

Herbrand's theorem as higher order recursion

Bahareh Afshari^{a,b,*}, Stefan Hetzl^c, Graham E. Leigh^b

^a *Institute for Logic, Language and Computation, University of Amsterdam, PO Box 94242, 1090 GE Amsterdam, The Netherlands*

^b *Department of Philosophy, Linguistics and Theory of Science, University of Gothenburg, Box 200, 405 30 Göteborg, Sweden*

^c *Institut für Diskrete Mathematik und Geometrie, TU Wien, Wiedner Hauptstraße 8–10, 1040 Wien, Austria*



ARTICLE INFO

Article history:

Received 4 February 2018

Received in revised form 1 February 2020

Accepted 13 February 2020

Available online 19 February 2020

MSC:

03F05

03F07

03D05

03F30

Keywords:

Classical sequent calculus

Herbrand's theorem

Cut elimination

Higher order recursion schemes

Computational content

ABSTRACT

This article examines the computational content of the classical Gentzen sequent calculus. There are a number of well-known methods that extract computational content from first-order logic but applying these to the sequent calculus involves first translating proofs into other formalisms, Hilbert calculi or Natural Deduction for example. A direct approach which mirrors the symmetry inherent in sequent calculus has potential merits in relation to proof-theoretic considerations such as the (non-)confluence of cut elimination, the problem of cut introduction, proof compression and proof equivalence. Motivated by such applications, we provide a representation of sequent calculus proofs as higher order recursion schemes. Our approach associates to an LK proof π of $\Rightarrow \exists vF$, where F is quantifier free, an acyclic higher order recursion scheme \mathcal{H} with a finite language yielding a Herbrand disjunction for $\exists vF$. More generally, we show that the language of \mathcal{H} contains all Herbrand disjunctions computable from π via a broad range of cut elimination strategies.

© 2020 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

The property of being a valid first-order formula is intimately tied to the consideration of the ground, i.e., variable-free, instances of that formula. This connection is apparent in most, if not all, proofs of the completeness theorem which, in one way or another, rely on the construction of a term model. The phenomenon is plainly visible in Herbrand's theorem which states that a formula is valid if, and only if, there is a finite expansion (of existential quantifiers to disjunctions and universal quantifiers to conjunctions of instances) which is tautological. This feature of classical first-order logic is in contrast to both classical

* Corresponding author at: Institute for Logic, Language and Computation, University of Amsterdam, PO Box 94242, 1090 GE Amsterdam, The Netherlands.

E-mail addresses: bahareh.afshari@gu.se (B. Afshari), stefan.hetzl@tuwien.ac.at (S. Hetzl), graham.leigh@gu.se (G.E. Leigh).

second-order logic, whose standard semantics goes beyond the ground instances of a countable language, and intuitionistic first-order logic, which exhibits a more complicated interaction between quantifiers and propositional connectives.

Proof-theoretically, the use of instances of a formula naturally leads to analytic, cut-free, proofs. Gentzen's mid-sequent theorem makes the close connection between Herbrand expansions and cut-free proofs apparent. Taking this perspective on the cut-elimination theorem, and thereby keeping the well-known complexity bounds in mind, shows that, in essence, cut-elimination consists of the computation of a Herbrand expansion. One may ask, however, whether it is possible to compute Herbrand expansions in a more direct way, circumventing the cumbersome process of cut-elimination. There are a number of formalisms that do just that, the historically first being Hilbert's ε -calculus [25] (see [30] for a contemporary exposition of the ε -theorems in English). In [17], Gerhardy and Kohlenbach adapt Shoenfield's variant [38] of Gödel's functional interpretation [18,6] to a system of pure predicate logic. A recent adaptation of the functional interpretation, utilising star types to interpret contraction, is given by Ferreira and Ferreira [14]. Related to proof nets is the work of Heijltjes [19] and McKinley [28], and a similar approach, in the formalism of expansion trees [29], can be found in [5]. A different method with similar aims is cut-elimination by resolution [9].

The present work is motivated by two follow-up questions: *Which Herbrand expansions are implicit in (i.e., can be computed from) a sequent calculus proof with cut? What is a minimal amount of information needed to describe these expansions?* These questions are closely related to the issue of (non-)confluence of cut elimination. For instance, the number of distinct Herbrand expansions computed by Gentzen-style cut-elimination can be non-elementary in the size of the starting proof [7]. In other words, the choice of reduction strategy affects which Herbrand expansion is computed by cut elimination. In light of this the motivating questions become a matter of *representation*, namely how to express the Herbrand expansions embedded in a proof with cut while abstracting away the propositional structure of the original proof and avoiding direct computation of cut-free proofs.

In this article, we provide a representation of Herbrand's theorem as languages of higher order recursion schemes. Higher order recursion schemes are a generalisation of regular tree grammars (which correspond to order-0 recursion schemes) to finite types. They have their origin in Park's program schemes [34] and are widely used for verification of higher-order functional programs [33].

More specifically, sequents in a given proof with cut are interpreted as non-terminals whose production rules follow the local instantiation structure of quantifiers. The type of a given non-terminal is completely determined by the quantifier complexity of formulæ in the corresponding sequent in such a way that a sequent comprising $\Sigma_n \cup \Pi_n$ formulæ is represented by a non-terminal of order n . These types correspond closely to the types arising in Shoenfield's version of Gödel's functional interpretation. Concerning production rules, cut corresponds to composition of non-terminals and contraction gives rise to non-determinism.

Our representation of Herbrand's theorem is specifically tailored for the classical sequent calculus in the sense that it remains faithful to the non-deterministic process of computing Herbrand expansions via (reductive) cut elimination. In this respect, we believe the present work marks the first method of Herbrand extraction that operates directly on sequent calculus proofs. The framework of higher order recursion schemes opens the door to applying techniques and results from formal language theory directly to structural proof theory. An example of the latter is an upper bound on the size (and, therefore, number) of Herbrand expansions which can be obtained via a broad array of cut elimination strategies. On the other hand, our adaptation of Beckmann's theorem on the length of β -reductions for the simply-typed λ -calculus [11] to language bounds on acyclic higher order recursion schemes, may be of independent interest in formal language theory.

The main result of this article can be summarised as follows and was announced in [3].

Theorem 1.1. *Let F be a quantifier-free formula and π a first-order proof of $\exists \vec{v} F$ in which cut-formulæ are prenex Π_n or Σ_n . There exists an acyclic order n recursion scheme \mathcal{H} with language $L(\mathcal{H})$ such that:*

i) $\forall_{\vec{t} \in L(\mathcal{H})} F(\vec{t})$ is valid; ii) $|L(\mathcal{H})| \leq 2^{4|\pi|^3}$ where $|\pi|$ is the number of inference rules in π ; iii) $L(\mathcal{H})$ subsumes the Herbrand set extracted from any cut-free proof that can be obtained from π via a sequence of Gentzen-style cut reductions that always reduces to the weak (quantifier) side of a cut before the strong side.

Outline

We begin by rehearsing the sequent calculus and reductive cut elimination. Higher order recursion schemes are introduced in Section 3. Theorem 3.16 establishes the upper bound on the size of languages for acyclic recursion schemes by generalising Beckmann’s theorem on the length of reduction sequences for simply typed λ -calculus [11]. Section 4 concerns our representation of sequent calculus proofs as higher order recursion schemes. Connections with the functional interpretation of classical logic are discussed and a correspondence between the two type hierarchies established (Theorem 4.4). The upper bound on language size stated in Theorem 1.1 is proved in Theorem 4.15. Section 5 provides a case study in which we analyse the Herbrand scheme for a proof of the pigeonhole principle. In Section 6 we establish the technical machinery necessary to relate reductive cut elimination to derivations in Herbrand schemes. The analysis of languages of Herbrand schemes in relation to reductive cut elimination is undertaken in Section 7, from which Theorem 1.1 follows. The article concludes with a discussion of the results and potential extensions.

2. Sequent calculus for classical first-order logic

Terms and formulæ of first-order logic are defined as usual using the connectives \wedge, \vee and quantifiers \forall, \exists , as well as a selection of predicate and function symbols. We assume two sets of variable symbols, *free* variables, denoted α, β , etc., and *bound* variables, v, w , etc. Upper-case Roman letters, A, B , etc. denote formulæ and upper-case Greek letters Γ, Δ , etc. range over *sequents*, namely finite sequences of formulæ. We abbreviate by Γ, Δ the concatenation of Γ and Δ ; and Γ, A is shorthand for $\Gamma, \{A\}$. The *length* of a sequent Γ is denoted $|\Gamma|$. As the order of the formulæ in a sequent is often (though not always) unimportant, we will frequently identify sequents with (finite) multisets. We write \bar{A} to denote the dual of the formula A obtained by de Morgan laws. Given a sequence of variable symbols $\vec{v} = (v_0, \dots, v_{k-1})$ of length k , we write $\forall \vec{v} A$ and $\exists \vec{v} A$ as shorthand for $\forall v_0 \dots \forall v_{k-1} A$ (resp. $\exists v_0 \dots \exists v_{k-1} A$). If $\vec{t} = (t_0, \dots, t_{k-1})$ is a sequence of terms of the same length, $A(\vec{t}/\vec{v})$ is the formula obtained from A by replacing each v_i by the corresponding term t_i , where bound variables in A are renamed as necessary to avoid variable capture.

The following abbreviations will be used in later sections. For a formula A , we write A_{qf} to indicate that A is quantifier-free, and $u(A)$ (resp. $e(A)$) for the number of consecutive universal (existential) quantifiers in A before encountering an existential (universal) quantifier:

$$\begin{aligned} u(\forall v A) &= u(A) + 1 & e(\exists v A) &= e(A) + 1 \\ u(\exists v A) &= u(A_{qf}) = 0 & e(\forall v A) &= e(A_{qf}) = 0 \end{aligned}$$

For notational simplicity, we work in one-sided sequent calculus with explicit structural rules for weakening (w), contraction (c) and permutation (p), though the results presented apply equally to two-sided (so-called Gentzen-style) sequent calculi and either form of calculus without explicit structural rules. The axioms and rules of the calculus are laid out in Fig. 1. The quantifier introduction rules $\forall_{\vec{\alpha}}$ and $\exists_{\vec{r}}$ introduce a sequence of quantifiers in one application. Applications of $\forall_{\vec{\alpha}}$ are subject to an eigenvariable condition that if $\vec{\alpha} = (\alpha_0, \dots, \alpha_{k-1})$ then α_i does not occur in the sequent Γ, A for any $i < k$. In each inference rule, the formulæ which are explicitly mentioned in the premise(s) (usually the right-most formula in the sequent) are said to be *active* in the rules applied. For example, A and B are active in \wedge rule, both copies of A are active in contraction, and there are no active formulæ in the weakening rule. Active formulæ of cut

Axioms:	A, \bar{A}	for A quantifier-free
Inference rules:	$\vee \frac{\Gamma, A, B}{\Gamma, A \vee B}$	$\wedge \frac{\Gamma, A \quad \Delta, B}{\Gamma, \Delta, A \wedge B}$
	$\forall_{\vec{\alpha}} \frac{\Gamma, A(\vec{\alpha}/\vec{v})}{\Gamma, \forall \vec{v} A}$	$\exists_r \frac{\Gamma, A(\vec{r}/\vec{v})}{\Gamma, \exists \vec{v} A}$
	$\text{w} \frac{\Gamma}{\Gamma, A}$	$\text{c} \frac{\Gamma, A, A}{\Gamma, A}$
		$\text{p} \frac{\Gamma, B, A, \Delta}{\Gamma, A, B, \Delta}$
		$\text{cut} \frac{\Gamma, A \quad \Delta, \bar{A}}{\Gamma, \Delta}$

Fig. 1. Axioms and rules of sequent calculus.

are referred to as *cut formulæ*. We often leave the applications of the permutation rule implicit, writing, for instance,

$$\forall_{\vec{\alpha}} \frac{\Gamma, A(\vec{\alpha}/\vec{v}), \Delta}{\Gamma, \forall \vec{v} A, \Delta} \quad \text{cut} \frac{\Gamma, A, \Gamma' \quad \Delta, \bar{A}, \Delta'}{\Gamma, \Gamma', \Delta, \Delta'}$$

to abbreviate derivations

$$\begin{array}{c} \text{p}^* \frac{\Gamma, A(\vec{\alpha}/\vec{v}), \Delta}{\Gamma, \Delta, A(\vec{\alpha}/\vec{v})} \\ \forall_{\vec{\alpha}} \frac{\Gamma, \Delta, A(\vec{\alpha}/\vec{v})}{\Gamma, \Delta, \forall \vec{v} A} \\ \text{p}^* \frac{\Gamma, \Delta, \forall \vec{v} A}{\Gamma, \forall \vec{v} A, \Delta} \end{array} \quad \begin{array}{c} \text{p}^* \frac{\Gamma, A, \Gamma'}{\Gamma, \Gamma', A} \quad \text{p}^* \frac{\Delta, \bar{A}, \Delta'}{\Delta, \Delta', \bar{A}} \\ \text{cut} \frac{\Gamma, \Gamma', A \quad \Delta, \Delta', \bar{A}}{\Gamma, \Gamma', \Delta, \Delta'} \end{array}$$

where in each case p^* denotes a sequence of permutation inferences p , a notation we also extend to the other structural rules.

