitution

We restate proposition 15:
Proposition. *Let π be any regular proof and σ a regular substitution stack. Then*

$$\mathsf{F}_i^{\pi}\sigma \equiv \mathsf{F}_i^{\pi[\mathsf{Val}(\sigma)]}\bot.$$

The proof proceeds by induction on the height of the proof π. Assuming that the induction hypothesis holds for all proofs smaller than π, we show by a subsidiary induction that, given a context $C[z_0, \ldots, z_n]$ where each z_i has type of $\mathsf{F}_i^{\pi}\sigma$.

Left-to-right If a term of the form $C[\mathsf{F}_0^{\pi}\sigma, \ldots, \mathsf{F}_n^{\pi}\sigma]$ rewrites to a final term t in k steps then there is a final term t' such that $C[\mathsf{F}_0^{\pi[\mathsf{Val}(\sigma)]}\bot, \ldots, \mathsf{F}_n^{\pi[\mathsf{Val}(\sigma)]}\bot] \longrightarrow t'$ and $\mathsf{Exp}(t) = \mathsf{Exp}(t')$.
Right-to-left If a term of the form $C[\mathsf{F}_0^{\pi[\mathsf{Val}(\sigma)]}\bot, \ldots, \mathsf{F}_n^{\pi[\mathsf{Val}(\sigma)]}\bot]$ rewrites to a final term t in k steps then there is a final term t' such that $C[\mathsf{F}_0^{\pi}\sigma, \ldots, \mathsf{F}_n^{\pi}\sigma] \longrightarrow t'$ and $\mathsf{Exp}(t) = \mathsf{Exp}(t')$.

For the induction step of this subsidiary induction, we focus on the most interesting cases of the first rewrite in a rewriting sequence of length $k+1$. We focus on the left-to-right direction since the converse is similar. To simplify notation we tacitly assume at most one of $z_0, \ldots, z_n$ occur in C. For most parts of the argument this assumption can be made without loss of generality as the individual non-terminals $\mathsf{F}_0^{\pi}, \ldots, \mathsf{F}_n^{\pi}$ can be treated pointwise in all but the case of the Peirce rewrite which can, in principle, rewrite an occurrence of the non-terminal F_i^{π} to a term containing F_j^{π} for some $j \neq i$. Thus, in the following we work under the assumption that $C = C[z_i]$ and leave the full formulation for the reader to verify.

If the applied rewrite is to a non-terminal occurrence in C then the result follows directly from the induction hypothesis provided this was not a Peirce rewrite. Thus we are left to consider the case of an inference-induced rewrite applied to the non-terminal F_i^{π} directly, or a Peirce rewrite applied to either F_i^{π} or a non-terminal occurrence in C. We treat these in turn.

Reduction is left $\exists$-rule

Then π ends with an application of $\exists\mathsf{L}$. Let π_0 be the immediate subproof of π. First, suppose $i = 0$ is the index of the principal formula of the end sequent. Then the applied rewrite is:

$$\mathsf{F}_i^{\pi}\sigma\vec{u} \longrightarrow \mathsf{a}(\lambda x^{\iota}.\, \mathsf{F}_i^{\pi_0}([x/\alpha]\sigma)\vec{u})$$

and a parallel reduction is available for $\mathsf{F}_i^{\pi[\mathsf{Val}(\sigma)]}$:

$$\mathsf{F}_i^{\pi[\mathsf{Val}(\sigma)]}\bot\vec{u} \longrightarrow \mathsf{a}(\lambda x^{\iota}.\, \mathsf{F}_i^{\pi_0[\mathsf{Val}(\sigma)]}([x/\alpha]\bot)\vec{u})$$

The induction hypothesis on π_0 gives for every regular $t : \iota$:

$$\begin{aligned}\mathsf{F}_i^{\pi_0}([t/\alpha]\sigma)\vec{u} &\equiv \mathsf{F}_i^{\pi_0[\mathsf{Val}([t/\alpha]\sigma)]}\bot\vec{u} \\ &= \mathsf{F}_i^{\pi_0[\mathsf{Val}(\sigma)][\mathsf{Val}(t)/\alpha]}\bot\vec{u} && \text{as } \mathsf{FV}(t) = \mathsf{FV}(\sigma) = \emptyset. \\ &\equiv \mathsf{F}_i^{\pi_0[\mathsf{Val}(\sigma)]}([t/\alpha]\bot)\vec{u}\end{aligned}$$

If $i \neq 0$ is not the index of the principal formula, then the rule application uses pattern matching so that $\vec{u} = \mathsf{e}tv, \vec{w}$ for terms t, v where $\vec{w} = \vec{u}_{>0}$. The rewrite applied to the term is then:

$$\mathsf{F}_i^{\pi}\sigma(\mathsf{e}tv)\vec{w} \longrightarrow \mathsf{F}_i^{\pi_0}([t/\alpha]\sigma)v\vec{w}$$

Correspondingly, we have:

$$\mathsf{F}_i^{\pi[\mathsf{Val}(\sigma)]}\bot(\mathsf{e}tv)\vec{w} \longrightarrow \mathsf{F}_i^{\pi_0[\mathsf{Val}(\sigma)]}([t/\alpha]\bot)v\vec{w}$$

The induction hypothesis on π_0 similarly yields:

$$\begin{aligned}\mathsf{F}_i^{\pi_0}([t/\alpha]\sigma)v\vec{w} &\equiv \mathsf{F}_i^{\pi_0[\mathsf{Val}([t/\alpha]\sigma)]}\bot v\vec{w} \\ &= \mathsf{F}_i^{\pi_0[\mathsf{Val}(\sigma)][\mathsf{Val}(t)/\alpha]}\bot v\vec{w} \\ &\equiv \mathsf{F}_i^{\pi_0[\mathsf{Val}(\sigma)]}([t/\alpha]\bot)v\vec{u}\end{aligned}$$

from which the desired conclusion follows.

Reduction is right $\exists$-rule

Then π ends with an application of $\exists\mathsf{R}$. Let π_0 be the immediate subproof of π. If i is the index of the principal formula of the endsequent, then the corresponding rewrite is

$$\mathsf{F}_i^{\pi}\sigma\vec{u} \longrightarrow \mathsf{e}\big(t \cdot \sigma\big)\big(\mathsf{F}_i^{\pi_0}\sigma\vec{u}\big)$$

where t is the term discharged from the minor formula. In response to the rewrite, we have:

$$\begin{aligned}\mathsf{F}_i^{\pi[\mathsf{Val}(\sigma)]}\bot\vec{u} &\longrightarrow \mathsf{e}\big(t[\mathsf{Val}(\sigma)] \cdot \bot\big)\big(\mathsf{F}_i^{\pi_0[\mathsf{Val}(\sigma)]}\bot\vec{u}\big) \\ &\equiv \mathsf{e}\big(t \cdot \sigma\big)\big(\mathsf{F}_i^{\pi_0[\mathsf{Val}(\sigma)]}\bot\vec{u}\big) && \text{Proposition 9} \\ &\equiv \mathsf{e}\big(t \cdot \sigma\big)\big(\mathsf{F}_i^{\pi_0}\sigma\vec{u}\big) && \text{IH}\end{aligned}$$

as required.

Next suppose i is not the index of the principal formula. Then $\vec{u} = \vec{v}(\mathsf{a}w)$ where $\vec{v}$ has length $|\vec{u}| - 1$ and the rewrite applied to this term has the form:

$$\mathsf{F}_i^{\pi}\sigma\vec{v}(\mathsf{a}w) \longrightarrow \mathsf{F}_i^{\pi_0}\sigma\vec{v}\big(w(t \cdot \sigma)\big)$$

Respond with the rewrite:

$$\mathsf{F}_i^{\pi[\mathsf{Val}(\sigma)]}\bot\vec{v}(\mathtt{a}w) \longrightarrow \mathsf{F}_i^{\pi_0[\mathsf{Val}(\sigma)]}\bot\vec{v}(w(t[\mathsf{Val}(\sigma)]\cdot\bot))$$
$$\equiv \mathsf{F}_i^{\pi_0}\sigma\vec{v}(w(t\cdot\sigma))$$

by the same reasoning as above.

Other inference-induced rewrite

These are straightforward.

Rewrite is a Peirce reduction

The rewrite is on some subterm of $C[\mathsf{F}_i^\pi\sigma/z]$ of the form $C' = \mathsf{F}_j^{\pi'}\rho\vec{v}(\mathtt{p}s)\vec{w}$. To simplify notation we assume that j is the index of the final formula occurrence in the end sequent of π'. It suffices to consider $\mathsf{F}_i^\pi\sigma$ as a subterm of C' as otherwise the applied rewrite occurs wholly within the context C and the desired equivalence follows directly from the induction hypothesis. This leaves two cases to consider depending on whether $\mathsf{F}_j^{\pi'}\rho$ is one of the substituted occurrences of $\mathsf{F}_i^\pi\sigma$. If so then $\vec{u}$ has the form $\vec{v}(\mathtt{p}s)\vec{w}$ for appropriate $\vec{v}$, $\vec{w}$ and s, and the applied rewrite is:

$$\mathsf{F}_i^\pi\sigma\vec{v}(\mathtt{p}s)\vec{w} \longrightarrow \mathtt{p}\big(\lambda y.\,\mathsf{F}_i^\pi\sigma\vec{v}\big(s\big(\mathsf{F}_k^\pi\sigma\vec{v}\vec{w}y\big)\big)\vec{w}\big)$$

where k is the index of the argument $\mathtt{p}s$, namely $k = |\vec{v}|$. In the term $C[\mathsf{F}_i^{\pi[\mathsf{Val}(\sigma)]}\bot/z]$ we respond with the rewrite:

$$\mathsf{F}_i^{\pi[\mathsf{Val}(\sigma)]}\bot\vec{v}(\mathtt{p}s)\vec{w} \longrightarrow \mathtt{p}\big(\lambda y.\,\mathsf{F}_i^{\pi[\mathsf{Val}(\sigma)]}\bot\vec{v}\big(s\big(\mathsf{F}_k^{\pi[\mathsf{Val}(\sigma)]}\bot\vec{v}\vec{w}y\big)\big)\vec{w}\big)$$

The induction hypothesis with the appropriate choice of context completes this part of the argument.