Definition 2.1. A *proof* is a finite tree labelled by sequents obtained from the axioms and rules of the calculus with the restriction that cuts apply to prenex formulæ only. Without loss of generality, we assume all proofs are *regular*, by which we mean:

1. every eigenvariable in the proof appears in exactly one $\forall_{\vec{\alpha}}$ inference in the proof and does not occur in any sequent outside the sub-proof of this inference,
2. if A appears as the active formula of a quantifier inference \forall (\exists) then $u(A) = 0$ (resp. $e(A) = 0$).

A proof that does not contain the rule *cut* is *cut-free*. A proof in which every cut formula is quantifier-free is called *quasi cut-free*.

We write $\pi \vdash \Gamma$ to express that π is a regular proof with Γ being the sequent appearing at the root of π . $\text{EV}(\pi)$ denotes the set of eigenvariables in a proof π , and for sequences $\vec{\alpha} = (\alpha_0, \dots, \alpha_{k-1})$ and $\vec{t} = (t_0, \dots, t_{k-1})$ of variable symbols and terms, $\pi^{(\vec{t}/\vec{\alpha})}$ is the result of replacing throughout the proof π each occurrence of the variable symbol α_i by the term t_i .

2.1. Cut reduction and normal forms

The standard cut reduction and cut permutation steps are given in Figs. 2 and 3. For the sake of a concise presentation, the axioms and rules are stated with implicit permutation in place. We assume all the proofs drawn in Figs. 2 and 3 are regular. Hence, in the case of contraction reduction where the sub-proof π_1 is duplicated it is assumed that the eigenvariables are renamed in the copy, which is emphasised by annotating




Fig. 2. One-step cut reduction rules.





Fig. 3. One-step cut permutation rules.

the sub-proof with an asterisk i.e. π_1^* . In the two reductions of Fig. 3, r represents an arbitrary unary or binary inference rule. An example of the binary inference permutation rule for $r = \text{cut}$ is



For proofs π and π' we write $\pi \rightsquigarrow \pi'$ to express that π' is obtained from π by application of a reduction or permutation rule to a sub-proof of π , and let \rightsquigarrow^* denote the reflexive transitive closure of \rightsquigarrow . If $\pi \rightsquigarrow \pi'$ then the reduced cut either no longer exists, is replaced by cuts on formulæ with either lower logical complexity or fewer applied contractions, or permuted to a subproof. In any given proof there may, however, be many cuts and eliminating one can (through duplicating a sub-proof) result in introducing several copies of other cuts. To obtain a cut-free proof, it is necessary to provide a (terminating) *cut elimination strategy* i.e. a procedure that given any proof $\pi \vdash \Gamma$ induces a sequence of cut reduction and permutation steps $\pi \rightsquigarrow^* \pi'$ such that $\pi' \vdash \Gamma$ and the rule cut is not used in π' .

Theorem 2.2 (*Gentzen's Hauptsatz*). *There is a cut elimination strategy that transforms any proof in first-order logic to a cut-free proof.*

There are many cut elimination strategies such as top-most reduction strategy or the elimination of the cut with highest logical complexity. Different strategies provide different cut-free proofs, commonly also referred to as *normal forms*. In fact, there exist proofs with infinitely many normal forms (see e.g. [41, Example 2.1.3]). We now turn to the relationship between cut elimination and Herbrand's theorem.

2.2. Herbrand's theorem and cut elimination

Herbrand's theorem is considered a classic result in proof theory. It can be thought of as reducing validity in first-order logic to validity in propositional logic. From the modern perspective it can also be seen as extracting computational content to first-order proofs. A simple case of the theorem is the following.

Theorem 2.3 (*Herbrand's theorem*). *A formula $\exists \vec{v} A_{qf}$ is valid if and only if there exists a finite set of sequences of terms $\{\vec{t}_0, \vec{t}_1, \dots, \vec{t}_k\}$ such that $\bigvee_{i=0}^k A_{qf}(\vec{t}_i/\vec{v})$ is valid.*

If a formula $\exists \vec{v} A_{qf}$ is valid then any set of terms $\{\vec{t}_0, \dots, \vec{t}_k\}$ that validate the disjunction $\bigvee_{i=0}^k A_{qf}(\vec{t}_i/\vec{v})$ is called a *Herbrand set*, and the disjunction itself a *Herbrand disjunction* for the formula.

Herbrand's theorem pre-dates Gentzen's Hauptsatz but the latter readily provides an instructive proof of the theorem: Suppose $\exists \vec{v} A_{qf}$ is valid and fix a quasi cut-free proof $\pi \vdash \exists \vec{v} A$. It is possible to permute the rules applied in π so that no quantifier inference occurs above a purely propositional rule (Gentzen's mid-sequent theorem [15]). Once the proof is partitioned into a *propositional part* and a *quantifier part*, the terms that validate the formula can be directly read off from the *mid-sequent*, the sequent separating the two parts.

Herbrand's original statement is much more general than that stated above and applies to any formula of first-order logic thanks to *Herbrandisation*, the dual notion of Skolemization. Given an arbitrary formula A , by introducing suitable constant and function symbols it is possible to remove universal quantifiers in A and obtain a Σ_1 prenex-formula which is equi-valid to A . Herbrandisation can also be applied to a proof of a sequent Γ transforming it to a proof of the Herbrandisation of Γ .

If a Herbrand set (disjunction) is obtained via cut elimination it is customary to refer to it as a Herbrand set (disjunction) of the proof. Note that these are not unique: different reduction strategies can lead to non-elementary many pairwise distinct Herbrand disjunctions [7]. For both computing and representing Herbrand disjunctions it is therefore desirable to bypass cut elimination. There has been a number of

successful approaches such as via Herbrand nets [28], proof forests [19], expansion trees with cut [5] and functional interpretation [17]. In the next section we introduce a fresh approach based on higher order recursion schemes. The aim of this approach is twofold. On the one hand, we wish to provide a representation of Herbrand's theorem specifically tailored for the classical sequent calculus which is faithful to the non-deterministic process of computing Herbrand expansions via (reductive) cut elimination. On the other hand, the framework of higher order recursion schemes opens the door to applying techniques and results from formal language theory directly to structural proof theory.

3. Recursion schemes

We begin this section introducing the type system and terms that will be used throughout the paper. In Sections 3.2 and 3.3, higher order recursion schemes for this type hierarchy are introduced. The upper bounds on language size we establish in Theorem 3.16 allow us to deduce the numerical bound claimed in Theorem 1.1. The association of proofs with recursion schemes is given in Section 4.

3.1. Types and terms

The type system we utilise extends the hierarchy of simple types (over a type of individuals ι) by pair types and two additional type constants. These are the unit type, denoted ϵ , and a type ς of (stacks of) substitutions, elements of which are finite sequences of pairs (α, r) where α and r are elements of some (and the same) type. We are interested specifically in the case that α is a constant symbol (of simple type) from a particular ranked alphabet Σ , and refer to the type ς as the type of *substitution stacks (over Σ)*, or simply Σ -*substitutions*.

The informal reading behind the type ς is that of an accumulator for a sequence of substitutions that are generated by reading a particular thread through a formal proof: when a witness to an existential quantifier is encountered along such a thread, the witness is outputted accompanied by the current stack of substitutions. The substitutions are not evaluated at the formal level but recorded as an element of ς .

We begin with a formal definition of the types and conventions for their representation, followed by ranked alphabets and the recursive definition of (typed) terms including the precise form of inhabitants of the type of substitution stacks.

Definition 3.1. The *types* are defined in the following way.

- ι is a type, called the type of *individuals*.
- ϵ is a type, called the *unit type*.
- ς is a type, called the type of *substitution stacks*.
- *Function types*: if ρ, σ are types then $\rho \rightarrow \sigma$ is a type.
- *Pair types*: if ρ, σ are types then $\rho \times \sigma$ is a type.

A type formed without reference to ς is called *basic*, and one formed only out of ι and \rightarrow is *simple*. The types ι and ϵ are referred to collectively as *ground types* and any type that is not a function type is called *prime*. The *sequence types* are the types of the form ι^n for any n , where $\iota^0 = \epsilon$ and $\iota^{n+1} = \iota \times \iota^n$. The set of all types is denoted Type .

We follow the convention that the two type forming operations \times and \rightarrow associate to the right, and that \rightarrow binds more strongly than \times , so for ρ_0, \dots, ρ_k types we have

$$\rho_0 \times \rho_1 \times \dots \times \rho_k = \rho_0 \times (\rho_1 \times \dots \times \rho_k)$$

$$\begin{aligned} \rho_0 \rightarrow \rho_1 \rightarrow \cdots \rightarrow \rho_k &= \rho_0 \rightarrow (\rho_1 \rightarrow \cdots \rightarrow \rho_k) \\ \rho_0 \times \cdots \times \rho_i \rightarrow \rho_{i+1} \times \cdots \times \rho_k &= (\rho_0 \times \cdots \times \rho_i) \rightarrow (\rho_{i+1} \times \cdots \times \rho_k) \end{aligned}$$

Every type ρ has a unique decomposition $\rho = \rho_1 \rightarrow \rho_2 \rightarrow \cdots \rightarrow \rho_k \rightarrow co(\rho)$ where $co(\rho)$ is a prime type. Given such a decomposition of ρ we refer to $co(\rho)$ as the *co-domain* of ρ , to k as the *arity* of ρ , and to ρ_i ($1 \leq i \leq k$) as the *i -th domain* of ρ . Note that the co-domain of a ground type or pair type is the type itself.

The *order* of a type ρ generalises the usual definition of order for simple types. Motivated by later technicalities however, it convenient to assume any function type whose co-domain is the unit type has order 0, which in our Herbrand schemes will always be observationally equivalent to a single constant function, has order 0.

Definition 3.2 (*Order*). The *order* of a type ρ , $ord(\rho)$, is defined as follows.

$$\begin{aligned} ord(\iota) = ord(\epsilon) = ord(\varsigma) &= 0 & ord(\rho \times \sigma) &= \max\{ord(\rho), ord(\tau)\} \\ ord(\rho \rightarrow \sigma) &= \begin{cases} 0, & \text{if } co(\sigma) = \epsilon, \\ \max\{ord(\rho) + 1, ord(\sigma)\}, & \text{otherwise.} \end{cases} \end{aligned}$$

Definition 3.3 (*Ranked alphabet*). A (ranked) *alphabet* is a pair $\mathcal{A} = \langle S, \lambda \rangle$ where S is a set, called the *carrier* of \mathcal{A} , and $\lambda: S \rightarrow \text{Type}$ is a type assignment for elements of S . If $\lambda(S)$ is a simple (basic) type for every $S \in S$ we call \mathcal{A} *simple* (resp. *basic*). Two ranked alphabets are *disjoint* just in case their carriers are disjoint sets.

Given an alphabet $\mathcal{A} = \langle S, \lambda \rangle$, we write $\alpha^\rho \in \mathcal{A}$ if $\alpha \in S$ and $\lambda(\alpha) = \rho$, and hence frequently identify \mathcal{A} with the set $\{\alpha^{\lambda(\alpha)} \mid \alpha \in S\}$ of symbols with type annotations. For alphabets $\mathcal{A} = \langle S, \lambda \rangle$ and $\mathcal{B} = \langle S', \lambda' \rangle$, we write $\mathcal{A} \subset \mathcal{B}$ if $S \subseteq S'$ and $\lambda = \lambda' \upharpoonright S$. In case \mathcal{A} and \mathcal{B} are disjoint, $\mathcal{A} \cup \mathcal{B}$ denotes the alphabet formed by the union of \mathcal{A} and \mathcal{B} , namely $\langle S \cup S', \lambda \cup \lambda' \rangle$. The empty alphabet is denoted \emptyset .

Definition 3.4 (*Terms and substitutions*). Fix alphabets $\Sigma \subset \mathcal{A}$ where Σ is simple. The *\mathcal{A} -terms over Σ* (henceforth *\mathcal{A} -terms*) and the types they inhabit are defined inductively as follows, where $r : \rho$ expresses that r is an \mathcal{A} -term of type ρ .

1. $\langle \rangle$ is an \mathcal{A} -term of type ϵ .
2. If $\alpha^\rho \in \mathcal{A}$ then α is an \mathcal{A} -term of type ρ .
3. If $r : \rho$ and $s : \sigma$ then $\langle r, s \rangle$ is a \mathcal{A} -term of type $\rho \times \sigma$.
4. If $r : \sigma \rightarrow \tau$ and $s : \sigma$ then rs is a \mathcal{A} -term of type τ .
5. \perp is an \mathcal{A} -term of type ς .
6. If $a : \varsigma$ and $r : \rho$, and $\alpha^\rho \in \Sigma$ then $[\alpha \leftarrow r]a$ is an \mathcal{A} -term of type ς .
7. If $r : \rho$, $a : \varsigma$ and ρ is basic then $r \cdot a$ is an \mathcal{A} -term of type ρ .

Note that λ -abstraction is not present in the term calculus, so the existence of terms of function type depends the \mathcal{A} -symbols.

In addition to the notation $r : \rho$ used above, we may also write r^ρ to express that r is an \mathcal{A} -term of type ρ . We drop mention of Σ and \mathcal{A} if they can be inferred from the context or are not important to the given setting, in which case \mathcal{A} -terms are referred to simply as *terms*. Terms arising from clauses 1, 2 and 5 are called *constants*; terms arising from 3 and 4 are called *pairs* and *applications* respectively; terms of type ς are called *substitution stacks*; and terms of the form in 7 are called *explicit substitutions* (or simply *substitutions* if there is no cause for confusion). A *basic term* is any term constructed via the rules 1 to 4

only, i.e. a \mathcal{B} -term for some basic alphabet \mathcal{B} . A term of a sequence type is called a *sequence*. Application is assumed to associate to the left, and pairing and the formation rule for substitution stacks both associate to the right.

The sub-term relation is defined as usual over the basic terms, and is extended to terms containing substitutions by defining the sub-terms of \perp to be $\{\perp\}$, the sub-terms of $a = [\alpha \leftarrow r]b$ to be a and any sub-term of r or b , and the sub-terms of $r = s \cdot a$ to be r and any sub-term of s or a . Thus the basic terms are precisely those terms that do not have a substitution stack as a sub-term.

Given a finite sequence of terms $(r_i : \rho_i)_{i \leq k}$, let $\langle r_0, r_1, \dots, r_k \rangle$ be the term r_0 if $k = 0$ and, otherwise, the pair $\langle r_0, \langle r_1, \dots, \langle r_{k-1}, r_k \rangle \dots \rangle$ of type $\rho_0 \times \rho_1 \times \dots \times \rho_k$. The *order* of a term is the order of its type.

Proposition 3.5. *If $\Sigma \subset \Sigma'$ are simple alphabets and \mathcal{A} is an alphabet extending Σ' , then every \mathcal{A} -term over Σ is an \mathcal{A} -term over Σ' .*

In addition to the term-level explicit substitutions, there is of course the usual operation of substituting given symbols by terms of corresponding type which we refer to as *implicit substitution*. Explicit substitutions can be interpreted as implicit substitutions by reading terms $r \cdot a$ as the image of r under the (implicit) substitution described by a , a process we call *evaluation*. The following definitions explicate these two operations. Fix an alphabet \mathcal{A} .

Definition 3.6 (*Implicit substitution*). For \mathcal{A} -terms $r : \rho$, $t_0 : \tau_0$, \dots , $t_k : \tau_k$ and distinct symbols $\alpha_0^{\tau_0}, \dots, \alpha_k^{\tau_k} \in \mathcal{A}$, the term $r(\vec{t}/\vec{\alpha})$ is the \mathcal{A} -term given by simultaneously replacing every occurrence of α_i (for $i \leq k$) in r by t_i , defined recursively by:

$$\beta(\vec{t}/\vec{\alpha}) = \begin{cases} \beta, & \text{if } \beta \in \mathcal{A} \text{ and } \beta \notin \{\alpha_i \mid i \leq k\}, \\ t_i, & \text{if } \beta = \alpha_i, \end{cases}$$

$$\langle \rangle(\vec{t}/\vec{\alpha}) = \langle \rangle \qquad (rs)(\vec{t}/\vec{\alpha}) = r(\vec{t}/\vec{\alpha})s(\vec{t}/\vec{\alpha})$$

$$\perp(\vec{t}/\vec{\alpha}) = \perp \qquad \langle r, s \rangle(\vec{t}/\vec{\alpha}) = \langle r(\vec{t}/\vec{\alpha}), s(\vec{t}/\vec{\alpha}) \rangle$$

$$([\beta \leftarrow s]a)(\vec{t}/\vec{\alpha}) = [\beta \leftarrow s(\vec{t}/\vec{\alpha})](a(\vec{t}/\vec{\alpha})) \qquad (r \cdot a)(\vec{t}/\vec{\alpha}) = r(\vec{t}/\vec{\alpha}) \cdot (a(\vec{t}/\vec{\alpha}))$$

If the choice of $\vec{\alpha}$ can be inferred from context, we write $r(\vec{t})$ in place of $r(\vec{t}/\vec{\alpha})$.

Definition 3.7 (*Evaluating substitutions*). Given an \mathcal{A} -term r and a substitution stack $a = [\alpha_1 \leftarrow s_1] \dots [\alpha_k \leftarrow s_k] \perp : \varsigma$ over some simple alphabet $\Sigma \subset \mathcal{A}$, the *evaluation of r relative to a* is the \mathcal{A} -term over Σ given by

$$r^a := r(s_1/\alpha_1) \dots (s_k/\alpha_k).$$

The *evaluation* of r is the term r° given by recursively evaluating relative to each substitution in r , namely evaluation leaves basic terms unchanged, commutes with application and pairing, is defined by $\perp^\circ = \perp$ and $([\alpha \leftarrow r]a)^\circ = [\alpha \leftarrow r^\circ]a^\circ$ on substitution stacks, and by $(r \cdot a)^\circ = (r^\circ)^{a^\circ}$ for explicit substitutions.

Note that the evaluation of a substitution stack on a term is well-defined due to the typing constraints on their formation.

An alphabet generally specifies a set of symbols which are associated certain re-write rules in a recursion scheme. In this context, an explicit substitution acts as a delayed substitution which is not evaluated until no further re-writes to sub-terms are possible. For instance, over the alphabet $\{F^{\iota \rightarrow \iota}, G^\iota, e^\iota, o^{\iota \rightarrow \iota \rightarrow \iota}, \alpha^\iota\}$ with associated re-write rules $Fx \rightarrow (G \circ x) \cdot [\alpha \leftarrow e] \perp$ (for any instantiation of x) and $G \rightarrow \alpha$, a derivation starting from the term $t = F\alpha$ is

$$t \rightarrow (\mathbf{G} \circ \alpha) \cdot ([\alpha \leftarrow \mathbf{e}] \perp) \rightarrow (\alpha \circ \alpha) \cdot ([\alpha \leftarrow \mathbf{e}] \perp).$$

The final term evaluates to $\mathbf{e} \circ \mathbf{e}$. Attempting to read the explicit substitution implicitly leads also to the derivation

$$t \rightarrow (\mathbf{G} \circ \alpha) \cdot ([\alpha \leftarrow \mathbf{e}] \perp) = \mathbf{G} \circ \mathbf{e} \rightarrow \alpha \circ \mathbf{e}.$$

Lemma 3.8. *If $r : \rho$ is a Σ -term for some simple alphabet Σ and ρ is a basic type then r° is a basic Σ -term of type ρ .*

Proof. All substitution stacks that may occur in a term of basic type built from an alphabet of simply-typed symbols must be within the context of an explicit substitution. As evaluation replaces every explicit substitution by an implicit one, the result is a basic term of the same type. \square

Lemma 3.9. *If r and $a = [\alpha \leftarrow s]b : \zeta$ are \mathcal{A} -terms such that α does not occur in r , then $r^a = r^b$.*

Definition 3.10 (Σ -length). Given alphabets $\Sigma \subset \mathcal{A}$ and an \mathcal{A} -term s , the Σ -length of s , written $|s|_\Sigma$, is the number of occurrences of symbols in s that are not Σ -terms, formally: $|s|_\Sigma = 0$ if $s \in \Sigma \cup \{\langle \rangle\}$; $|s|_\Sigma = 1$ if $s \in \mathcal{A}$ and $s \notin \Sigma$; and $|r|_\Sigma = |s|_\Sigma + |t|_\Sigma$ if $r \in \{\langle s, t \rangle, st, s \cdot t, [\alpha \leftarrow s]t\}$.

In particular, the Σ -length of a Σ -term is 0 and if Σ is the empty alphabet then Σ -length of any term is the number of leaves in the tree representation of the term not labelled by $\langle \rangle$.

Notational conventions Symbols ρ, σ and τ (also with indices) range over types. We commonly notate alphabets by upper-case Roman symbols in calligraphic typeface: \mathcal{A}, \mathcal{B} , etc., though Greek symbols Σ and Σ' will be used for simple alphabets. Sans-serif typeface (f, F, s, S, etc.) and lowercase Greek symbols α, β , etc. range over elements of ranked alphabets, with the latter particularly used for constants of simple type. In the Herbrand schemes we introduce in Section 4, the constant symbols of simple type will be precisely the eigenvariables of sequent calculus proofs, hence our use of the same symbols. Italicised letters r, s, t, R, S , etc. range over terms and a, b over substitution stacks, i.e. terms of type ζ .

3.2. Higher order recursion schemes

Higher order recursion schemes (HORS) are a generalisation of regular and context-free grammars to the simple type hierarchy. Their origin lies in Park's *program schemes* [34] from the late 1960s. More recently, HORS have found notable applications in the verification of higher-order functional programs [31,26] (see, e.g. [33], for an overview).

In this section we recall the notion of higher order recursion schemes, which we generalise to the type system introduced above. We establish bounds on the size of languages of acyclic HORS (Corollary 3.17), which will be utilised later to deduce upper bounds on the length of Herbrand disjunctions.

Definition 3.11 (*Higher order recursion scheme*). A (*non-deterministic*) *higher order recursion scheme*, or simply *recursion scheme*, is a tuple $\mathcal{R} = \langle \Sigma, \mathcal{N}, \mathcal{S}, \mathcal{P} \rangle$ where Σ is a simple alphabet, \mathcal{N} is a alphabet of *non-terminals* disjoint from Σ , $\mathcal{S} \subseteq \mathcal{N}$ is a designated finite set of *starting symbols* of sequence type, \mathcal{P} is a set of pairs (F^ρ, t) , called *production rules*, such that $F^\rho \in \mathcal{N}$ and $t : \text{co}(\rho)$ is a $(\Sigma \cup \mathcal{N} \cup \{x_1^{\rho_1}, \dots, x_k^{\rho_k}\})$ -term over Σ where x_i is a fresh symbol not in \mathcal{N} and ρ_i is the i -th domain of ρ . A production rule (F^ρ, t) where the arity of ρ is k is written as

$$F x_1 \cdots x_k \rightarrow_{\mathcal{R}} t,$$

or $F\vec{x} \rightarrow_{\mathcal{R}} t$. Notice that by definition the term $Fx_1 \cdots x_k$ is of type $co(\rho)$.

A non-terminal $F \in \mathcal{N}$ of \mathcal{R} is *determined* if there is a unique production rule (F^ρ, t) in \mathcal{P} . By an \mathcal{R} -term we mean a $(\Sigma \cup \mathcal{N})$ -term over Σ . The *order* of \mathcal{R} is the supremum over orders of the types of non-terminals of \mathcal{R} .