This leaves the case where $\mathsf{F}_i^\pi\sigma$ is a subterm of s, i.e., the Peirce rewrite applies to the term C' above wherein $s = s_0[\mathsf{F}_i^\pi\sigma\vec{u}/z']$ for some term s_0. Then the applied rewrite is of the form:

$$\mathsf{F}_j^{\pi'}\rho\vec{v}\big(\mathtt{p}\big(s_0[\mathsf{F}_i^\pi\sigma/z']\big)\vec{w} \longrightarrow \mathtt{p}\big(\lambda y.\,\mathsf{F}_j^{\pi'}\rho\vec{v}\big(\big(s_0\big[\mathsf{F}_i^\pi\sigma/z'\big]\big)\big(\mathsf{F}_k^{\pi'}\rho\vec{v}\vec{w}y\big)\big)\vec{w}\big)$$

and we can respond with the analogous rewrite:

$$\mathsf{F}_j^{\pi'}\rho\vec{v}\big(\mathtt{p}\big(s_0\big[\mathsf{F}_i^{\pi[\mathsf{Val}(\sigma)]}\bot/z'\big]\big)\big)\vec{w} \longrightarrow \mathtt{p}\big(\lambda y.\,\mathsf{F}_j^{\pi'}\rho\vec{v}\big(\big(s_0\big[\mathsf{F}_i^{\pi[\mathsf{Val}(\sigma)]}\bot/z'\big]\big)\big(\mathsf{F}_k^{\pi'}\rho\vec{v}\vec{w}y\big)\big)\vec{w}\big)$$

and an application of the subsidiary induction hypothesis.

Other rewrites

These are rewrites for extractor, helper or generic evidence non-terminals, and are straightforward.

A.2 Multicut link index

Proposition 16:

Proposition. *Given two formula occurrences* (S, i) *and* (U, j) *with* $S \neq U$*, we have* $(S, i) \prec^+ (U, j)$ *if and only if* $i = \mathsf{link}(S, U)$ *and* $j \neq \mathsf{link}(U, S)$*.*

The proof proceeds by induction on the length of the unique path from S to U. For the base case we have $S \bigtriangledown U$, where $\mathsf{link}(S, U)$ and $\mathsf{link}(U, S)$ are the unique indices for which $(S, \mathsf{link}(S, U)) \bigtriangledown (U, \mathsf{link}(U, S))$. It follows directly from the definition of the direct dependency relation that $(S, i) \prec (U, j)$ iff $i = \mathsf{link}(S, U)$ and $j \neq \mathsf{link}(U, S)$.

For the induction step, suppose the unique path from S to U has length $k+1$, and suppose this path contains the unique path from V to U of length k where $S \bigtriangledown V$. Suppose $(S, i) \prec^+ (U, j)$. Then there must be some index k such that $(S, i) \prec (V, k)$ and $(V, k) \prec^+ (U, j)$. By the induction hypothesis this means that $\mathsf{link}(V, U) = k$ and $\mathsf{link}(U, V) \neq j$. Note that $\mathsf{link}(U, V) = \mathsf{link}(U, S)$, so $\mathsf{link}(U, S) \neq j$. Furthermore, by the base case of the induction we have $\mathsf{link}(S, V) = i$ and $\mathsf{link}(V, S) \neq k$. But $\mathsf{link}(S, V) = \mathsf{link}(S, U)$ so $\mathsf{link}(S, U) = i$ as required.

Conversely, suppose $\mathsf{link}(S, U) = i$ and $\mathsf{link}(U, S) \neq j$. Then $\mathsf{link}(U, V) \neq j$ since $\mathsf{link}(U, V) = \mathsf{link}(U, S)$. By the induction hypothesis we get $(V, \mathsf{link}(V, U)) \prec^+ (U, j)$. Clearly we must have $\mathsf{link}(V, U) \neq \mathsf{link}(V, S)$, so:

$$(S, i) = (S, \mathsf{link}(S, U)) = (S, \mathsf{link}(S, V)) \prec (V, \mathsf{link}(V, U))$$

by the base case of the induction. So we get

$$(S, i) \prec (V, \mathsf{link}(V, U)) \prec^+ (U, j)$$

hence $(S, i) \prec^+ (U, j)$ as required.

A.3 Language subsumption for the guarded version

Proposition 18:

Proposition. *Let* π *be a proof with* Σ_1 *end sequent* S *of the form* $\Gamma \Rightarrow A$*, and let* π' *be its guarded version. Then:*

$$\mathsf{S}_\pi \sqsupseteq T(\pi')$$

We prove the proposition by well-founded leaf-to-root induction in the end piece of π: for any formula occurrence (S, i) in the end piece of π, and for any regular substitution stack σ such that every eigenvariable that occurs in σ occurs in S, we have

$$\mathsf{F}_i^{\pi_S} \sigma \vec{\mathsf{C}} \sqsupseteq T(S, i, \pi[\mathsf{Val}(\sigma)])$$

where π_S is the generated subproof of π at the sequent S and $\vec{\mathsf{C}}$ consists of one term $\mathsf{C}_{\neg B}$ for each left formula occurrence B and one term C_B for each right formula occurrence B.

Consider the case where π_S ends with an application of a right $\exists$-rule:

$$\dfrac{\begin{array}{c}\pi_1\\ \Gamma \Rightarrow \Theta, B[t/y]\end{array}}{\pi_0 : \Gamma \Rightarrow \Theta, \exists y B}$$

Let S_1 denote the premise. We have the following non-principal rewrite:

$$\begin{aligned}
\mathsf{F}_i^{\pi_0}\sigma\vec{\mathsf{C}}\mathsf{C}_{\neg\exists yB} &\longrightarrow \mathsf{F}_i^{\pi_0}\sigma\vec{\mathsf{C}}\mathsf{a}(\lambda w.\mathsf{C}_{\neg B}w)\\
&\longrightarrow \mathsf{F}_i^{\pi_1}\sigma\vec{\mathsf{C}}((\lambda w.\mathsf{C}_{\neg B}w)(t\cdot\sigma))\\
&\longrightarrow \mathsf{F}_i^{\pi_1}\sigma\vec{\mathsf{C}}\mathsf{C}_{\neg B}(t\cdot\sigma)\\
&\longrightarrow \mathsf{F}_i^{\pi_1}\sigma\vec{\mathsf{C}}\mathsf{C}_{\neg B}\\
&= \mathsf{F}_i^{\pi_1}\sigma\vec{\mathsf{C}}\mathsf{C}_{\neg B}
\end{aligned}$$

The induction hypothesis gives:

$$\mathsf{F}_i^{\pi_1}\sigma\vec{\mathsf{C}}\mathsf{C}_{\neg B} \sqsupseteq T(S_1, i, \pi[\mathsf{Val}(\sigma)] = T(S, i, \pi[\mathsf{Val}(\sigma)]))$$

as required. The principal rewrite is simpler.

For a left $\exists$-rule:

$$\dfrac{\begin{array}{c}\pi_1\\ B[\alpha/y], \Gamma \Rightarrow \Theta\end{array}}{\pi_0 : \exists y B, \Gamma \Rightarrow \Theta}$$

Let S_1 denote the premise. Since α cannot occur in the sequent S, by our assumption on the substitution stack σ we have $\alpha \notin \mathsf{FV}(\sigma) \cup \mathsf{BV}(\sigma)$. We have the non-principal rewrite:

$$\begin{aligned}
\mathsf{F}_i^{\pi_0}\sigma\mathsf{C}_{\exists yB}\vec{\mathsf{C}} &\longrightarrow \mathsf{F}_i^{\pi_0}\sigma\mathsf{e}(\mathsf{c}, \mathsf{C}_B)\vec{\mathsf{C}}\\
&\longrightarrow \mathsf{F}_i^{\pi_1}[\mathsf{c}/\alpha]\sigma\mathsf{C}_B\vec{\mathsf{C}}
\end{aligned}$$

Note that the substitution stack $[\mathsf{c}/\alpha]\sigma$ is still regular provided that σ is, and every variable that occurs in it also occurs in the sequent S_1. Applying the induction hypothesis we get:

$$\begin{aligned}
\mathsf{F}_i^{\pi_1}[\mathsf{c}/\alpha]\sigma\mathsf{C}_B\vec{\mathsf{C}} &\sqsupseteq T(S_1, i, \pi[\mathsf{Val}([\mathsf{c}/\alpha]\sigma)])\\
&= T(S_1, i, \pi[\mathsf{Val}(\sigma)][\mathsf{c}/\alpha]) \quad \alpha \notin \mathsf{FV}(\sigma) \cup \mathsf{BV}(\sigma)\\
&= T(S, i, \pi[\mathsf{Val}(\sigma)])
\end{aligned}$$

as required. The principal rewrite is simpler.

The remaining cases are left to the reader. To finish the proof of the proposition, apply the claim that we proved by induction to the root sequent with substitution stack $\bot$.

A.4 Correctness of external multicut reductions

A.4.1 Disjunction case

Let $(\mathbb{S}, \bigtriangledown, \pi) \Rightarrow (\mathbb{S}^*, \bigtriangledown^*, \pi^*)$ result via the external reduction for disjunction on page 36. We prove Proposition 21:

Proposition. *For any formula occurrence (S, i) of the multicut, we have:*

1. $O(P_0, p) \sqsupseteq \mathsf{j}^j(O^*_L(P_1, p), O^*_R(P_2, p))$.
2. *If $i \neq p$ then:* $O(P_0, i) \sqsupseteq \{O^*_L(P_1, i), O^*_R(P_2, i)\}$
3. *If $S \neq P_0$ and (S, i) is a side formula occurrence then* $O(S, i) \sqsupseteq \{O^*_L(S, i), O^*_R(S, i)\}$.

Here, the terms $O^*_L(S, i)$ and $O^*_R(S, i)$ refer to the output terms in the left and right reduced multicut respectively.

Say that a formula occurrence (S, i) *depends on* P_0 if $S = P_0$ or $(P_0, k) \prec^+ (S, i)$ for some k.

Claim 1. *If S does not depend on P_0 then $O(S, i) \sqsupseteq O^*_L(S, i)$ and $O(S, i) \sqsupseteq O^*_R(S, i)$.*

Proof. By well-founded induction on $\prec$. Suppose $O(S, i) = F_i^{\pi_S} \bot \vec{u}$. Each term $\vec{u}$ is either of the form C_D or $\mathsf{C}_{\neg D}$ for some formula D, or of the form $O(U, k)$ for some (U, k). In the latter case the formula occurrence (U, k) also does not depend on P_0, for if it were dependent on P_0 we would have either $P_0 = U$ or $(P_0, k') \prec^+ (U, k)$ for some k'. But $(U, k) \prec (S, i)$, so in either case we get $(P_0, k') \prec^+ (S, i)$ for some k', contradicting our assumption that (S, i) does not depend on P_0. Hence we can apply the induction hypothesis to (U, k) and get $O(U, k) \sqsupseteq O^*_L(U, k)$ as well as $O(U, k) \sqsupseteq O^*_R(U, k)$.