Notice that we do not require that \mathcal{R} contains only finitely many non-terminals, nor that the set of start symbols is non-empty. This is for technical convenience as it allows us to consider the recursion schemes of the next section as finitely generated ‘sub-schemes’ of a single infinite recursion scheme. Moreover, higher order recursion schemes are traditionally presented in the context of simple types, wherein start symbols are all of type ι (and indeed a single start symbol suffices) and production rules have the form $F\vec{x} \rightarrow t$ with $t : \iota$. The above definition is a direct extension of recursion schemes that accommodates non-trivial prime types.

A given non-terminal may be assigned multiple production rules, leading to non-determinism. To simplify presentation of production rules in this case we adopt the convention of writing

$$F\vec{x} \rightarrow_{\mathcal{R}} t_0 \mid \cdots \mid t_k$$

to express that \mathcal{R} contains exactly the production rule $F\vec{x} \rightarrow t_i$ for each $i \leq k$, i.e. $F\vec{x} \rightarrow_{\mathcal{R}} t_i$ for each $i \leq k$ and if $F\vec{x} \rightarrow_{\mathcal{R}} t$ then $t = t_i$ for some $i \leq k$.

Definition 3.12 (*Derivations and language*). Let $\mathcal{R} = \langle \Sigma, \mathcal{N}, \mathcal{S}, \mathcal{P} \rangle$ be a higher order recursion scheme. We extend the relation $\rightarrow_{\mathcal{R}}$ to a relation on \mathcal{R} -terms defined by setting $r \rightarrow_{\mathcal{R}} s$ if either

- $r = Fr_1 \cdots r_k$ for some $F^\rho \in \mathcal{N}$ with arity k and there exists a production rule $F\vec{x} \rightarrow_{\mathcal{R}} t$ such that $s = t(\vec{r}/\vec{x})$;
- $r = t(r_0/x)$, $s = t(s_0/x)$ and $r_0 \rightarrow_{\mathcal{R}} s_0$.

A *derivation* of s from r is a sequence $r = r_0 \rightarrow_{\mathcal{R}} \cdots \rightarrow_{\mathcal{R}} r_k = s$, the *length* of which is k . We say s is *derivable from* r in \mathcal{R} , in symbols $r \rightarrow_{\mathcal{R}}^* s$ (or $r \rightarrow^* s$ if \mathcal{R} is clear from the context), if there exists a derivation of s from r , and s is *derivable in* \mathcal{R} if s is derivable from some $S \in \mathcal{S}$. The *language* of \mathcal{R} , written $L(\mathcal{R})$, is the set of pairs (S, t) such that $S \in \mathcal{S}$, t is a basic Σ -term and $S \rightarrow_{\mathcal{R}}^* t$.

Definition 3.13. Let $\mathcal{R} = \langle \Sigma, \mathcal{N}, \mathcal{S}, \mathcal{P} \rangle$ be a higher order recursion scheme. \mathcal{R} is *finite* if \mathcal{N} and \mathcal{P} are both finite sets, and is *acyclic* if there exists a transitive, irreflexive relation $<$ on \mathcal{N} such that for every production rule $F\vec{x} \rightarrow_{\mathcal{R}} t$ and every non-terminal G occurring in t , $G < F$.

Lemma 3.14. *A finite acyclic recursion scheme induces a finite language.*

An upper bound on the size of the language of acyclic recursion schemes can be obtained by reducing the problem to the length of reduction sequences for the simply-typed λ -calculus. Bounds on normalisation in the simply-typed λ -calculus have been given by Schwichtenberg [36] and improved to exact bounds by Beckmann [11]. In the following we use Beckmann’s result to obtain concrete bounds for acyclic recursion schemes. Let $2_0^n = n$ and $2_{k+1}^n = 2^{2_k^n}$ and extend the length function of the previous section to include λ -abstractions by setting $|\lambda xs|_{\Sigma} = |s|_{\Sigma} + 1$.

Theorem 3.15 (*Beckmann [11]*). *Let t be a term in the simply-typed λ -calculus over a simple alphabet Σ . The length of any β -reduction sequence starting from t is bounded by $2_{d(t)}^{|t|_{\Sigma}}$ where $d(t)$ denotes the maximum order among sub-terms of t .*

Beckmann’s bound still applies if t is an arbitrary λ -term over the calculus of Σ -terms given in Definition 3.4 subject to the restriction that $\langle \rangle$ is the only Σ -term of type ϵ (a necessary restriction due to our requirement that $ord(\sigma \rightarrow \epsilon) = 0$). Non-deterministic reductions can also be incorporated via a fresh operator $|$ and permitting β -reductions of the form $(\lambda x. t_0 | \cdots | t_k) s \rightarrow_{\beta} t_i(s/x)$ for each $i \leq k$. In this case the length and the function d is given by $|s|_{\Sigma} = \max\{|s|_{\Sigma}, |t|_{\Sigma}\}$ and $d(s|t) = \max\{d(s), d(t)\}$. Finally, we wish to allow for so-called η -long reductions, i.e., reductions $(\lambda x_0 \cdots x_k. s) t_0 \cdots t_k \rightarrow_{\beta} s(\vec{t}/\vec{x})$ where s is not an abstraction. Provided that only η -long reductions are permitted and each counts as one step in a β -reduction sequence, Beckmann’s bound holds with the analogous change to the length function: $|\lambda \vec{x} s|_{\Sigma} = |s|_{\Sigma} + 1$ if s not an abstraction.

From these observations we may deduce the following result. We restrict ourselves to recursion schemes over basic types (i.e. without substitution stacks) as this will suffice for later use.

Theorem 3.16. *Let $\mathcal{R} = \langle \Sigma, \mathcal{N}, \mathcal{S}, \mathcal{P} \rangle$ be a finite acyclic order n recursion scheme such that every non-terminal has basic type, and for every production rule $F\vec{x} \rightarrow t$ in \mathcal{R} , $|t|_{\Sigma} < k$. The length of every derivation in \mathcal{R} is bounded by $2^{|N|(k+1)}_{n+1}$.*

Proof. Let $\mathcal{R} = \langle \Sigma, \mathcal{N}, \mathcal{S}, \mathcal{P} \rangle$ be an order n recursion scheme fulfilling the requirements in the statement. Without loss of generality we may assume that \mathcal{S} is a singleton, that every non-terminal is associated at least one production rule, and that \mathcal{X} is the alphabet of variable symbols disjoint from both Σ and \mathcal{N} such that every term occurring in a production rule in \mathcal{R} is an $(\Sigma \cup \mathcal{N} \cup \mathcal{X})$ -term.

Fix an enumeration $F_N^{\rho_N} < \cdots < F_2^{\rho_2} < F_1^{\rho_1}$ of the non-terminals of \mathcal{R} according to a total ordering $\langle \rangle$ witnessing acyclicity of \mathcal{R} . We may assume $\mathcal{S} = \{F_1\}$, so ρ_1 is prime. Let $\mathcal{Y} = \{y_1^{\rho_1}, y_2^{\rho_2}, \dots, y_N^{\rho_N}\}$ be a set of fresh variable symbols of marked type. We define by recursion a sequence s_1, \dots, s_N of well-typed λ -terms all of type ρ_1 such that s_i contains only the variables y_{i+1}, \dots, y_N free and the length of every derivation from F_1 which only re-writes non-terminals F_j for $j \leq i$ is bounded by the length of the longest β -reduction sequence starting from s_i . Suppose $\{F_i \vec{x} \rightarrow T_j : j \leq m\}$ is the set of production rules associated to F_i in \mathcal{R} . For each $j \leq m$, let $t_j = T_j(y_{i+1}, \dots, y_N / F_{i+1}, \dots, F_N)$ be the $(\Sigma \cup \mathcal{X} \cup \mathcal{Y})$ -term resulting from T_j by substituting the non-terminals F_{i+1}, \dots, F_N by variables y_{i+1}, \dots, y_N respectively. It follows that $|t_j|_{\Sigma} \leq |T_j|_{\Sigma} < k$. Finally, define $s_i = \lambda \vec{x}. t_0 | \cdots | t_m$ if $i = 1$, and $s_i = (\lambda y_i. s_{i-1})(\lambda \vec{x}. t_0 | \cdots | t_m)$ otherwise. Notice $|s_N|_{\Sigma} \leq N(k+1)$ and the maximal order among sub-terms of s_N is no greater than $n+1$. Every \mathcal{R} -derivation from F_1 can be replicated as a sequence of one-step η -long β -reductions starting from s_N , the length of which, by Beckmann’s bound, is no greater than $2^{|N|(k+1)}_{n+1}$. \square

As a corollary we obtain we obtain bounds on the size of languages.

Corollary 3.17. *Let \mathcal{R} and k be as in the previous theorem and suppose every non-terminal in \mathcal{R} is associated at most two production rules. Then the size of $L(\mathcal{R})$ is bounded by $2^{|N|(k+1)}_{n+2}$.*

Proof. Given a recursion scheme \mathcal{R} all terms in $L(\mathcal{R})$ can be derived via the leftmost reduction strategy. By the previous theorem, the length of these derivations is bounded by $2^{|N|(k+1)}_{n+1}$, leading to a bound of $2^{|N|(k+1)}_{n+2}$ on the size of $L(\mathcal{R})$. \square

The bound given in Corollary 3.17 is optimal in the parameter n as the next lemma demonstrates.

Lemma 3.18. *Let Σ be the ranked alphabet $\{a^t, b^t, d^{t \rightarrow t \rightarrow t}\}$. There exists a sequence of acyclic higher order recursion schemes $\mathcal{R}_n = \langle \Sigma, \mathcal{N}_n, \mathcal{S}_n, \mathcal{P}_n \rangle$ such that*

1. the order of \mathcal{R}_n is n ,
2. $|\mathcal{N}_n|, |\mathcal{P}_n| = O(n)$,

3. $\max\{|t|_{\Sigma} : F\vec{x} \rightarrow_{\mathcal{R}_n} t\} = O(n)$,
4. $|L(\mathcal{R}_n)| \geq 2^1_{n+2}$.

Proof. It suffices to translate Beckmann’s lower bounds from [11] to the context of recursion schemes. Define $\tau_0 = \iota$ and $\tau_{i+1} = \tau_i \rightarrow \tau_i$ for each $i < \omega$. So τ_i has order and arity i for each i . Fix $n > 0$. The recursion scheme \mathcal{R}_n comprises a single start symbol $S_n : \iota$ and a non-terminal $F_i : \tau_i$ for each $i \leq n$. The production rules are

$$\begin{aligned} F_0 &\rightarrow a \mid b & S_n &\rightarrow F_n(F_n F_{n-1})F_{n-2} \dots F_1 F_0 \\ F_1 x_0 &\rightarrow d F_0 x_0 x_0 & F_{i+2} x_0 x_1 \dots x_{i+1} &\rightarrow x_0(x_0 x_1)x_2 \dots x_{i+1} \end{aligned}$$

Requirements 1–3 are clearly satisfied. To deduce 4, observe that applying deterministic production rules only, $S_n \rightarrow^* F_1^{(2^1_n)} F_0$, where $X^{(k)}$ denotes the k -fold iteration of X . Thus we see that $L(\mathcal{R}_n)$ is the set of complete binary trees of height $2^1_n + 1$ with each leaf and inner node labelled by either a or b , i.e. $|L(\mathcal{R}_n)| \geq 2^1_{n+2}$. \square

3.3. Recursion schemes with pattern-matching

To control the space of derivations we will utilise recursion schemes equipped with pattern-matching, introduced in [32]. In their full generality pattern-matching recursion schemes form a Turing complete model of computation [32], though we will require only the decidable subclass in which pattern-matching is restricted to decomposing sequences. The following definition presents the particular schemes we utilise.

Definition 3.19 (*Pattern-matching recursion scheme*). A *pattern-matching recursion scheme* is a tuple $\mathcal{R} = \langle \Sigma, \mathcal{N}, \mathcal{S}, \mathcal{P} \rangle$ where Σ , \mathcal{N} and \mathcal{S} are as in Definition 3.11 and \mathcal{P} may include type-preserving production rules of the form

$$F x_0 \dots x_{k-1} \langle x_k, \dots, x_{k+l} \rangle \rightarrow_{\mathcal{R}} t$$

where t is a $\Sigma \cup \mathcal{N} \cup \{x_i \mid i \leq k+l\}$ -term over Σ of prime type.

The associated reduction relation $r \rightarrow_{\mathcal{R}} s$ is defined by the two conditions in Definition 3.12 and an additional clause:

- $r = F r_0 \dots r_{k-1} \langle r_k, \dots, r_{k+l} \rangle$ for some $F \in \mathcal{N}$ of arity k and terms $\vec{r} = (r_i)_{i \leq k+l}$, and there is a production rule $F x_0 \dots x_{k-1} \langle x_k, \dots, x_{k+l} \rangle \rightarrow_{\mathcal{R}} t$ such that $s = t(\vec{r}/\vec{x})$.

The definition of a derivation and language for pattern-matching recursion schemes are analogous.

Pattern-matching recursion schemes can be simulated by higher order recursion schemes using constants representing projection functions for pairs in place of pattern-matching. In particular, the upper-bounds given by Theorem 3.16 and Corollary 3.17 apply to pattern-matching recursion schemes without change. There is, however, a subtle difference between the two in the presence of non-determinism and this will be exploited heavily in the next section. In the remainder of this paper *recursion scheme* refers to pattern-matching recursion schemes unless otherwise stated.

4. Herbrand schemes

We now turn to the task of associating to each sequent calculus proof π with Σ_1 end-sequent a non-deterministic higher order recursion scheme \mathcal{H}_π . The recursion scheme \mathcal{H}_π , which we term the *Herbrand*

scheme of π , contains a non-terminal N_π^i for each sub-proof $\pi \vdash A_0, \dots, A_k$ and each $i \leq k$. The interpretation of N_π^i is of a function which returns a witness (possibly involving explicit substitutions) for each weak quantifier in A_i given input for every strong quantifier in the sequent. The arity of N_π^i is $k + 2$, namely one greater than the length of the sequent: the first argument is a substitution stack and the $(j + 1)$ -th argument is the ‘input’ for the formula A_j . The type of N_π^i depends only on the quantifier structure of the formulæ in the sequent. In particular, the types of N_π^i and N_π^j differ only in their co-domain. Reduction rules governing the non-terminal N_π^i are determined by the final inference in π and value of i , and re-write the non-terminal to a term built from non-terminals for the immediate sub-proofs of π . Hence, they are independent of the particular starting proof. This property implies that the typing and re-write rules for a non-terminal N_π^i are invariant across all Herbrand schemes for proofs that feature π as a sub-proof, whence we may consider two Herbrand schemes as comprising identical sets of non-terminals and production rules and differing only in the selection of start symbols.

For proofs with only $\Pi_2 \cup \Sigma_2$ cuts, Herbrand schemes closely resemble the context-free grammars introduced in [4]. The ‘generic’ case of the representation, enabling the interpretation of cuts of arbitrary quantifier complexity, is when the cut formula on both sides of a cut feature weak/strong quantifier alternations, i.e. $\Pi_n \cup \Sigma_n$ for $n \geq 3$.

We start by introducing the types that occur most prominently in Herbrand schemes. To each prenex formula F we assign two types, the *output* type, τ_F , and the *input* type, $\hat{\tau}_F$, representing the ‘existential’ and ‘universal’ structure of F respectively. These types are determined by the quantifiers in F and are defined as follows. For quantifier-free F , $\tau_F = \hat{\tau}_F = \epsilon$; otherwise,

$$\begin{aligned} \tau_{\forall v F} &= \tau_F & \tau_{\exists v F} &= \begin{cases} \iota \times \tau_F, & \text{if } u(F) = 0, \\ \iota \times (\tau_{\bar{F}} \rightarrow \tau_F), & \text{if } u(F) > 0, \end{cases} \\ \hat{\tau}_{\forall v F} &= \tau_{\exists v \bar{F}} & \hat{\tau}_{\exists v F} &= \tau_{\exists v F} \rightarrow \tau_{\bar{F}} \end{aligned}$$

Example 4.1. We compute the input and output types for prenex Π_2 and Σ_2 formulæ. Let \vec{u} and \vec{v} be sequences of variables of non-zero length $|\vec{u}|$ and $|\vec{v}|$ respectively and C_{qf} any quantifier-free formula. Then

$$\begin{aligned} \tau_{\exists \vec{v} C_{qf}} &= \iota^{|\vec{v}|} & \tau_{\forall \vec{v} C_{qf}} &= \epsilon & \hat{\tau}_{\exists \vec{v} C_{qf}} &= \iota^{|\vec{v}|} \rightarrow \epsilon & \hat{\tau}_{\forall \vec{v} C_{qf}} &= \iota^{|\vec{v}|} \\ \tau_{\forall \vec{u} \exists \vec{v} C_{qf}} &= \iota^{|\vec{v}|} & \tau_{\exists \vec{u} \forall \vec{v} C_{qf}} &= \underbrace{\iota \times \dots \times \iota}_{|\vec{u}|} (\iota^{|\vec{v}|} \rightarrow \epsilon) & \hat{\tau}_{\forall \vec{u} \exists \vec{v} C_{qf}} &= \iota^{|\vec{u}|} & \hat{\tau}_{\exists \vec{u} \forall \vec{v} C_{qf}} &= \iota^{|\vec{u}|} \rightarrow \iota^{|\vec{v}|} \end{aligned}$$

Lemma 4.2. Let F be a prenex formula and \vec{v} a sequence of variables of length k . Then

1. τ_F is a non-simple basic prime type.
2. $\tau_F = \tau_{F(\bar{\tau}/\bar{\alpha})}$ and $\hat{\tau}_F = \hat{\tau}_{F(\bar{\tau}/\bar{\alpha})}$.
3. If $e(F) > 0$ then $\hat{\tau}_F = \hat{\tau}_{\bar{F}} \rightarrow \tau_{\bar{F}}$.
4. $\tau_{\exists \vec{v} F} = \underbrace{\iota \times \dots \times \iota}_k \times \hat{\tau}_{\bar{F}}$ and $\hat{\tau}_{\forall \vec{v} F} = \underbrace{\iota \times \dots \times \iota}_k \times \hat{\tau}_F$.

Proof. By definition. \square

Lemma 4.3. Fix a prenex formula A . The order of τ_A , $\hat{\tau}_A$ are as presented in Table 1 where $\dot{-}$ denotes subtraction truncated at 0, i.e. $n \dot{-} m = \max\{n - m, 0\}$.

Proof. By induction on complexity of A . If A is quantifier free then $\tau_A = \epsilon = \hat{\tau}_A$ so $\text{ord}(\tau_A) = \text{ord}(\hat{\tau}_A) = 0$. Moreover, by Example 4.1, the lemma holds for $A \in (\Sigma_1 \setminus \Pi_1) \cup (\Pi_1 \setminus \Sigma_1)$. Suppose $n > 1$. For $A = \exists \vec{v} B$ where $B \in \Pi_{n-1} \setminus \Sigma_{n-1}$,

Table 1
Order of types τ_A and $\hat{\tau}_A$.

A	Σ_0	$\Sigma_n \setminus \Pi_n$	$\Pi_n \setminus \Sigma_n$
$ord(\tau_A)$	0	$n \dot{-} 2$	$n \dot{-} 3$
$ord(\hat{\tau}_A)$	0	$n \dot{-} 1$	$n \dot{-} 2$

$$\begin{aligned}
 ord(\tau_A) &= ord(\hat{\tau}_{\bar{B}}) && \text{(Lemma 4.2(4))} \\
 &= n - 2 && \text{(induction hypothesis)} \\
 ord(\hat{\tau}_A) &= \max\{ord(\tau_A) + 1, ord(\tau_{\bar{B}})\} && \text{(definition)} \\
 &= n - 1 && \text{(induction hypothesis)}
 \end{aligned}$$

For $A = \forall \bar{v} B$ where $B \in \Sigma_{n-1} \setminus \Pi_{n-1}$,

$$\begin{aligned}
 ord(\tau_A) &= ord(\tau_B) && \text{(definition)} \\
 &= n \dot{-} 3 && \text{(induction hypothesis)} \\
 ord(\hat{\tau}_A) &= ord(\tau_{\bar{A}}) && \text{(definition)} \\
 &= n - 2 && \square
 \end{aligned}$$

We have been describing the types τ_F and $\hat{\tau}_F$ as representing, respectively, the existential and universal structure of F . Beyond the basic case of $\Pi_1 \cup \Sigma_1$ formulæ this view may not be obvious from the definition and requires explanation. Consider first the case of a Π_2 formula $F = \forall x \forall y \exists z G_{qf}(x, y, z)$. (We assume two universal quantifiers to provide contrast between the ‘type’ of the universal quantifiers, $\forall x \forall y$, corresponding to the pair $\iota \times \iota$, and the existential, $\exists z$, whose type is ι .) As was seen in the example above, $\hat{\tau}_F = \iota \times \iota$ and $\tau_F = \iota$ represent the ‘type’ of the two quantifier kinds in F . However, thinking of the ‘computational content’ of a Π_2 formula such as F suggests an alternate explanation, namely a function $f: \hat{\tau}_F \rightarrow \tau_F$ such that $\forall x \forall y G_{qf}(x, y, f(x, y))$. The type of this function is precisely the type we associate to the universal structure of the dual of F . That is, for a formula $H = \exists x \exists y \forall z I_{qf}(x, y, z)$ we have $\hat{\tau}_H = \iota \times \iota \rightarrow \iota$ which, by the above, is simply the type of the universal quantifier in the Skolemised form $\forall g \exists x \exists y I_{qf}(x, y, g(x, y))$. The existential structure of H is defined as $\tau_H = \iota \times \iota \times (\iota \rightarrow \epsilon)$. Morally, this is no different from the type of the existential quantifiers $\iota \times \iota$ (as ϵ is the unit type, any function $a: \sigma \rightarrow \epsilon$ is observationally equivalent to the constant function $\lambda x^\sigma \langle \rangle$). The additional structure arises because, for the purposes of defining its existential ‘content’, we treat the Σ_2 formula H as a Σ_3 formula with a vacuous inner existential block, $H' = \exists x' \exists y' \forall z' \exists w^\epsilon I'_{qf}(x, y, z, w)$. Skolemising the inner quantifier yields $\exists x' \exists y' \exists f^{\iota \rightarrow \epsilon} \forall z' I'_{qf}(x, y, z, fz)$ which has the desired type.

The astute reader will have observed the similarity between the types and formulæ above and those arising in Gödel’s functional interpretation. Indeed, the pattern extends to the whole prenex quantifier hierarchy.

Theorem 4.4. *Shoenfield’s variant of the functional interpretation [38] translates a prenex formula A in the language of (classical) first-order logic to a prenex Π_2 formula which is provably equivalent to a formula $\forall x^{\hat{\tau}_A} \exists y^{\tau_A} F_{qf}$ over a system of first-order logic in finite types extended with pairing and unit.*

Proof. Shoenfield’s functional interpretation [38] associates to each formula A in the language of arithmetic a prenex Π_2 formula $A^S = \forall x \exists y A_S(x, y)$ of HA^ω such that if $\text{PA} \vdash A$ then for some term t , $\text{HA}^\omega \vdash \forall x A_S(x, tx)$. In [17], Gerhardy and Kohlenbach apply the interpretation to pure predicate logic yielding Π_2 formulæ in the language of “extensional predicate logic in all finite types”, denoted E-PL^ω . For the present lemma we utilise the expansion of E-PL^ω to the hierarchy of basic types (including new constant

symbols). Utilising pair types to collapse blocks of like quantifiers, we may assume A^S always has the form $\forall x^{\rho_A} \exists y^{\sigma_A} A_S$ where A_S is quantifier free and ρ_A, σ_A are basic types. If A is quantifier-free we may take $\rho_A = \sigma_A = \epsilon$, in which case A^S is already of the desired form. If $A = \forall v B$ and $B^S = \forall x^{\rho_B} \exists y^{\sigma_B} B_S$ then by construction $A^S = \forall x^{\iota \times \rho_B} \exists y^{\sigma_B} A_S$. Since $\hat{\tau}_A = \iota \times \hat{\tau}_B$ and $\tau_A = \tau_B$, the induction hypothesis shows A^S is equivalent to a formula $\forall x^{\hat{\tau}_A} \exists y^{\tau_A} A_{qf}^*$. For the existential case, let $A = \exists v B$ with $u(B) > 0$ and suppose $B^S = \forall x^{\rho_B} \exists y^{\sigma_B} B_S$. Then $A^S = \forall x^{\rho_A} \exists y^{\iota \times \sigma_A} A_S$ where $\sigma_A = \rho_B \rightarrow \sigma_B$ and $\rho_A = \iota \times \sigma_A \rightarrow \rho_B$. Using the fact that $\tau_A = \iota \times (\hat{\tau}_B \rightarrow \tau_B)$ and $\hat{\tau}_A = \tau_A \rightarrow \hat{\tau}_B$, the induction hypothesis shows A^S is logically equivalent to a formula $\forall x^{\hat{\tau}_A} \exists y^{\tau_A} A_{qf}^*$ for appropriate A_{qf}^* . The case that A leads with a block of existential quantifiers is analogous. \square

We now present the definition of the Herbrand scheme associated to LK proofs.