Let $\vec{u}^*_L$ be the result of replacing each term of the form $O(U, k)$ in $\vec{u}$ by $O^*_L(U, k)$. Then we get:

$$\begin{aligned} O(S, i) &= \mathsf{F}_i^{\pi_S} \bot \vec{u} \\ &\sqsupseteq \mathsf{F}_i^{\pi_S} \bot \vec{u}^*_L \\ &= O^*_L(S, i) \end{aligned}$$

Similarly we get $O(S, i) \sqsupseteq O^*_R(S, i)$. □

We now prove all three items of the proposition by simultaneous well-founded induction on $\prec$.

Item (1) is a straightforward application of the principal rewrite rule for a proof ending with a left $\vee$-rule together with Claim 1. Item (2) follows by applying the non-principal rewrite rule for a proof ending with a left $\vee$-rule, observing that the input term $I(P_0, p)$ for the principal index p is $\mathsf{C}_{A \vee B}$ since the principal formula is a side formula of the multicut. So writing $O(P_0, i) = \mathsf{F}_i^{\pi_{P_0}} \bot \vec{u} \mathsf{C}_{A \vee B} \vec{v}$ we can rewrite the term as follows:

$$\begin{aligned} &\mathsf{F}_i^{\pi_{P_0}} \bot \vec{u} \mathsf{C}_{A \vee B} \vec{v} \\ \longrightarrow\ &\mathsf{F}_i^{\pi_{P_0}} \bot \vec{u} (\mathtt{i}_0 \mathsf{C}_A) \vec{v} \parallel \mathsf{F}_i^{\pi_{P_0}} \bot \vec{u} (\mathtt{i}_1 \mathsf{C}_B) \vec{v} \\ \longrightarrow\ &\mathsf{F}_i^{\pi_{P_1}} \bot \vec{u} \mathsf{C}_A \vec{v} \parallel \mathsf{F}_i^{\pi_{P_2}} \bot \vec{u} \mathsf{C}_B \vec{v} \end{aligned}$$

The desired conclusion now follows by applying Claim 1 to the arguments $\vec{u}, \vec{v}$; note that if one of these arguments are of the form $O(U, k)$ for some formula occurrence

(U, k), then this formula occurrence does not depend on P_0 by the constraints on a multicut.

For item (3), we note that any side formula occurrence (U, i) depends on P_0: if $U = P_0$ this is immediate, and for $U \neq P_0$ it suffices to note that the $\triangledown$-path from U to P_0 has the form: $(U, j) \triangledown (V_0, k_0) \neq (V_0, k'_0) \triangledown \ldots \triangledown (V_n, k_n) \neq (V_n, k'_n) \triangledown (P_0, m)$ Since (U, i) is a side formula occurrence we get $i \neq j$, so $(P_0, m) \prec^+ (U, i)$. With this in mind, we prove by a subsidiary induction on the length of the $\triangledown$-path from U to P_0 that, if (U, i) depends on P_0 and $U \neq P_0$, then $O(U, i) \sqsupseteq O^*_L(U, i) \parallel O^*_R(U, i)$. Item (3) follows from this stronger claim.

The base case of the induction is given by $(U, j) \triangledown (P_0, k)$ for some j, k where that $j \neq i$ since (S, i) is assumed to depend on P_0. Furthermore $k \neq p$ since (P_0, p) is a side formula occurrence. The term $O(U, i)$ is of the form $\mathsf{F}_i^{\pi_U} \bot \vec{u}(O(P_0, j))\vec{v}$, and we can rewrite:

$$\begin{aligned} &\mathsf{F}_i^{\pi_U} \bot \vec{u}(O(P_0, j))\vec{v} \\ \longrightarrow\ &\mathsf{F}_i^{\pi_U} \bot \vec{u}(O^*_L(P_1, j))\vec{v} \parallel \mathsf{F}_i^{\pi_U} \bot \vec{u}(O^*_R(P_2, j))\vec{v} \quad \text{Item (2) of IH} \end{aligned}$$

Note that if a term in $\vec{u}, \vec{v}$ is of the form $O(U', i')$ for some formula occurrence (U', i'), then this formula occurrence does not depend on P_0 by the constraints on multicuts. With this in mind we can apply Claim 1 to get:

$$\mathsf{F}_i^{\pi_U} \bot \vec{u}(O^*_L(P_1, j))\vec{v} \longrightarrow O^*_L(U, i)$$

$$\mathsf{F}_i^{\pi_U} \bot \vec{u}(O^*_R(P_2, j))\vec{v} \longrightarrow O^*_R(U, i)$$

Hence:

$$O(U, j) \sqsupseteq \{O^*_L(U, i), O^*_R(U, i)\}$$

as required.

For the induction step, suppose the $\triangledown$-path from U to P_0 has length $n + 1$. Then since (U, i) depends on P_0 there is some $j \neq i$ and some formula occurrence (U', j') with $(U, j) \triangledown (U', j')$, where (U', j') depends on P_0 and the length of the $\triangledown$-path from U' to P_0 has length n. So (U, i) is of the form $\mathsf{F}_i^{\pi_U} \bot \vec{u}(O(U', j'))\vec{v}$. Using the induction hypothesis on n we can now get:

$$\mathsf{F}_i^{\pi_U} \bot \vec{u}(O(U', j'))\vec{v} \sqsupseteq \{\mathsf{F}_i^{\pi_U} \bot \vec{u}(O^*_L(U', j'))\vec{v}, \mathsf{F}_i^{\pi_U} \bot \vec{u}(O^*_R(U', j'))\vec{v}\}$$

If a term in $\vec{u}, \vec{v}$ is of the form $O(V, m)$ for some formula occurrence (V, m), then this formula occurrence does not depend on P_0 by the constraints on multicuts, since (U', j') does depend on P_0. With this in mind we can apply Claim 1 to get $O(U, i) \sqsupseteq O^*_L(U, i) \parallel O^*_R(U, i)$ as required.

A.4.2 Existential case

Let $(\mathbb{S}, \triangledown, \pi) \Rightarrow (\mathbb{S}^*, \triangledown^*, \pi^*)$ arise from the external existential reduction on page 37. We prove Proposition 22:

Proposition. *For any formula occurrence (S, i), we have:*

1. *If* $S = P_0$ *and* $i \neq p$ *then* $O(P_0, i) \sqsupseteq O^*(P_1, i)$.
2. $O(P_0, p) \sqsupseteq \mathsf{e}(t \cdot \bot, O^*(P_1, p))$.
3. *If* $S \neq P_0$ *then* $O(S, i) \sqsupseteq O^*(S, i)$.

We prove the three items by simultaneous well-founded induction on $\prec$.

For item (1), note that since (P_0, p) is a right side formula occurrence we have $I(P_0, p) = \mathsf{C}_{\neg\exists x A}$. So we can write $O(P_0, i) = \mathsf{F}_i^{\pi_{P_0}} \bot \vec{u} \mathsf{C}_{\neg\exists A} \vec{v}$. For each index $j \notin \{i, p\}$, $I(P_0, j)$ is either $\mathsf{C}_{P_0[j]}$, or $\mathsf{C}_{\neg P_0[j]}$, or of the form $O(U, k)$ for some $(U, k) \bigtriangledown (P_0, j)$. Then $U \neq P_0$ so item (3) of the induction hypothesis gives $O(U, k) \sqsupseteq O^*(U, k)$. Let $\vec{u}^*, \vec{v}^*$ be the result of replacing each input term $O(U, k)$ by $O^*(U, k)$. We have:

$$\begin{aligned} O(P_0, i) &= \mathsf{F}_i^{\pi_{P_0}} \bot \vec{u} \mathsf{C}_{\neg\exists x A} \vec{v} \\ &\sqsupseteq \mathsf{F}_i^{\pi_{P_1}} \bot \vec{u} \mathsf{C}_{\neg A} \vec{v} \\ &\sqsupseteq \mathsf{F}_i^{\pi_{P_1}} \bot \vec{u}^* \mathsf{C}_{\neg A} \vec{v}^* \\ &= O^*(P_1, i) \end{aligned}$$

as required.

For item (2), write $O(P_0, p) = \mathsf{F}_p^{\pi_{P_0}} \bot \vec{u}$. For each index $i \neq p$, $I(P_0, i)$ is either $\mathsf{C}_{P_0[i]}$, or $\mathsf{C}_{\neg P_0[i]}$, or of the form $O(U, k)$ for some $(U, k) \bigtriangledown (P_0, i)$. In the latter case $U \neq P_0$ so item (3) of the induction hypothesis gives $O(U, k) \sqsupseteq O^*(U, k)$. Let $\vec{u}^*$ be the result of replacing each input term $O(U, k)$ by $O^*(U, k)$. We have:

$$\begin{aligned} O(P_0, p) &= \mathsf{F}_p^{\pi_{P_0}} \bot \vec{u} \\ &\sqsupseteq \mathsf{F}_p^{\pi_{P_0}} \bot \vec{u}^* \\ &\sqsupseteq \mathsf{e}(t \cdot \bot, \mathsf{F}_p^{\pi_{P_1}} \bot \vec{u}^*) \\ &= \mathsf{e}(t \cdot \bot, O^*(P_1, p) \end{aligned}$$

as required.

The proof of item (3) is straightforward.

A.4.3 Contraction case

Let $(\mathbb{S}, \bigtriangledown, \pi) \Rightarrow (\mathbb{S}^*, \bigtriangledown^*, \pi^*)$ result via the external contraction reduction on page 38: We prove Proposition 23:

Proposition. *For any formula occurrence* (S, i) *we have:*

1. *If* $S = P_0$ *and* $i < p$ *then* $O(P_0, j) \sqsupseteq O^*(P_1, j)$.
2. *If* $S = P_0$ *and* $i > p$ *then* $O(P_0, j) \sqsupseteq O^*(P_1, j + 1)$.
3. $\mathsf{E}(O(P_0, p)) \sqsupseteq \{\mathsf{E}(O^*(P_1, p)), \mathsf{E}(O^*(P_1, p + 1))\}$.
4. *If* $S \neq P_0$ *then* $O(S, i) \sqsupseteq O^*(S, i)$.

We prove all four items by a simultaneous well-founded induction on $\prec$.

For item (1), note that since (P_0, p) is a side formula occurrence we have $I(P_0, p) = \mathsf{C}_A$ or $\mathsf{C}_{\neg A}$. Let's assume that (P_0, p) is a left formula occurrence so that $I(P_0, p) = \mathsf{C}_A$; the other case is proved in the same way. We may write $O(P_0, i) = \mathsf{F}_i^{\pi_{P_0}} \bot \vec{u} \mathsf{C}_A \vec{v}$. For each index $j \notin \{i, p\}$, either $I(P_0, j)$ is $\mathsf{C}_{P_0[j]}$, or $\mathsf{C}_{\neg P_0[j]}$, or it is of the form $O(U, k)$

for some $(U,k) \bigtriangledown (P_0,j)$. In the latter case we get $O(U,k) \sqsupseteq O^*(U,k)$ by item (4) of the induction hypothesis. Let $\vec{u}^*, \vec{v}^*$ be the result of replacing input terms $O(U,k)$ by $O^*(U,k)$. We get:

$$\begin{aligned} O(P_0,i) &= \mathsf{F}_i^{\pi_{P_0}} \bot \vec{u} \mathsf{C}_A \vec{v} \\ &\sqsupseteq \mathsf{F}_i^{\pi_{P_1}} \bot \vec{u} \mathsf{C}_A \mathsf{C}_A \vec{v} \\ &\sqsupseteq \mathsf{F}_i^{\pi_{P_1}} \bot \vec{u}^* \mathsf{C}_A \mathsf{C}_A \vec{v}^* \\ &= O^*(P_1,i) \end{aligned}$$

as required. Item (2) is similar so we omit it.

For item (3), we write $O(P_0,p) = \mathsf{F}_p^{\pi_{P_0}} \bot \vec{u}\vec{v}$ where $\vec{u}$ are the input terms for indices below p. For each index $j \neq p$, either $I(P_0,j)$ is $\mathsf{C}_{P_0[j]}$, or $\mathsf{C}_{\neg P_0[j]}$, or it is of the form $O(U,k)$ for some $(U,k) \bigtriangledown (P_0,j)$. In the latter case we get $O(U,k) \sqsupseteq O^*(U,k)$ by item (4) of the induction hypothesis. Let $\vec{u}^*, \vec{v}^*$ be the result of replacing input terms $O(U,k)$ by $O^*(U,k)$. We get:

$$\begin{aligned} \mathsf{E}(O(P_0,p)) &= \mathsf{E}(\mathsf{F}_p^{\pi_{P_0}} \bot \vec{u}\vec{v}) \\ &\sqsupseteq \mathsf{E}(\mathsf{F}_p^{\pi_{P_0}} \bot \vec{u}^*\vec{v}^*) \\ &\sqsupseteq \{\mathsf{E}(\mathsf{p}(\lambda x\, \mathsf{F}_p^{\pi_{P_1}} \bot \vec{u}^* x \vec{v}^*)), \mathsf{E}(\mathsf{p}(\lambda x\, \mathsf{F}_{p+1}^{\pi_{P_1}} \bot \vec{u}^* x \vec{v}^*))\} \\ &\sqsupseteq \{\mathsf{E}(\mathsf{F}_p^{\pi_{P_1}} \bot \vec{u}^* \mathsf{C}_A \vec{v}^*), \mathsf{E}(\mathsf{p}(\mathsf{F}_{p+1}^{\pi_{P_1}} \bot \vec{u}^* \mathsf{C}_A \vec{v}^*))\} \\ &= \{\mathsf{E}(O^*(P_1,p)), \mathsf{E}(O^*(P_1,p+1))\} \end{aligned}$$

as required.

For item (4), write $O(S,i) = \mathsf{F}_i^{\pi_S} \bot \vec{u}$. For each index $j \neq i$, $I(S,j)$ is either $\mathsf{C}_{S[i]}$, or $\mathsf{C}_{\neg S[i]}$, or it is $O(U,k)$ for some $(U,k) \bigtriangledown (S,j)$. If $U \neq P_0$ then item (4) of the induction hypothesis gives $O(U,k) \sqsupseteq O^*(U,k)$. If $U = P_0$, then $k \neq p$ since (P_0,p) was a side formula occurrence. If $k < p$ then item (1) of the induction hypothesis gives $O(P_0,k) \sqsupseteq O^*(P_1,k)$, and if $k > p$ then item (1) of the induction hypothesis gives $O(P_0,k) \sqsupseteq O^*(P_1,k+1)$. Let $\vec{u}^*$ be the result of replacing each input term $O(U,k)$ for $U \neq P_0$ by $O^*(U,k)$, each input term $O(P_0,k)$ with $k < p$ by $O^*(P_1,k)$ and each input term $O(P_0,k)$ with $k > p$ by $O^*(P_1,k+1)$. We get:

$$\begin{aligned} O(S,i) &= \mathsf{F}_i^{\pi_S} \bot \vec{u} \\ &\sqsupseteq \mathsf{F}_i^{\pi_S} \bot \vec{u}^* \\ &= O^*(S,i) \end{aligned}$$

as required.

A.5 Correctness of internal multicut reductions

A.5.1 Weakening case

We restate proposition 25.

Proposition. *Let $(\mathbb{S}, \bigtriangledown, \pi) \Rightarrow (\mathbb{S}^*, \bigtriangledown^*, \pi^*)$ be the multicut reduction via an internal weakening reduction as on page 41. Let (S,i) be a side formula occurrence in one of the premises.*

1. *If S is unaffected then $O(S,i) \sqsupseteq O^*(S,i)$.*
2. *If $S = P$ then $O(P,i) \sqsupseteq O^*(P',i')$, where i' is the index associated with i via the weakening inference.*
3. *If S is affected then $O(S,i) \sqsupseteq \bot_{S[i]}$.*

Items (1) and (2) are fairly straightforward to prove. Item (3) is proved by an induction on the length of the unique $\triangledown$-path from S to P. More precisely, we prove the following stronger statement by such an induction: for any formula occurrence (S,i) for which S is affected, if i is *not* the index of the formula occurrence in S that is linked to the formula occurrence (P,p) via the unique $\triangledown$-path from S to P, then $O(S,i) \sqsupseteq \bot_{S[i]}$.

For the base step of this induction we have $(S,j) \triangledown (P,p)$ for some $j \neq i$, which means that $O(S,i)$ is of the form:

$$\mathsf{F}_i^{\pi_S} \bot \vec{u}(\mathsf{F}_p^{\pi_P} \bot \vec{v})\vec{w}$$

But since $\mathsf{F}_p^{\pi_P} \bot \vec{w} \longrightarrow \bot_{P[p]}$ by the rewrite rules for weakening (recalling that (P,p) was the principal formula occurrence of the weakening inference), we can apply the rewrite rule:

$$\mathsf{F}_i^{\pi_S} \bot \vec{u}(\bot_{P[p]})\vec{w} \longrightarrow \bot_{S[i]}$$

to get $O(S,i) \longrightarrow \bot_{S[i]}$, hence $O(S,i) \sqsupseteq \bot_{S[i]}$ as required.

For the induction step, suppose the length of the $\triangledown$-path from S to P is $n+1$. Then there is some U and indices j,k such that $(S,j) \triangledown (U,k)$ such that the length of the $\triangledown$-path from U to P is n. By assumption $i \neq j$, and furthermore k cannot be the index of the cut formula occurrence that is linked to (P,p) via the $\triangledown$-path from U to P, since that would contradict $(S,j) \triangledown (U,k)$. So our induction hypothesis gives $O(U,k) \sqsupseteq \bot_{U[k]}$. Since $O(S,i)$ is of the form $\mathsf{F}_i^{\pi_S} \bot \vec{u}(O(U,k))\vec{v}$ (because $i \neq j$), we get:

$$\begin{aligned} O(S,i) &= \mathsf{F}_i^{\pi_S} \bot \vec{u}(O(U,k))\vec{v} \\ &\sqsupseteq \mathsf{F}_i^{\pi_S} \bot \vec{u} \bot_{U[k]} \vec{v} \\ &\sqsupseteq \bot_{S[i]} \end{aligned}$$

as required.

A.5.2 Contraction case

We prove proposition 26:

Proposition. *Let $(\mathbb{S}, \triangledown, \pi)$ and $(\mathbb{S}^*, \triangledown^*, \pi^*)$ be an instance of the internal contraction reduction (p. 42), let (S,i) be a side formula occurrence with $S \in \mathbb{S}$.*

- *If $S \neq R_0$ is affected by the reduction then*

$$\mathsf{E}(O(S,i)) \sqsupseteq \{E(O^*(S_L,i)), E(O^*(S_R,i))\}$$

- *If $S = R_0$ we have $O(R_0,i) \sqsupseteq O^*(R_1,i)$ if $i < p$, and $O(R_0,i) \sqsupseteq O^*(R_1,i+1)$ if $i > p$.*

- *In all other cases we have* $\mathsf{E}(O(S,i)) \sqsupseteq \mathsf{E}(O^*(S,i))$.

We shall need to prove a number of claims. The first claim will help us with the case distinctions we have to make:

Claim 2. *If* $S \in \mathbb{S}$ *is affected by the reduction, then there is some cut formula occurrence* (S,i) *such that* $(R_0,p) \prec^+ (S,j)$ *for each* $j < l(S)$ *with* $j \neq i$.

Proof of Claim 2. If S is affected by the reduction we get $\mathsf{link}(R_0,S) = \mathsf{link}(R_0,Q) = p$. The claim now follows from Proposition 16. □

With Claim 2, we see that for an arbitrary *side* formula occurrence (S,i) appearing in the multicut to be reduced, there are the following cases: (1) $S = R_0$ and $i \neq p$, (2) $S = Q$ and $i \neq q$, (3) $S \notin \{R_0, Q\}$ and S is unaffected by the reduction, and (4) $S \notin \{R_0, Q\}$ and S is affected by the reduction, in which case we have $(R_0,p) \prec^+ (S,i)$ by the Claim. However, this case distinction is only valid for side formula occurrences, and since we want to proceed by well-founded induction on the relation $\prec$ we need to say something about how the prior output term $O(S,i)$ relates to posterior output terms for any formula occurrence (S,i), including cut formula occurrences. The most fundamental case distinction is: either $(R_0,p) \prec^+ (S,i)$, or not. The latter case splits into precisely the following three different subcases: (a) $S = R_0$, (b) $S \neq R_0$ and S is unaffected by the reduction, or (c) $S \neq R_0$ and S is affected by the reduction. For the first case (a), the case where $i = p$ needs special attention and is treated later in Claim 4. Cases (b,c) and the remaining sub-case of (a) where $i \neq p$ are handled by the following claim.