Definition 4.5 (*Herbrand scheme*). Fix a proof $\pi \vdash A_0, \dots, A_k$ with Σ_1 end-sequent and let Σ_π be the simple alphabet consisting of a constant symbol c of type ι and the function symbols and eigenvariables occurring in π (typed accordingly). The *Herbrand scheme for π* is the higher order recursion scheme $\mathcal{H}_\pi = \langle \Sigma_\pi, \mathcal{N}_\pi, \mathcal{S}_\pi, \mathcal{P}_\pi \rangle$ with the following non-terminals and production rules.

1. A non-terminal $c_\rho : \rho$ for each basic type $\rho \notin \{\iota, \epsilon\}$ that occurs as a sub-type of a type τ_B or $\hat{\tau}_B$ for a formula B occurring in π , with production rules

$$\begin{aligned} c_\rho &\rightarrow \langle c_{\tau_0}, c_{\tau_1} \rangle && \text{if } \rho = \tau_0 \times \tau_1, \\ c_\rho x_0^{\rho_0} \cdots x_k^{\rho_k} &\rightarrow c_{co(\rho)} && \text{if } \rho = \rho_0 \rightarrow \cdots \rightarrow \rho_k \rightarrow co(\rho), \end{aligned}$$

with c_ι and c_ϵ defined to be the constants c and $\langle \rangle$ respectively.

2. A non-terminal $N_{\pi'}^i$ for each sub-proof $\pi' \vdash B_0, \dots, B_l$ of π and for each $i \leq l$, with type

$$N_{\pi'}^i : \varsigma \rightarrow \hat{\tau}_{B_0} \rightarrow \cdots \rightarrow \hat{\tau}_{B_l} \rightarrow \tau_{B_i}$$

and production rule(s) as given in Table 2, determined in each case by the final inference of π' .

3. A start symbol $S_{\pi,i} : \tau_{A_i}$ for each $i \leq k$ with associated production rules

$$S_{\pi,i} \rightarrow N_{\pi'}^i \perp c_{\tau_{A_0}}^* \cdots c_{\tau_{A_k}}^*$$

The *language of π* is the set $L(\pi) = \{(i, r^\circ) \mid i \leq k \text{ and } (S_{\pi,i}, r) \in L(\mathcal{H}_\pi)\}$.

It remains to check that the production rules of Herbrand schemes are well-typed. This task will be taken up later in Lemma 4.13. For now we take for granted the fact that Herbrand schemes are well-defined and continue with some basic properties of them (Lemmas 4.7 to 4.10) followed by the intended interpretation of the schemes as generating Herbrand disjunctions (Definition 4.11) and the observation that this interpretation coincides with the Herbrand set for quasi cut-free proofs (Lemma 4.12). We start, however, with a brief explanation of some of the production rules from Table 2.

Remark 4.6. We comment on some of the rules from Table 2.

- **Axiom.** We are restricting axioms to quantifier-free formulæ only, which motivates the simple production rule given in the table. One may wish to permit axioms $\pi \vdash A_0, A_1$ where $A_1 = \bar{A}_0$ has arbitrary (prenex) complexity. These can be accommodated by the production rules

Table 2

Production rules for Herbrand schemes. \vec{x} and \vec{y} are sequences of distinct variable symbols of length $m := |\Gamma|$ and $n := |\Delta|$ respectively.

Inference deriving π	Production rule(s)
$\text{ax}: \pi \vdash A, \bar{A}$	$\text{N}_{\pi}^i ax_0x_1 \rightarrow \langle \rangle$
$\vee \frac{\pi_0 \vdash \Gamma, A, B}{\pi \vdash \Gamma, A \vee B}$	$\text{N}_{\pi}^i a\vec{x}z \rightarrow \begin{cases} \text{N}_{\pi_0}^i a\vec{x}zz, & \text{if } i < m, \\ \langle \rangle, & \text{otherwise.} \end{cases}$
$\wedge \frac{\pi_0 \vdash \Gamma, A \quad \pi_1 \vdash \Delta, B}{\pi \vdash \Gamma, \Delta, A \wedge B}$	$\text{N}_{\pi}^i a\vec{x}\vec{y}z \rightarrow \begin{cases} \text{N}_{\pi_0}^i a\vec{x}z, & \text{if } i < m, \\ \text{N}_{\pi_1}^{i-m} a\vec{y}z, & m \leq i < m+n, \\ \langle \rangle, & \text{otherwise.} \end{cases}$
$\forall_{\alpha} \frac{\pi_0 \vdash \Gamma, A(\vec{\alpha}/\vec{v})}{\pi \vdash \Gamma, \forall \vec{v}A}$	$\text{N}_{\pi}^i a\vec{x}\langle z_0, \dots, z_{p+1} \rangle \rightarrow$
$\exists_{\vec{r}} \frac{\pi_0 \vdash \Gamma, A(\vec{r}/\vec{v})}{\pi \vdash \Gamma, \exists \vec{v}A}$	$\text{N}_{\pi}^i a\vec{x}z \rightarrow \begin{cases} \vec{r} \cdot a \star (\text{N}_{\pi_0}^m a\vec{x}), & \text{if } i = m \text{ and } u(A) > 0, \\ \vec{r} \cdot a \star \langle \rangle, & \text{if } i = m \text{ and } u(A) = 0, \\ \text{N}_{\pi_0}^i a\vec{x}(z(\text{N}_{\pi}^m a\vec{x}z)) & \text{if } i \neq m, \end{cases}$
$\text{cut} \frac{\pi_0 \vdash \Gamma, A \quad \pi_1 \vdash \Delta, \bar{A}}{\pi \vdash \Gamma, \Delta}$	$\text{N}_{\pi}^i a\vec{x}\vec{y} \rightarrow \begin{cases} \text{N}_{\pi_0}^i a\vec{x}((\text{N}_{\pi_1}^n a\vec{y}) \circ_A (\text{N}_{\pi_0}^m a\vec{x})), & \text{if } i < m, \\ \text{N}_{\pi_1}^{i-m} a\vec{y}((\text{N}_{\pi_0}^m a\vec{x}) \circ_{\bar{A}} (\text{N}_{\pi_1}^n a\vec{y})), & \text{if } m \leq i. \end{cases}$
$\text{w} \frac{\pi_0 \vdash \Gamma}{\pi \vdash \Gamma, A}$	$\text{N}_{\pi}^i a\vec{x}z \rightarrow \begin{cases} c_{\tau_A}, & \text{if } i = m, \\ \text{N}_{\pi_0}^i a\vec{x}, & \text{otherwise.} \end{cases}$
$\text{c} \frac{\pi_0 \vdash \Gamma, A, A}{\pi \vdash \Gamma, A}$	$\text{N}_{\pi}^i a\vec{x}z \rightarrow \begin{cases} \text{N}_{\pi_0}^i a\vec{x}zz, & \text{if } i < m, \\ \text{N}_{\pi_0}^i a\vec{x}zz \mid \text{N}_{\pi_0}^{i+1} a\vec{x}zz, & \text{if } i = m. \end{cases}$
$\text{p} \frac{\pi_0 \vdash \Gamma, B, A, \Delta}{\pi \vdash \Gamma, A, B, \Delta}$	$\text{N}_{\pi}^i a\vec{x}z_0z_1\vec{y} \rightarrow \begin{cases} \text{N}_{\pi_0}^{i+1} a\vec{x}z_1z_0\vec{y}, & \text{if } i = m, \\ \text{N}_{\pi_0}^{i-1} a\vec{x}z_1z_0\vec{y}, & \text{if } i = m+1, \\ \text{N}_{\pi_0}^i a\vec{x}z_1z_0\vec{y}, & \text{otherwise.} \end{cases}$

$$\vec{\alpha} = (\alpha_0, \dots, \alpha_p)$$

$$\vec{r} = (r_0, \dots, r_p)$$

$$\vec{r} \cdot a = (r_0 \cdot a, \dots, r_p \cdot a)$$

$$(u_j)_{j \leq q} \star t = \langle u_0, \dots, u_q, t \rangle$$

$$r \circ_A s = \begin{cases} r, & \text{if } e(A) > 0, \\ rs, & \text{if } u(A) > 0, \\ \langle \rangle, & \text{otherwise.} \end{cases}$$

$$\text{N}_{\pi}^i ax_0x_1 \rightarrow \begin{cases} x_{1-i}, & \text{if } u(A_i) = 0, \\ x_{1-i}x_i, & \text{if } u(A_i) > 0, \end{cases}$$

which the interested reader can check are well-typed. This definition mimics the behaviour of the Herbrand scheme for the natural proof of A_0, A_1 that uses only quantifier-free instances of axioms and alternate applications of \exists and \forall inferences. Our reason for favouring quantifier-free axioms is that, as a consequence, production rules never return their arguments as output, a fact that simplifies some technical aspects of the later analysis (specifically Lemma 6.9).

- \wedge and \vee . As proofs involve prenex formulæ only, conjunctions and disjunctions are necessarily quantifier-free with associated type ϵ , and therefore possess no computational content relevant to the construction of a Herbrand disjunction. When focusing on such formulæ, the production rule in each case returns the empty sequence.
- $\exists_{\vec{r}}$. The production rule in this case depends on both i and the quantifier form of the active formula. Consider the instance of $\exists_{\vec{r}}$ given in Table 2. As π is assumed regular, the active formula (A in the table) is either quantifier-free or universally quantified. If i marks the active formula (i.e. $i = m$) then the production rule for N_{π}^i directly outputs the witness terms provided by the proof and the current substitution (the sequence $(r_0 \cdot a, \dots, r_p \cdot a)$) as the first $p + 1$ components of a nested pair. The final component is either trivial (in case A is quantifier-free) or, if A is universally quantified, the continuation

of the trace to the immediate sub-proof in the form of a function. If $i \neq m$, the production rule instead passes the above term to the corresponding argument.

- $\forall_{\vec{\alpha}}$. This is the only case that involves pattern matching in Herbrand schemes. Although it can be simulated by a recursion scheme without pattern-matching using projection functions for pair types, doing so introduces a duplication of arguments that is avoided in the chosen formulation. For instance, the production rule for $\forall_{\vec{\alpha}}$ where $\vec{\alpha}$ consists of the single eigenvariable α and the sequent Γ is empty yields the production rule

$$\mathbf{N}_{\pi}^0 a \langle z_0, z_1 \rangle \rightarrow \mathbf{N}_{\pi_0}^0 ([\alpha \leftarrow z_0] a) z_1$$

which may be simulated by the rule

$$\mathbf{N}_{\pi}^0 a z \rightarrow \mathbf{N}_{\pi_0}^0 ([\alpha \leftarrow \mathbf{p}_0 z] a) (\mathbf{p}_1 z) \quad (1)$$

where \mathbf{p}_0 and \mathbf{p}_1 are constants representing the two projection functions for pair types. If s is a term such that $s \rightarrow_{\mathcal{H}}^* \langle r_0, s_0 \rangle \mid \langle r_1, s_1 \rangle$ and the four sub-terms are pairwise distinct then the reduction in (1) permits the derivation $\mathbf{N}_{\pi}^0 \perp s \rightarrow^* \mathbf{N}_{\pi_0}^0 ([\alpha \leftarrow \mathbf{p}_0 \langle r_0, s_0 \rangle] \perp) (\mathbf{p}_1 \langle r_1, s_1 \rangle)$, essentially the term $\mathbf{N}_{\pi_0}^0 ([\alpha \leftarrow r_0] \perp) s_1$, which is forbidden in the Herbrand scheme due to pattern-matching. In this sense pattern matching plays a role analogous to the rigidity conditions utilised in [21,1,2] for representing first-order proofs with Π_1/Π_2 cut complexity.