Claim 3. *Let* (S,i) *be any formula occurrence, and suppose it is* not *the case that* $(R_0,p) \prec^+ (S,i)$. *By Proposition 16 this is equivalent to saying that either* $S = R_0$, *or* $\mathsf{link}(R_0,S) \neq p$, *or* $\mathsf{link}(S,R_0) = i$. *Then the following statements hold:*

1. *If* $S = R_0$ *and* $i \neq p$, *we have* $O(R_0,i) \sqsupseteq O^*(R_1,i)$ *if* $i < p$, *and* $O(R_0,i) \sqsupseteq O^*(R_1,i+1)$ *if* $i > p$.
2. *If* $S \neq R_0$ *is not affected by the reduction then* $O(S,i) \sqsupseteq O^*(S,i)$.
3. *If* $S \neq R_0$ *is affected by the reduction then* $O(S,i) \sqsupseteq O^*(S_L,i)$ *and* $O(S,i) \sqsupseteq O^*(S_R,i)$.

Proof. We prove all three statements by a simultaneous well-founded induction on the dependency relation $\prec$.

For item (1), suppose $S = R_0$ and $i \neq p$, and suppose the induction hypothesis holds for all formula occurrences (U,j) with $(U,j) \prec^+ (R_0,i)$. We assume $i < p$ since the other case is similar. Let $O(R_0,i) = F_i^{\pi R_0} \bot \vec{u}\vec{v}(O(Q,q))\vec{w}$, where $\vec{u}$ are the inputs for indices below i and $\vec{v}$ the inputs for indices between i and p. For each index $k < l(R_0)$, $k \notin \{i,p\}$, the corresponding input term is either the constant $\mathsf{C}_{R_0[k]}$, or $\mathsf{C}_{\neg R_0[k]}$, or a canonical output term $O(U,j)$ where $(U,j) \prec (R_0,i)$ and U is unaffected by the reduction. By item (2) of the induction hypothesis, which is currently assumed for all formula occurrences $\prec$-smaller than (R_0,i), we get $O(U,j) \sqsupseteq O^*(U,j)$ for each such formula occurrence. Let $\vec{u}^*, \vec{v}^*, \vec{w}^*$ be the result of replacing each input term $O(U,j)$ by $O^*(U,j)$. Since it is not the case that $(R_0,p) \prec^+ (Q,q)$, item (3) of the induction hypothesis gives $O(Q,q) \sqsupseteq O^*(Q_L,q)$ and $O(Q,q) \sqsupseteq O^*(Q_R,q)$. So we get:

$$\begin{aligned} O(R_0,i) &= \mathsf{F}_i^{\pi_{R_0}} \bot \vec{u}\vec{v}(O(Q,q))\vec{w} \\ &\sqsupseteq \mathsf{F}_i^{\pi_{R_1}} \vec{u}\vec{v}(O(Q,q))(O(Q,q))\vec{w} \\ &\sqsupseteq \mathsf{F}_i^{\pi_{R_1}} \vec{u}^*\vec{v}^*(O^*(Q_L,q))(O^*(Q_R,q))\vec{w}^* \quad \text{IH} \\ &= O^*(R_0,i) \end{aligned}$$

as required.

For item (2), suppose $S \neq R_0$ and S is unaffected by the reduction. We write $O(S,i) = F_i^{\pi_S} \bot \vec{u}\vec{v}$ where $\vec{u}$ are the input terms for indices below i. For each index $k \neq i$, the input term is either $\mathsf{C}_{S[k]}$, or $\mathsf{C}_{\neg S[k]}$, or $O(U,j)$ for some $(U,j) \bigtriangledown (S,k)$.There are two possibilities: either $U = R_0$, or $U \neq R_0$ and U is also unaffected by the reduction. In the former case we must have $j \neq p$ since otherwise we would have $(S,i) = (Q,q)$ and therefore S would be affected by the reduction. Note that we have $(R_0,j) \prec (S,i)$ since $k \neq i$. The induction hypothesis is currently assumed for all formula occurrences $\prec$-smaller than (S,i), hence item (1) of the induction hypothesis gives $O(R_0,j) \sqsupseteq O^*(R_0,j)$ if $j < p$ and $O(R_0,j) \sqsupseteq O^*(R_0,j+1)$ if $j > p$. In the latter case (i.e. $U \neq R_0$), we again have $(U,j) \prec (S,i)$ since $k \neq i$, hence item (2) of the induction hypothesis (currently assumed for all formula occurrences $\prec$-smaller than (S,i)) gives $O(U,j) \sqsupseteq O^*(U,j)$. Let $\vec{u}^*, \vec{v}^*$ be the result of replacing each input term $O(R_0,j)$ for $j < p$ by $O^*(R_0,j)$, each input term $O(R_0,j)$ for $j > p$ by $O^*(R_0,j+1)$ and each input term $O(U,j)$ for $U \neq R_0$ by $O^*(U,j)$. Then we get:

$$\begin{aligned} O(S,i) &= \mathsf{F}_i^{\pi_S} \bot \vec{u}\vec{v} \\ &\sqsupseteq \mathsf{F}_i^{\pi_S} \bot \vec{u}^*\vec{v}^* \\ &= O^*(S,i) \end{aligned}$$

as required.

For item (3), suppose $S \neq R_0$ is affected by the reduction. We write $O(S,i) = \mathsf{F}_i^{\pi_S} \bot \vec{u}\vec{v}$ where $\vec{u}$ are the input terms for indices below i. For each index $k \neq i$, the input term $I(S,k)$ is either $\mathsf{C}_{S[k]}$, or $\mathsf{C}_{\neg S[k]}$, or $O(U,j)$ for some $(U,j) \bigtriangledown (S,k)$. In the latter case it must be that $U \neq R_0$ and U is also affected by the reduction. For suppose $U = R_0$; then we must have $j \neq p$ since otherwise we would have $(R_0,p) \prec (S,i)$ contradicting our assumption. But since we assumed S was affected by the reduction, its unique path to R_0 in the cut-connectedness graph of the multicut must contain Q, which means that in fact $S = Q$, $k = q$ and $j = p$, contradiction. So we have $U \neq R_0$, and U is affected by the reduction since S is. Since $k \neq i$ we get $(U,j) \prec (S,i)$, and the induction hypothesis is currently assumed for all formula occurrences $\prec$-smaller than (S,i). So item (3) of the induction hypothesis gives $O(U,j) \sqsupseteq O^*(U_L,j)$ and $O(U,j) \sqsupseteq O^*(U_R,j)$. Let $\vec{u}_L^*, \vec{v}_L^*$ be the result of replacing each input term $O(R_0,j)$ for $j < p$ by $O^*(R_0,j)$, each input term $O(R_0,j)$ for $j > p$ by $O^*(R_0,j+1)$ and each input term $O(U,j)$ for $U \neq R_0$ by $O^*(U_L,j)$. Then we get:

$$\begin{aligned} O(S,i) &= \mathsf{F}_i^{\pi_S} \bot \vec{u}\vec{v} \\ &\sqsupseteq \mathsf{F}_i^{\pi_{S_L}} \bot \vec{u}_L^* \vec{v}_L^* \\ &= O^*(S_L,i) \end{aligned}$$

A similar argument shows that $O(S,i) \sqsupseteq O^*(S_R,i)$. □

Now we proceed with the remaining sub-case of (a) that was left out in the previous claim.

Claim 4. *$O(R_0,p) \sqsupseteq \{\mathtt{p}(\lambda x\, t), \mathtt{p}(\lambda x\, s)\}$ for some terms t,s such that $O^*(R_1,p) = t[I^*(R_1,p+1)/x]$ and $O^*(R_1,p+1) = s[I^*(R_1,p)/x]$.*

Proof. Write $O(R_0,p) = \mathsf{F}_p^{\pi_{R_0}} \bot \vec{u}\vec{v}$ where $\vec{u}$ are the input terms for indices below p. For each index $k \neq p$, $k < l(R_0)$, the input term is either $\mathsf{C}_{R_0[k]}$ or $\mathsf{C}_{\neg R_0[k]}$ or $O(U,j)$ for some $(U,j) \bigtriangledown (R_0,k)$. In the latter case U is unaffected by the reduction since $k \neq p$, so item (2) of Claim 3 gives $O(U,j) \sqsupseteq O^*(U,j)$. Let $\vec{u}^*\vec{v}^*$ be the result of replacing each input term $O(U,j)$ by $O^*(U,j)$. Then we get:

$$\begin{aligned} O(R_0,p) &= \mathsf{F}_p^{\pi_{R_0}} \bot \vec{u}\vec{v} \\ &\sqsupseteq \{\mathtt{p}(\lambda x\, \mathsf{F}_p^{\pi_{R_1}} \bot \vec{u}x\vec{v}), \mathtt{p}(\lambda x\, \mathsf{F}_{p+1}^{\pi_{R_1}} \bot \vec{u}x\vec{v})\} \\ &\sqsupseteq \{\mathtt{p}(\lambda x\, \mathsf{F}_p^{\pi_{R_1}} \bot \vec{u}^* x\vec{v}^*), \mathtt{p}(\lambda x\, \mathsf{F}_{p+1}^{\pi_{R_1}} \bot \vec{u}^* x\vec{v}^*)\} \end{aligned}$$

We set $t = \mathsf{F}_p^{\pi_{R_1}} \bot \vec{u}^* x\vec{v}^*$ and $s = \mathsf{F}_{p+1}^{\pi_{R_1}} \bot \vec{u}^* x\vec{v}^*$. Clearly then $O^*(R_1,p) = t[I^*(R_1,p+1)/x]$ and $O^*(R_1,p+1) = s[I^*(R_1,p)/x]$. □

We now proceed with the remaining case of a formula occurrence (S,i) such that $(R_0,p) \prec^+ (S,i)$. This case will be handled using a subsidiary induction on the length of the path from S to Q, where as before we let (Q,q) denote the unique formula occurrence such that $(Q,q) \bigtriangledown (R_0,p)$. The base case is thus $S = Q$ and $i \neq q$, which we treat in a separate claim.

Claim 5. *For $i < l(Q)$ with $i \neq q$, there are terms t,s such that $O(Q,i) \sqsupseteq \mathtt{p}(\lambda x\, t) \parallel \mathtt{p}(\lambda x\, s)$ for some terms t,s, such that:*

$$O^*(Q_L,i) = t[I^*(Q_R,i)/x] \text{ and } O^*(Q_R,i) = s[I^*(Q_L,i)/x].$$

Proof. We assume $i < q$ since the other case is treated the same way. The term $O(Q,i)$ is then of the form:

$$F_i^{\pi_Q} \bot \vec{u}\vec{v} O(R_0,p)\vec{w}$$

where $\vec{u}$ are the canonical input terms for formula occurrences (Q,j) with $j < i$, $\vec{v}$ are the canonical input terms for formula occurrences (Q,j) with $i < j < q$ and $\vec{w}$ are the canonical input terms for formula occurrences (Q,j) with $j > q$. Note that given an input term $I(Q,j)$ for $j \notin \{q,i\}$, there are two possibilities for (Q,j): it is a side formula occurrence and, in that case, $I(Q,j) = \mathsf{C}_{Q[j]} = I^*(Q_L,j) = I^*(Q_R,j)$ or $I(Q,j) = \mathsf{C}_{\neg Q[j]} = I^*(Q_L,j) = I^*(Q_R,j)$, or (Q,j) is a cut formula occurrence and $I(Q,j) = O(U,k)$ for some $(U,k) \bigtriangledown (Q,j)$. Note that, in this case, U is affected by the reduction but it is not the case that $(R_0,p) \prec^+ (U,k)$ (specifically, we have $(R_0,p) \prec (Q,i) \bigtriangledown (U,k)$). So by Claim 3, we get $O(U,k) \sqsupseteq O^*(U_L,k) = I^*(Q_L,j)$ and $O(U,k) \sqsupseteq O^*(U_R,k) = I^*(Q_R,j)$. Hence it follows that in each case, $I(Q,j) \sqsupseteq I^*(Q_L,j)$ and $I(Q,j) \sqsupseteq I^*(Q_R,j)$. Let $\vec{u}^*_L, \vec{v}^*_L, \vec{w}^*_L$ denote the result of replacing each input term $I(Q,j)$ among $\vec{u},\vec{v},\vec{w}$ by $I^*(Q_L,j)$, and similarly let $\vec{u}^*_R, \vec{v}^*_R, \vec{w}^*_R$ denote the result of replacing each input term $I(Q,j)$ among $\vec{u},\vec{v},\vec{w}$ by $I^*(Q_R,j)$.

By Claim 4, $O(R_0,p) \sqsupseteq \{\mathsf{p}(\lambda x\, t_0), \mathsf{p}(\lambda x\, s_0)\}$ for some terms t_0, s_0 such that $O^*(R_1,p) = t_0[I^*(R_1,p+1)/x]$ and $O^*(R_1,p+1) = s_0[I^*(R_1,p)/x]$. So we get:

$$
\begin{aligned}
O(Q,i) &= \mathsf{F}_i^{\pi_Q} \bot \vec{u}\vec{v} O(R_0,p)\vec{w} \\
&\sqsupseteq \{\mathsf{F}_i^{\pi_Q} \bot \vec{u}\vec{v}(\mathsf{p}(\lambda x\, t_0)\vec{w}, \mathsf{F}_i^{\pi_Q} \bot \vec{u}\vec{v}(\mathsf{p}(\lambda x\, s_0))\vec{w}\} \\
&\sqsupseteq \{\mathsf{p}\big(\lambda z\, \mathsf{F}_i^{\pi_Q} \bot \vec{u}\vec{v}\big(t_0[\mathsf{F}_q^{\pi_Q} \bot \vec{u} z \vec{v}\vec{w}/x]\big)\vec{w}\big), \mathsf{p}\big(\lambda z\, \mathsf{F}_i^{\pi_Q} \bot \vec{u}\vec{v}\big(s_0[\mathsf{F}_q^{\pi_Q} \bot \vec{u} z \vec{v}\vec{w}/x]\big)\vec{w}\big)\} \\
&\sqsupseteq \{\mathsf{p}\big(\lambda z\, \mathsf{F}_i^{\pi_{Q_L}} \bot \vec{u}^*_L \vec{v}^*_L \big(t_0[\mathsf{F}_q^{\pi_{Q_R}} \bot \vec{u}^*_R z \vec{v}^*_R \vec{w}^*_R/x]\big)\vec{w}^*_L\big), \\
&\qquad\qquad \mathsf{p}\big(\lambda z\, \mathsf{F}_i^{\pi_{Q_R}} \bot \vec{u}^*_R \vec{v}^*_R \big(s_0[\mathsf{F}_q^{\pi_{Q_L}} \bot \vec{u}^*_L z \vec{v}^*_L \vec{w}^*_L/x]\big)\vec{w}^*_R\big)\}
\end{aligned}
$$

Now set

$$
\begin{aligned}
t &:= \mathsf{F}_i^{\pi_{Q_L}} \bot \vec{u}^*_L \vec{v}^*_L (t_0[\mathsf{F}_q^{\pi_{Q_R}} \bot \vec{u}^*_R z \vec{v}^*_R \vec{w}^*_R/x])\vec{w}^*_L \\
s &:= \mathsf{F}_i^{\pi_{Q_R}} \bot \vec{u}^*_R \vec{v}^*_R (s_0[\mathsf{F}_q^{\pi_{Q_L}} \bot \vec{u}^*_L z \vec{v}^*_L \vec{w}^*_L/x])\vec{w}^*_R
\end{aligned}
$$

We then have:

$$
\begin{aligned}
t[I^*(Q_R,i)/z] &= \mathsf{F}_i^{\pi_{Q_L}} \bot \vec{u}^*_L \vec{v}^*_L (t_0[\mathsf{F}_q^{\pi_{Q_R}} \bot \vec{u}^*_R (I^*(Q_R,i)) \vec{v}^*_R \vec{w}^*_R/x]) \\
&= \mathsf{F}_i^{\pi_{Q_L}} \bot \vec{u}^*_L \vec{v}^*_L (t_0[O^*(Q_R,q)/x]) \\
&= \mathsf{F}_i^{\pi_{Q_L}} \bot \vec{u}^*_L \vec{v}^*_L (t_0[I^*(R_1,p+1)/x]) \\
&= \mathsf{F}_i^{\pi_{Q_L}} \bot \vec{u}^*_L \vec{v}^*_L (O^*(R_1,p)) \\
&= \mathsf{F}_i^{\pi_{Q_L}} \bot \vec{u}^*_L \vec{v}^*_L (I^*(Q_L,q)) \\
&= O^*(Q_L,q)
\end{aligned}
$$

Similarly, we get $s[I^*(Q_L,i)/z] = O^*(Q_R,q)$. □

Claim 6. *If* $(R_0,p) \prec^+ (S,i)$ *then* $O(S,i) \sqsupseteq \{\mathsf{p}(\lambda x\, t), \mathsf{p}(\lambda x\, s)\}$ *for some terms* t, s, *such that:*

$$O^*(S_L,i) = t[I^*(S_R,i)/x] \text{ and } O^*(S_R,i) = s[I^*(S_L,i)/x].$$

Proof. By induction on the length of the path from S to Q. For the base case we have $S = Q$, and this is established by Claim 5.

For the induction step, suppose the length of the path from S to Q is $n+1$. Then there is some U on this path such that the length of the path from U to Q is n, and formula occurrences (U,k) and (S,j) such that $(U,k) \bigtriangledown (S,j)$. Since $(R_0,p) \prec^+ (S,i)$, by Proposition 16 we must have $i \neq \mathsf{link}(S,R_0) = j$ and $(R_0,p) \prec^+ (U,k)$ (since $k = \mathsf{link}(U,S)$ hence $k \neq \mathsf{link}(U,R_0)$). By the induction hypothesis, there are terms t_0, s_0 such that $O(U,k) \sqsupseteq \{\mathsf{p}(\lambda x\, t_0), \mathsf{p}(\lambda x\, s_0)\}$, $O^*(U_L,k) = t_0[I^*(U_R,k)/x]$ and

$O^*(U_R, k) = s_0[I^*(U_L, k)/x]$. We assume $i < j$ since the other case is similar, and write $O(S, i) = \mathsf{F}_i^{\pi_S} \vec{u}\vec{v}(O(U, k))\vec{w}$ where $\vec{u}$ are the input terms for indices below i. We define the tuples of terms $\vec{u}^*_L, \vec{v}^*_L, \vec{w}^*_L$ and $\vec{u}^*_R, \vec{v}^*_R, \vec{w}^*_R$, by appeal to Claim 3, as in the proof of Claim 5. We have:

$$\begin{aligned}
O(S,i) &= \mathsf{F}_i^{\pi_S}\vec{u}\vec{v}(O(U,j))\vec{w} \\
&\sqsupseteq \{\mathsf{F}_i^{\pi_S}\vec{u}\vec{v}(\mathtt{p}(\lambda x\, t_0))\vec{w}, \mathsf{F}_i^{\pi_S}\vec{u}\vec{v}(\mathtt{p}(\lambda x\, s_0))\vec{w}\} \\
&\sqsupseteq \{\mathtt{p}\big(\lambda z\, \mathsf{F}_i^{\pi_S}\vec{u}\vec{v}(t_0[\mathsf{F}_j^{\pi_S}\vec{u}z\vec{v}\vec{w}/x])\vec{w}\big), \mathtt{p}\big(\lambda z\, \mathsf{F}_i^{\pi_S}\vec{u}\vec{v}(s_0[\mathsf{F}_j^{\pi_S}\vec{u}z\vec{v}\vec{w}/x])\vec{w}\big)\} \\
&\sqsupseteq \{\mathtt{p}\big(\lambda z\, \mathsf{F}_i^{\pi_{S_L}}\vec{u}^*_L\vec{v}^*_L(t_0[\mathsf{F}_j^{\pi_{S_R}}\vec{u}^*_R z\vec{v}^*_R\vec{w}^*_R/x])\vec{w}_L\big), \\
&\qquad \mathtt{p}\big(\lambda z\, \mathsf{F}_i^{\pi_{S_R}}\vec{u}^*_R\vec{v}^*_R(s_0[\mathsf{F}_j^{\pi_{S_L}}\vec{u}^*_L z\vec{v}^*_L\vec{w}^*_L/x])\vec{w}^*_R\big)\}
\end{aligned}$$

Set

$$\begin{aligned}
t &:= \mathsf{F}_i^{\pi_{S_L}}\vec{u}^*_L\vec{v}^*_L(t_0[\mathsf{F}_j^{\pi_{S_R}}\vec{u}^*_R z\vec{v}^*_R\vec{w}^*_R/x])\vec{w}_L \\
s &:= \mathsf{F}_i^{\pi_{S_R}}\vec{u}^*_R\vec{v}^*_R(s_0[\mathsf{F}_j^{\pi_{S_L}}\vec{u}^*_L z\vec{v}^*_L\vec{w}^*_L/x])\vec{w}^*_R
\end{aligned}$$