- **cut.** For each choice of i , the rule provides exactly one reduction for the non-terminal \mathbf{N}_{π}^i : for $i < m$ this is

$$\mathbf{N}_{\pi}^i a \vec{x} \vec{y} \rightarrow \begin{cases} \mathbf{N}_{\pi_0}^i a \vec{x} (\mathbf{N}_{\pi_1}^n a \vec{y} (\mathbf{N}_{\pi_0}^m a \vec{x})), & \text{if } u(A) > 0, \\ \mathbf{N}_{\pi_0}^i a \vec{x} (\mathbf{N}_{\pi_1}^n a \vec{y}), & \text{if } e(A) > 0, \\ \mathbf{N}_{\pi_0}^i a \vec{x} \langle \rangle, & \text{if } A \text{ is q.f.} \end{cases}$$

Note that, in the case $e(A) = 0$ the type $\hat{\tau}_A$ (which marks the final argument to $\mathbf{N}_{\pi_0}^i$) is prime, and is otherwise the function type $\hat{\tau}_A \rightarrow \tau_{\bar{A}}$. Moreover, the case distinction above is independent of i . For instance, if $A = \forall v B$ exactly the following production rules arise from the cut.

$$\mathbf{N}_{\pi}^j a \vec{x} \vec{y} \rightarrow \begin{cases} \mathbf{N}_{\pi_0}^j a \vec{x} (\mathbf{N}_{\pi_1}^n a \vec{y} (\mathbf{N}_{\pi_0}^m a \vec{x})), & \text{if } j < m, \\ \mathbf{N}_{\pi_1}^{m-j} a \vec{y} (\mathbf{N}_{\pi_0}^m a \vec{x}), & \text{if } m \leq j < m + n. \end{cases}$$

In the following let $\mathcal{H} = \langle \Sigma, \mathcal{N}, \mathcal{S}, \mathcal{P} \rangle$ be the Herbrand scheme for a regular proof π with prenex Σ_1 end-sequent.

Lemma 4.7. *\mathcal{H} is an acyclic recursion scheme. Hence, $L(\pi)$ is finite.*

Proof. Let $<$ be the transitive relation on non-terminals in \mathcal{H} generated by the equations: $\mathbf{c}_{\rho} < \mathbf{c}_{\sigma}$ if ρ is a proper sub-type of σ ; $\mathbf{c}_{\rho} < \mathbf{N}_{\pi_0}^i$ for every ρ , sub-proof π_0 of π and i ; $\mathbf{N}_{\pi_0}^i < \mathbf{N}_{\pi_1}^j$ if either π_0 is a proper sub-proof of π_1 or $\pi_0 = \pi_1$ and $j < i$; and $\mathbf{N}_{\pi}^i < \mathbf{S}_{\pi, i}$ for any i . Clearly $<$ is acyclic and irreflexive. Moreover, for every production rule $F \vec{x} \rightarrow_{\mathcal{H}} t$ and any non-terminal \mathbf{G} occurring in t we have $\mathbf{G} < F$. \square

Lemma 4.8. *Every \mathcal{H} -term of simple type is a Σ -term, and every \mathcal{H} -term of substitution stack type has the form either \perp or $[\alpha \leftarrow s] b$ for some $\alpha \in \Sigma$, Σ -term s and $b : \varsigma$.*

Proof. The non-terminals of \mathcal{H} all have type one of three forms: ϵ , pair type, or function type with non-simple co-domain. It therefore follows that the only \mathcal{H} -terms of simple type are the Σ -terms. Likewise, \perp

and $[\alpha \leftarrow s]b$ are the only kind of \mathcal{H} -terms of type ς . Given a substitution stack $[\alpha \leftarrow s]b$ however, as $\alpha \in \Sigma$ the first part of the lemma implies that s is a Σ -term. \square

Lemma 4.9. *If $r : \sigma \rightarrow \tau$ is a \mathcal{H} -term then τ is a basic type and σ is either basic or the type of substitution stacks. In the latter case, $r = N_\pi^i$ or \hat{N}_π^i for some π and i .*

Proof. By inspection of the types of non-terminals and terms. \square

Lemma 4.10. *Suppose r is an \mathcal{H} -term of type ϵ containing no explicit substitutions (i.e. having no sub-term of the form $t \cdot a$). If $r \rightarrow_{\mathcal{H}}^* s$ for some basic term s then $s = \langle \rangle$.*

Proof. By induction on the proof generating \mathcal{H} , on the composition of r and the length of the derivation $r \rightarrow^* s$. \square

We now describe how Herbrand schemes can be interpreted as ascribing existential content to first-order proofs.

Definition 4.11 (Herbrand expansion). Let $\pi \vdash \Gamma$ be a proof with $\Gamma = \exists \vec{v}_0 A_0, \dots, \exists \vec{v}_k A_k$ where A_i is quantifier-free for each $i \leq k$. Let k_i be the length of \vec{v}_i . The *Herbrand expansion* of π is the quantifier free sequent Γ^π given by

$$\Gamma^\pi := \{A_i(\vec{r}_i/\vec{v}_i) \mid \vec{r}_i = (r_j)_{j < k_i} \text{ and } (i, \langle r_0, \dots, r_{k_i-1}, \langle \rangle \rangle) \in L(\pi)\}.$$

Lemma 4.12. *If $\pi \vdash \Gamma$ is a quasi cut-free proof of a Σ_1 end-sequent then the Herbrand expansion of π is a valid sequent and $\bigvee \Gamma^\pi$ is a Herbrand disjunction in the sense of Theorem 2.3.*

Proof. Observe that in every production rule associated to a quantifier-free cut, the term $r \circ_A s$ becomes $\langle \rangle$. Derivations in π are therefore in 1-1 correspondence with traces following the breakdown of formulae in the end-sequent. As a result we observe that $L(\pi)$ simply outputs all literal witnesses to the existential quantifiers in the end-sequent. \square

The idea behind Herbrand schemes is to provide a generalisation of the above lemma to proofs containing quantified cuts. The analysis necessary for the result is carried out in Section 7. In the remainder of this section we prove the production rules of Herbrand schemes are well-typed and derive upper bounds on the size of Herbrand expansions.

Lemma 4.13. *The production rules of Herbrand schemes are type preserving.*

Proof. Fix a proof π with prenex end-sequent A_0, \dots, A_m and $i \leq m$. We establish type-preservation of the production rules for the non-terminals N_π^0, \dots, N_π^m via a case distinction on the final inference rule in π .

Suppose $\pi \vdash A_0, \dots, A_{m-1}, \exists \vec{v} A$ is obtained from proof π_0 by $\exists \vec{v}$. Thus $\pi_0 \vdash \Gamma, A(\vec{r}/\vec{v})$ for some sequence $\vec{r} = (r_j)_{j \leq k}$ of simple Σ -terms of type ι . By regularity, $e(A) = 0$, i.e. either A is quantifier-free or $u(A) > 0$. Let $\Gamma = A_0, \dots, A_{m-1}$ and fix a term $z : \hat{\tau}_{\exists \vec{v} A}$ and a sequence of terms \vec{x} of length m such that $N_\pi^i \vec{x} z$ is well-typed. By definition $N_{\pi_0}^i$ has type

$$N_{\pi_0}^i : \varsigma \rightarrow \hat{\tau}_{A_0} \rightarrow \dots \rightarrow \hat{\tau}_{A_{m-1}} \rightarrow \hat{\tau}_A \rightarrow \begin{cases} \tau_A, & \text{if } i = m, \\ \tau_{A_i}, & \text{otherwise.} \end{cases}$$

To check type preservation there are two cases to consider:

1. $i = m$. If $u(A) = 0$ then A is quantifier-free and $\hat{\tau}_{\bar{A}} = \epsilon$. If $u(A) > 0$ then $\hat{\tau}_{\bar{A}} = \hat{\tau}_A \rightarrow \tau_A$ by Lemma 4.2(3), so the type of $N_{\pi_0}^i a\vec{x}$ is $\hat{\tau}_{\bar{A}}$. Since also $\tau_{\exists\vec{v}A} = \underbrace{\iota \times \cdots \times \iota}_k \times \hat{\tau}_{\bar{A}}$ by Lemma 4.2(4), we are done.
2. $i \neq m$. In this case it is necessary to check that $\hat{\tau}_{\exists\vec{v}A} = \tau_{\exists\vec{v}A} \rightarrow \hat{\tau}_A$. But this follows directly from the definition and the fact that $\tau_{\forall\vec{v}\bar{A}} = \tau_{\bar{A}} = \hat{\tau}_A$ as $e(A) = 0$.

Suppose π is derived from π_0 via the inference $\forall_{\vec{\alpha}}$ and $A_m = \forall\vec{v}A$ with $u(A) = 0$ and $\vec{\alpha} = (\alpha_j)_{j < k}$. Let $i \leq m$ and fix terms $\vec{x}, \vec{z} = (z_0, \dots, z_k)$ such that $N_{\pi}^i a\vec{x}\langle z_0, \dots, z_k \rangle$ is well-typed. Lemma 4.2 implies that $z_j : \iota$ for each $j < k$, and $z_k : \hat{\tau}_A$. Thus $N_{\pi_0}^i b\vec{x}z_k$ is well-typed and has type $\tau_A = \tau_{A_m}$.

Suppose π is derived via cut from sub-proofs $\pi_0 \vdash \Gamma, A$ and $\pi_1 \vdash \Delta, \bar{A}$. Let $m = |\Gamma|$ and $n = |\Delta|$ and fix \vec{x} and \vec{y} suitably typed. Without loss of generality we may assume $i < m$, in which case we require to show $(N_{\pi_1}^n a\vec{y}) \circ_A (N_{\pi_0}^m a\vec{x}) : \hat{\tau}_A$ which reduces (via Remark 4.6) to proving

$$e(A) > 0 \text{ implies } \hat{\tau}_A = \hat{\tau}_{\bar{A}} \rightarrow \tau_{\bar{A}},$$

$$u(A) > 0 \text{ implies } \hat{\tau}_A = \tau_{\bar{A}} \text{ and } \hat{\tau}_{\bar{A}} = \hat{\tau}_A \rightarrow \tau_A,$$

both of which follow directly from Lemma 4.2.

The remaining cases are straightforward and omitted. \square

Lemma 4.14. *For a proof $\pi \vdash A_0, \dots, A_k$ and $i < k$, the order of the non-terminal N_{π}^i is the smallest n such that $\{A_j : j \leq k\} \subset \Pi_{n+1}$, unless A_i is Π_1 , in which case the order of N_{π}^i is zero.*

Proof. If A_i is Π_1 then $\tau_{A_i} = \epsilon$ and the order of N_{π}^i is 0 by definition. Otherwise, by Lemma 4.3, the order of N_{π}^i is one greater than the maximum among the orders of $\hat{\tau}_{A_0}, \dots, \hat{\tau}_{A_k}$ which is the smallest n such that every A_j is Π_{n+1} . \square

It is now possible to strengthen Lemma 4.7 to a concrete bound on the number of terms derivable from a Herbrand scheme. The idea is to eliminate occurrences of pattern-matching in a Herbrand scheme \mathcal{H} in a way that does not decrease the length of derivations so that Theorem 3.16 and Corollary 3.17 can be applied.