We have:

$$\begin{aligned}
t[I^*(S_R,i)/z] &= \mathsf{F}_i^{\pi_{S_L}}\vec{u}^*_L\vec{v}^*_L(t_0[\mathsf{F}_j^{\pi_{S_R}}\vec{u}^*_R(I^*(S_R,i))\vec{v}^*_R\vec{w}^*_R/x])\vec{w}_L \\
&= \mathsf{F}_i^{\pi_{S_L}}\vec{u}^*_L\vec{v}^*_L(t_0[O^*(S_R,j)/x])\vec{w}_L \\
&= \mathsf{F}_i^{\pi_{S_L}}\vec{u}^*_L\vec{v}^*_L(t_0[I^*(U_R,k)/x])\vec{w}_L \\
&= \mathsf{F}_i^{\pi_{S_L}}\vec{u}^*_L\vec{v}^*_L(O^*(U_L,k))\vec{w}_L \\
&= \mathsf{F}_i^{\pi_{S_L}}\vec{u}^*_L\vec{v}^*_L(I^*(S_L,j))\vec{w}_L \\
&= O^*(S_L,i)
\end{aligned}$$

Similarly, we have $s[I^*(S_L, i)/z] = O^*(S_R, i)$. □

We can now prove the proposition. Suppose (S, i) is a side formula occurrence of the original multicut. If S is unaffected by the reduction then $O(S, i) \sqsupseteq O^*(S, i)$ by Claim 3, hence $\mathsf{E}(O(S, i)) \sqsupseteq \mathsf{E}(O^*(S, i))$. The case where $S = R_0$ is handled by Claim 3. If S is affected by the reduction and $S \neq R_0$, then it follows from Claim 2 that there is some cut formula occurrence (S, j) such that $(R_0, p) \prec^+ (S, k)$ for all $k \neq j$. But $i \neq j$ since (S, i) was a side formula occurrence, hence $(R_0, p) \prec^+ (S, i)$. It therefore follows from Claim 6 that $O(S, i) \sqsupseteq \{\mathtt{p}(\lambda x.t), \mathtt{p}(\lambda x.s)\}$ for some terms t, s such that we have $O^*(S_L, i) = t[I^*(S_L, i)/x]$ and $O^*(S_R, i) = s[I^*(S_R, i)/x]$. But since (S, i) was a side formula occurrence, $I^*(S_L, i) = I^*(S_R, i) = \mathsf{C}_{S[i]}$ or $I^*(S_L, i) = I^*(S_R, i) = \mathsf{C}_{\neg S[i]}$. So we get:

$$\begin{aligned}
\mathsf{E}(O(S,i)) &\sqsupseteq \{\mathsf{E}(\mathtt{p}(\lambda x.t)), \mathsf{E}(\mathtt{p}\lambda x.s))\} \\
&\sqsupseteq \{\mathsf{E}(t[\mathsf{C}_{S[i]}/x]), \mathsf{E}(s[\mathsf{C}_{S[i]}/x])\} \\
&= \{\mathsf{E}(t[I^*(S_L,i)/x]), \mathsf{E}(s[I^*(S_R,i)/x])\} \\
&= \{\mathsf{E}(O^*(S_L,i)), \mathsf{E}(O^*(S_R,i))\}
\end{aligned}$$

if $I^*(S_L, i) = I^*(S_R, i) = \mathsf{C}_{S[i]}$. The case where $I^*(S_L, i) = I^*(S_R, i) = \mathsf{C}_{\neg S[i]}$ is identical. This completes the proof.

A.5.3 Existential case

Let $(\mathbb{S}, \triangledown, \pi) \Rightarrow (\mathbb{S}^*, \triangledown^*, \pi^*)$ result from the internal reduction for the existential quantifier, per page 43. We prove Proposition 27:

Proposition. *Let (S, i) be any side formula occurrence with $S \in \mathbb{S}$.*

1. *If $S = L_0$ then $O(L_0, i) \sqsupseteq O^*(L_1, i)$.*
2. *If $S = R_0$ then $O(R_0, i) \sqsupseteq O^* R_1, i)$.*
3. *If $S \notin \{L_0, R_0\}$ then $O(S, i) \sqsupseteq O^*(S, i)$.*

Let (S, i) be any formula occurrence, either a side formula or a cut formula. We prove the following items by a simultaneous well-founded induction on the dependency relation $\prec$:

1. If $S = L_0$ and $i \neq l$ then $O(L_0, i) \sqsupseteq O^*(L_1, i)$.
2. If $S = R_0$ and $i \neq r$ then $O(R_0, i) \sqsupseteq O^*(R_1, i)$.
3. If $S \notin \{L_0, R_0\}$ then $O(S, i) \sqsupseteq O^*(S, i)$.

The proposition follows immediately from these statements.

For item (1), we write $O(L_0, i) = \mathsf{F}_i^{\pi_{L_0}} \bot \vec{u}(\mathsf{F}_r^{\pi_{R_0}} \bot \vec{w})\vec{v}$ where $\vec{u}$ are the inputs for indices below l. The input for each index $j \neq l$ is either of the form $\mathsf{C}_{L_0[j]}$, or $\mathsf{C}_{\neg L_0[j]}$, or of the form $O(U, k)$ where $(L_0, j) \triangledown (U, k)$. In the latter case, $U \notin \{L_0, R_0\}$ and item (3) of the induction hypothesis gives $O(U, k) \sqsupseteq O^*(U, k)$. Similarly, in the term $\mathsf{F}_r^{\pi_{R_0}} \bot \vec{w}$, each input for an index $j \neq r$ is either of the form $\mathsf{C}_{R_0[j]}$, or $\mathsf{C}_{\neg R_0[j]}$, or of the form $O(U, k)$ for some $(U, k) \triangledown (R_0, r)$. In the latter case we have $U \notin \{L_0, R_0\}$ and item (3) of the induction hypothesis gives $O(U, k) \sqsupseteq O^*(U, k)$. Let $\vec{u}^*, \vec{v}^*, \vec{w}^*$ be the result of replacing each of these terms of the form $O(U, k)$ by $O^*(U, k)$. We then have:

$$
\begin{aligned}
O(L_0, i) &= \mathsf{F}_i^{\pi_{L_0}} \bot \vec{u}(\mathsf{F}_r^{\pi_{R_0}} \bot \vec{w})\vec{v} \\
&\sqsupseteq \mathsf{F}_i^{\pi_{L_0}} \bot \vec{u}(\mathsf{e}(t \cdot \bot, \mathsf{F}_r^{\pi_{R_1}} \bot \vec{w})\vec{v} \\
&\sqsupseteq \mathsf{F}_i^{\pi_{L_1}} \bot [t \cdot \bot / \alpha]\vec{u}(\mathsf{F}_r^{\pi_{R_1}} \bot \vec{w})\vec{v} \\
&\equiv \mathsf{F}_i^{\pi_{L_1}[t/\alpha]} \bot \vec{u}(\mathsf{F}_r^{\pi_{R_1}} \bot \vec{w})\vec{v} && \text{Proposition 15, } \mathsf{Val}(t \cdot \bot) = t \\
&\sqsupseteq \mathsf{F}_i^{\pi_{L_1}[t/\alpha]} \bot \vec{u}^*(\mathsf{F}_r^{\pi_{R_1}} \bot \vec{w}^*)\vec{v}^* \\
&= O^*(L_1, i)
\end{aligned}
$$

as required.

For item (2), we write $O(R_0, i) = \mathsf{F}_i^{\pi_{R_0}} \bot \vec{u}(\mathsf{F}_l^{\pi_{L_0}} \bot \vec{w})\vec{v}$ where $\vec{u}$ are the inputs for indices below r. The input for each index $j \neq r$ is either of the form $\mathsf{C}_{L_0[j]}$, or $\mathsf{C}_{\neg L_0[j]}$, or of the form $O(U, k)$ where $(R_0, j) \triangledown (U, k)$. In the latter case, $U \notin \{L_0, R_0\}$ and item (3) of the induction hypothesis gives $O(U, k) \sqsupseteq O^*(U, k)$. Similarly, in the term $\mathsf{F}_l^{\pi_{L_0}} \bot \vec{w}$, each input for an index $j \neq l$ is either of the form $\mathsf{C}_{R_0[j]}$, or $\mathsf{C}_{\neg R_0[j]}$, or of the form $O(U, k)$ for some $(U, k) \triangledown (L_0, l)$. In the latter case we have $U \notin \{L_0, R_0\}$ and item (3) of the induction hypothesis gives $O(U, k) \sqsupseteq O^*(U, k)$. Let $\vec{u}^*, \vec{v}^*, \vec{w}^*$ be the result of replacing each of these terms of the form $O(U,k)$ by $O^*(U,k)$. We then have:

$$\begin{aligned}
O(R_0,i) &= \mathsf{F}_i^{\pi_{R_0}}\bot\vec{u}(\mathsf{F}_l^{\pi_{L_0}}\bot\vec{w})\vec{v} \\
&\sqsupseteq \mathsf{F}_i^{\pi_{R_0}}\bot\vec{u}(\mathsf{a}(\lambda x.\mathsf{F}_l^{\pi_{L_1}}\bot[x/\alpha]\vec{w}))\vec{v} \\
&\sqsupseteq \mathsf{F}_i^{\pi_{R_1}}\bot\vec{u}(\mathsf{F}_l^{\pi_{L_1}}\bot[t\cdot\bot/\alpha]\vec{w})\vec{v} \\
&\equiv \mathsf{F}_i^{\pi_{R_1}}\bot\vec{u}(\mathsf{F}_l^{\pi_{L_1}[t/\alpha]}\bot\vec{w})\vec{v} \\
&\sqsupseteq \mathsf{F}_i^{\pi_{R_1}}\bot\vec{u}^*(\mathsf{F}_l^{\pi_{L_1}[t/\alpha]}\bot\vec{w}^*)\vec{v}^* \\
&= O^*(R_1,i)
\end{aligned}$$

as required.

For item (3), suppose $S \notin \{L_0, R_0\}$ and write $O(S,i) = \mathsf{F}_i^{\pi_S}\bot\vec{u}$. Each input for and index $j \neq i$ is either $\mathsf{C}_{S[j]}$, or $\mathsf{C}_{\neg S[j]}$, or of the form $O(U,k)$ for some $(U,k) \bigtriangledown (S,i)$. In the latter case, there are three possibilities. If $U \notin \{L_0, R_0\}$ then item (3) of the induction hypothesis gives $O(U,k) \sqsupseteq O^*(U,k)$. If $U = L_0$ then we must have $k \neq l$ since otherwise we would have $(U,k) \bigtriangledown (S,j)$, hence $(S,j) = (R_0, r)$, contradicting our assumption that $S \notin \{L_0, R_0\}$. Item (1) of the induction hypothesis gives $O(L_0,j) \sqsupseteq O^*(L_1,j)$. Similarly, if $U = R_0$ then we must have $k \neq r$ and item (2) of the induction hypothesis gives $O(R_0,k) \sqsupseteq O^*(R_1,k)$. Let $\vec{u}^*$ be the result of replacing each input term (U,k) for $U \notin \{L_0,R_0\}$ by $O^*(U,k)$, each input term of the form (L_0,k) by $O^*(L_1,k)$ and each input term of the form (R_0,k) by $O^*(R_1,k)$. Then we have:

$$\begin{aligned}
O(S,i) &= \mathsf{F}_i^{\pi_S}\bot\vec{u} \\
&\sqsupseteq \mathsf{F}_i^{\pi_S}\bot\vec{u}^* \\
&= O^*(S,i)
\end{aligned}$$

as required.

A.5.4 Disjunction case

Let $(\mathbb{S}, \bigtriangledown, \pi) \Rightarrow (\mathbb{S}^*, \bigtriangledown^*, \pi^*)$ result from the internal existential reduction on page 43. We prove Proposition 28:

Proposition. *For each side formula occurrence (S,i), we have:*

1. *If $S = L_0$ then $O(L_0,i) \sqsupseteq O^*(L_1,i)$.*
2. *If $S = R_0$ then $O(R_0,i) \sqsupseteq O^*(R_1,i)$.*
3. *If $S \notin \{L_0,R_0\}$ then $O(S,i) \sqsupseteq O^*(S,i)$.*

We prove the following items by simultaneous well-founded induction on the relation $\prec$, for any given formula occurrence (S,i):

1. If $S = L_0$ and $i \neq l$ then $O(L_0,i) \sqsupseteq O^*(L_1,i)$.
2. If $S = R_0$ and $i \neq r$ then $O(R_0,i) \sqsupseteq O^*(R_1,i)$.
3. If $S \notin \{L_0,R_0\}$ then $O(S,i) \sqsupseteq O^*(S,i)$.

For item (1), write $O(L_0,i) = \mathsf{F}_i^{\pi_{L_0}}\bot\vec{u}(\mathsf{F}_r^{\pi_{R_0}}\bot\vec{w})\vec{v}$. For each index $j \notin \{l,i\}$, $I(L_0,j)$ is either of the form $\mathsf{C}_{L_0[j]}$, or $\mathsf{C}_{\neg L_0[j]}$, or of the form $O(U,k)$ for some $(U,k) \bigtriangledown (L_0,j)$.

In the latter case $U \notin \{L_0, R_0\}$ so item (3) of the induction hypothesis gives $O(U,k) \sqsupseteq O^*(U,k)$. Similarly, for each index $j \neq r$, $I(R_0, j)$ is either of the form $\mathsf{C}_{R_0[j]}$, or $\mathsf{C}_{\neg R_0[j]}$, or of the form $O(U,k)$ for some $(U,k) \bigtriangledown (R_0, j)$. In the latter case $U \notin \{L_0, R_0\}$ so item (3) of the induction hypothesis gives $O(U,k) \sqsupseteq O^*(U,k)$. Let $\vec{u}^*, \vec{v}^*, \vec{w}^*$ be the result of replacing each input term of the form $O(U,k)$ by $O^*(U,k)$. We have:

$$\begin{aligned} O(L_0, i) &= \mathsf{F}_i^{\pi_{L_0}} \bot \vec{u}(\mathsf{F}_r^{\pi_{R_0}} \bot \vec{w})\vec{v} \\ &\sqsupseteq \mathsf{F}_i^{\pi_{L_0}} \bot \vec{u}(\mathtt{i}_0(\mathsf{F}_r^{\pi_{R_1}} \bot \vec{w}))\vec{v} \\ &\sqsupseteq \mathsf{F}_i^{\pi_{L_1}} \bot \vec{u}(\mathsf{F}_r^{\pi_{R_1}} \bot \vec{w})\vec{v} \\ &\sqsupseteq \mathsf{F}_i^{\pi_{L_1}} \bot \vec{u}^*(\mathsf{F}_r^{\pi_{R_1}} \bot \vec{w}^*)\vec{v}^* \\ &= O^*(L_1, i) \end{aligned}$$

as required.

For item (2), write $O(R_0, i) = \mathsf{F}_i^{\pi_{R_0}} \bot \vec{u}(\mathsf{F}_l^{\pi_{L_0}} \bot \vec{w})\vec{v}$. For each index $j \notin \{r, i\}$, $I(R_0, j)$ is either of the form $\mathsf{C}_{R_0[j]}$, or $\mathsf{C}_{\neg R_0[j]}$, or of the form $O(U,k)$ for some $(U,k) \bigtriangledown (R_0, j)$. In the latter case $U \notin \{L_0, R_0\}$ so item (3) of the induction hypothesis gives $O(U,k) \sqsupseteq O^*(U,k)$. Similarly, for each index $j \neq l$, $I(L_0, j)$ is either of the form $\mathsf{C}_{L_0[j]}$, or $\mathsf{C}_{\neg L_0[j]}$, or of the form $O(U,k)$ for some $(U,k) \bigtriangledown (L_0, j)$. In the latter case $U \notin \{L_0, R_0\}$ so item (3) of the induction hypothesis gives $O(U,k) \sqsupseteq O^*(U,k)$. Let $\vec{u}^*, \vec{v}^*, \vec{w}^*$ be the result of replacing each input term of the form $O(U,k)$ by $O^*(U,k)$. We have:

$$\begin{aligned} O(R_0, i) &= \mathsf{F}_i^{\pi_{R_0}} \bot \vec{u}(\mathsf{F}_l^{\pi_{L_0}} \bot \vec{w})\vec{v} \\ &\sqsupseteq \mathsf{F}_i^{\pi_{R_0}} \bot \vec{u}(\mathtt{j}(\mathsf{F}_l^{\pi_{L_1}} \bot \vec{w}, \mathsf{F}_l^{\pi_{L_2}} \bot \vec{w}))\vec{v} \\ &\sqsupseteq \mathsf{F}_i^{\pi_{R_1}} \bot \vec{u}(\mathsf{F}_l^{\pi_{L_1}} \bot \vec{w})\vec{v} \\ &\sqsupseteq \mathsf{F}_i^{\pi_{R_1}} \bot \vec{u}^*(\mathsf{F}_l^{\pi_{L_1}} \bot \vec{w}^*)\vec{v}^* \\ &= O^*(R_1, i) \end{aligned}$$

as required.

The argument for item (3) is exactly as in the case of $\exists$-reduction so we omit it.

References

[1] Afshari, B., Hetzl, S., Leigh, G.E.: Herbrand's theorem as higher order recursion. Ann. Pure Appl. Log. **171**(6), 102792 (2020) https://doi.org/10.1016/j.apal.2020.102792

[2] Hetzl, S.: Applying tree languages in proof theory. In: Dediu, A., Martín-Vide, C. (eds.) Language and Automata Theory and Applications - 6th International Conference, LATA 2012, ACoruña, Spain, March 5-9, 2012. Proceedings. Lecture Notes in Computer Science, vol. 7183, pp. 301–312. Springer, Berlin (2012). https://doi.org/10.1007/978-3-642-28332-1_26

[3] Afshari, B., Hetzl, S., Leigh, G.E.: Herbrand disjunctions, cut elimination and context-free tree grammars. In: Altenkirch, T. (ed.) 13th International Conference on Typed Lambda Calculi and Applications, TLCA 2015, July 1-3, 2015, Warsaw, Poland. LIPIcs, vol. 38, pp. 1–16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, Dagstuhl (2015). https://doi.org/10.4230/LIPIcs.TLCA.2015.1

[4] Constable, R., Murthy, C.: Finding computational content in classical proofs, pp. 341–362. Cambridge University Press, USA (1991)

[5] Downen, P., Ariola, Z.M.: A tutorial on computational classical logic and the sequent calculus. Journal of Functional Programming **28**, 3 (2018) https://doi.org/10.1017/S0956796818000023

[6] Coquand, T.: Computational content of classical logic. In: Pitts, A.M., Dybjer, P.E. (eds.) Semantics and Logics of Computation. Publications of the Newton Institute, pp. 33–78. Cambridge University Press, Cambridge (1997). https://doi.org/10.1017/CBO9780511526619.003

[7] Powell, T.: Computational interpretations of classical reasoning: From the Epsilon Calculus to stateful programs. In: Centrone, S., Negri, S., Sarikaya, D., Schuster, P.M. (eds.) Mathesis Universalis, Computability and Proof, pp. 255–290. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-20447-1_14

[8] Kohlenbach, U.: Applied Proof Theory: Proof Interpretations and Their Use in Mathematics. Springer monographs in mathematics. Springer, Berlin (2008)

[9] Krivine, J.-L.: Realizability in classical logic. Panoramas & Synthéses, 199–231 (2009)

[10] Plotkin, G.D.: Call-by-name, call-by-value and the λ-calculus. Theoretical Computer Science **1**(2), 125–159 (1975) https://doi.org/10.1016/0304-3975(75)90017-1

[11] Avigad, J., Zach, R.: The Epsilon Calculus. In: Zalta, E.N. (ed.) The Stanford Encyclopedia of Philosophy, Fall 2020 edn. Metaphysics Research Lab, Stanford

Univesity (2020)

[12] Alcolei, A., Clairambault, P., Hyland, M., Winskel, G.: The true concurrency of herbrand's theorem. In: 27th EACSL Annual Conference on Computer Science Logic (CSL 2018) (2018). Schloss-Dagstuhl-Leibniz Zentrum für Informatik

[13] Troelstra, A.S., Schwichtenberg, H.: Basic Proof Theory, 2nd edn. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, Cambridge (2000). https://doi.org/10.1017/CBO9781139168717

[14] Hague, M., Murawski, A.S., Ong, C.-H., Serre, O.: Collapsible pushdown automata and recursion schemes. In: 2008 23rd Annual IEEE Symposium on Logic in Computer Science, pp. 452–461 (2008)

[15] Saurin, A.: A linear perspective on cut-elimination for non-wellfounded sequent calculi with least and greatest fixed-points. In: Ramanayake, R., Urban, J. (eds.) Automated Reasoning with Analytic Tableaux and Related Methods, pp. 203–222. Springer, Cham (2023)