Theorem 4.15. *If $\pi \vdash \Gamma$ is a proof of a single prenex Σ_1 formula in which all cut formulae are contained in $\Pi_n \cup \Sigma_n$ then the size of the Herbrand expansion Γ^{π} is no greater than $2^{4|\pi|_{n+2}^3}$ where $|\pi|$ is the number of inference rules in π .*

Proof. The case $n = 0$ is covered by Lemma 4.12 so suppose $n > 0$. Let \mathcal{H} be the Herbrand scheme of π . Since the cut rank of π is bounded by n , Lemma 4.14 implies that the order of \mathcal{H} is no greater than n . To obtain the desired bounds we apply Theorem 3.16. However, this requires first eliminating the explicit substitutions introduced by the \forall inferences. Let \mathcal{H}' denote the higher order recursion scheme with non-terminals of basic type obtained from \mathcal{H} by removing all substitutions terms and types from non-terminals and production rules. In particular, the productions originating from $\forall_{\vec{\alpha}}$ and $\exists_{\vec{r}}$ inferences are replaced by following in \mathcal{H}' :

$$\forall_{\vec{\alpha}}: N_{\pi}^i \vec{x}\langle z_0, \dots, z_{p+1} \rangle \rightarrow N_{\pi_0}^i \vec{x}z_p$$

$$\exists_{\vec{r}}: N_{\pi}^i \vec{x}z \rightarrow \begin{cases} N_{\pi_0}^i \vec{x}(z(N_{\pi}^m \vec{x})), & i \neq m, \\ \langle c, \dots, c, N_{\pi_0}^m \vec{x} \rangle, & i = m \text{ and } u(A) > 0, \\ \langle c, \dots, c, \langle \rangle \rangle, & i = m \text{ and } u(A) = 0. \end{cases}$$

$$\boxed{
 \begin{array}{c}
 \frac{h : \Gamma \vdash I_0^{\alpha, m\alpha\bar{\alpha}}, I_1^{\bar{\alpha}, m\alpha\bar{\alpha}}}{\exists} \quad \frac{7 : \Delta \vdash \bar{I}_\gamma^{0, \beta}, \bar{I}_\gamma^{s\beta, \bar{\beta}}, T_{\beta, \bar{\beta}}}{\exists} \\
 \frac{g : \Gamma \vdash I_0^{\alpha, m\alpha\bar{\alpha}}, I_1^{\bar{\alpha}}}{\exists} \quad \frac{6 : \Delta \vdash \bar{I}_\gamma^{0, \beta}, \bar{I}_\gamma^{s\beta, \bar{\beta}}, T}{\forall_\beta} \\
 \frac{f : \Gamma \vdash I_0^\alpha, I_1^{\bar{\alpha}}}{\forall_\alpha} \quad \frac{5 : \Delta \vdash \bar{I}_\gamma^{0, \beta}, \bar{I}_\gamma^{s\beta}, T}{\exists} \\
 \frac{e : \Gamma \vdash I_0, I_1^{\bar{\alpha}}}{\forall_\alpha} \quad \frac{4 : \Delta \vdash \bar{I}_\gamma^{0, \beta}, \bar{I}_\gamma, T}{\forall_\beta} \\
 \frac{d : \Gamma \vdash I_0, I_1}{\exists} \quad \frac{3 : \Delta \vdash \bar{I}_\gamma^0, \bar{I}_\gamma, T}{\exists} \\
 \frac{c : \Gamma \vdash I, I_1}{\exists} \quad \frac{2 : \Delta \vdash \bar{I}_\gamma, \bar{I}_\gamma, T}{c} \\
 \frac{b : \Gamma \vdash I, I}{c} \quad \frac{1 : \Delta \vdash \bar{I}_\gamma, T}{\forall_\gamma} \\
 \frac{a : \Gamma \vdash I}{c} \quad \frac{0 : \Delta \vdash \bar{I}, T}{c} \\
 \text{cut} \quad \frac{\Gamma, \Delta \vdash T}{\Gamma, \Delta \vdash T}
 \end{array}
 }$$

Fig. 4. Proof π_∞ of pigeonhole principle.

The second part of Lemma 4.8 implies that derivations in \mathcal{H}' from the start symbol are in 1-1 correspondence with derivations in \mathcal{H} . Repeating the argument of Corollary 3.17, the size of $L(\pi)$ is therefore bounded by 2^K where K is the length of the longest derivation in \mathcal{H}' from the single start symbol. The order of \mathcal{H}' is no greater than n , the number of non-terminals is bounded by $|\pi|^2$, and for each production rule $F\bar{x} \rightarrow t$ in \mathcal{H}' , $|t|_\Sigma < 3 \times |\pi|$ where Σ is the ranked alphabet of function symbols and constants occurring in π . Theorem 3.16 then implies $K \leq 2_{n+1}^{4|\pi|^3}$. \square

5. A Herbrand disjunction for the pigeonhole principle

We consider a formal proof of the pigeonhole principle for two boxes via the infinite pigeonhole principle. The question of the computational content of this proof is attributed to G. Stolzenberg [13]. A variety of analytic methods have since been applied to this proof [20,10,41,8,1] and its generalisations [37,35]. The version we present here is a formal proof with a single Π_3 cut based on the proof with two Π_2 cuts given in [1,41].

Let $f : \mathbb{N} \rightarrow \{0, 1\}$ be a total Boolean function, let I_i (for $i = 0, 1$) express that there are infinitely many $m \in \mathbb{N}$ for which $f(m) = i$ and T express that there exists $m < n$ such that $f(m) = f(n)$. A consequence of the law of excluded middle is $\exists w I_w$. Moreover, I_i implies T for each $i \in \{0, 1\}$: assuming I_i there exists $m \geq 0$ and $n \geq m + 1$ for which $f(m) = f(n) = i$. Combining these observations we conclude T .

The following formalises the above argument into a proof with a single Π_3 cut. The formal language, Σ , comprises two unary function symbols f, s , one binary function symbol m , a constant symbol 0 and a binary relation \leq . We make the following definitions and abbreviations:

- $T = \exists u \exists v (u < v \wedge fu = fv)$,
- $I = \exists w I_w$ where $I_r = \forall u \exists v (u \leq v \wedge fv = r)$,
- $\Gamma = \{\forall u \forall v (u \leq muv \wedge v \leq muv), \forall u (fu = 0 \vee fu = s0)\}$,
- $\Delta = \{\forall u \forall v \forall w (u = v \wedge w = v \rightarrow u = w), \forall u \forall v (su \leq v \rightarrow u < v)\}$,
- I_r^s and $I_r^{s,t}$ denote, respectively, $\exists v (s \leq v \wedge fv = r)$ and $(s \leq t \wedge ft = r)$,
- $T_{s,t}$ denotes $(s < t \wedge fs = ft)$.

The intended interpretation of the symbols is: f represents the (arbitrary) function f , s the successor function on \mathbb{N} , \leq the standard ordering and m the binary max function.

A formal proof of the pigeonhole principle (namely $\Gamma, \Delta \vdash T$) is given in Fig. 4 which we name π_∞ . The proof is displayed in two-sided sequent calculus as this simplifies the presentation and following discussion. The intended interpretation of the two-sided sequent $A_1, \dots, A_k \vdash B_1, \dots, B_l$ is the sequent $\bar{A}_1, \dots, \bar{A}_k, B_1, \dots, B_l$. For brevity, only eigenvariables and witnesses of the quantifiers and instances of the existential formula T are displayed in π_∞ . The proof fully fleshed out uses about 50 application of the axioms and rules of the calculus but the only cut in π_∞ is the one displayed in the figure. Two normal forms

of the proof of size ~ 200 have been computed in a case study [41] from which one can read off the Herbrand sets for the formula T (also for formulæ in $\Gamma \cup \Delta$ but these are less interesting). Up to interpretation of the logical symbols by their intended semantics, the two Herbrand sets combined provide the witnesses $\{\langle 0, 1 \rangle, \langle 1, 2 \rangle, \langle 2, 3 \rangle, \langle 0, 2 \rangle, \langle 1, 3 \rangle\}$ to the existential quantifiers in T .¹ The Herbrand scheme \mathcal{H}_{π_∞} associated to the proof π_∞ computes the same Herbrand set, a fact we demonstrate in the following.

Types and terms The Herbrand scheme for π_∞ comprises a non-terminal for each sub-proof of π_∞ and each formula in the end-sequent of that sub-proof. Recall, for each sub-proof $p : \Pi \vdash \Lambda$ of π_∞ and each $i < |\Pi| + |\Lambda|$ there is a non-terminal N_p^i in \mathcal{H}_{π_∞} representing the existential content of the i -th formula in the sequent at position p . In the following, in place of N_p^i we will write N_p^A where A is the i -th formula in the sequent assuming this is unique. In case A occurs more than once in the sequent $\Pi \vdash \Lambda$ (such as at positions b and 2) the non-terminal N_p^A refers to the first occurrence of A and we use the notation $N_p^{A^+}$ for the second occurrence. Concerning the type of N_p^A , we recall the types τ_F and $\hat{\tau}_F$ for each formula F in π_∞ . Let $\hat{\iota} = \iota \times (\iota^1 \rightarrow \epsilon)$ and $\hat{\epsilon} = \iota^2 \rightarrow \epsilon$.

- For $F \in \Gamma \cup \Delta$ we have $\tau_{\bar{F}} \in \{\iota^1, \iota^2, \iota^3\}$, and $\hat{\tau}_{\bar{F}} = \tau_{\bar{F}} \rightarrow \epsilon$.
- T : $\tau_T = \iota^2$, $\hat{\tau}_T = \hat{\epsilon}$.
- I_r^s : $\tau_{I_r^s} = \iota^1 = \hat{\tau}_{I_r^s}$, $\hat{\tau}_{I_r^s} = \iota^1 \rightarrow \epsilon$ and $\tau_{\bar{I}_r^s} = \epsilon$.
- I_r : $\tau_{I_r} = \iota^1$, $\tau_{\bar{I}_r} = \hat{\tau}_{I_r} = \hat{\iota}$ and $\hat{\tau}_{\bar{I}_r} = \hat{\iota} \rightarrow \iota^1$.
- I : $\tau_I = \iota \times (\hat{\iota} \rightarrow \iota^1) = \hat{\tau}_{\bar{I}}$, $\tau_{\bar{I}} = \hat{\iota}$ and $\hat{\tau}_I = \tau_I \rightarrow \tau_{\bar{I}}$.
- The remaining formulæ that occur in π_∞ are quantifier-free and are assigned type ϵ in all cases.

According to the definition, the type of $N_{\pi_\infty}^T$ is $\varsigma \rightarrow \hat{\tau}_{\bar{F}} \rightarrow \hat{\tau}_{\bar{G}} \rightarrow \hat{\tau}_{\bar{C}} \rightarrow \hat{\tau}_{\bar{D}} \rightarrow \hat{\tau}_T \rightarrow \tau_T$ where ς is the type of substitution stacks, F and G are the two formulæ in Γ and C and D are the formulæ in Δ . As the formulæ in $\Gamma \cup \Delta$ are Σ_1 , their input type carries no computational content (cf. Lemma 7.3), and we can ignore these formulæ and identify the type above with $\varsigma \rightarrow \hat{\tau}_T \rightarrow \tau_T$, and the term $N_{\pi_\infty}^T ac_{\hat{\tau}_{\bar{F}}} c_{\hat{\tau}_{\bar{G}}} c_{\hat{\tau}_{\bar{C}}} c_{\hat{\tau}_{\bar{D}}}$ with $N_{\pi_\infty}^T a$. Likewise, the type of $N_c^{I_1}$ is assumed to be $\varsigma \rightarrow \hat{\tau}_I \rightarrow \hat{\tau}_{I_1} \rightarrow \tau_{I_1}$ and the type of $N_2^{I_1^+}$ is $\varsigma \rightarrow (\hat{\iota} \rightarrow \iota^1) \rightarrow (\hat{\iota} \rightarrow \iota^1) \rightarrow \hat{\epsilon} \rightarrow \hat{\iota}$.

Other abbreviations and simplifications we utilise are:

- $\langle r \rangle$ for either the sequence $\langle r, \langle \rangle \rangle$ or $\langle r, c_{\iota^1 \rightarrow \epsilon} \rangle$, depending on type, and $\langle r, s \rangle$ as a term of type ι^2 represents $\langle r, s, \langle \rangle \rangle$.
- $\hat{0} = m00$, $1 = s\hat{0}$, $\hat{1} = m01$, $2 = s(m10)$ and $\hat{2} = m02$.
- For each non-terminal N_p^A where A is the i -th formula at position p , an additional non-terminal \hat{N}_p^A with the same arity as N_p^A and associated production rule

$$\hat{N}_p^A ax_0 \cdots x_k \rightarrow N_p^A ax_0 \cdots x_{i-1} x_k x_i \cdots x_{k-1}$$

is included in the Herbrand scheme \mathcal{H}_{π_∞} . These non-terminals ease the computation in derivation steps involving permutation.

- The Herbrand scheme also includes explicit non-terminals for non-determinism at each type, which are represented via set notation: for terms $s_0, \dots, s_k : \rho$ of the same type, the set $S = \{s_i \mid i \leq k\}$ is a term of type ρ with reduction $S \rightarrow s_i$ for each $i \leq k$.
- An equivalence relation \asymp on terms of identical type defined as inducing the same language within all contexts. Formally, we set $r \asymp s$ iff $r \preceq s \preceq r$ where $r \preceq s$ holds just if $r, s : \rho$ and for every $\mathcal{H}_{\pi_\infty} \cup \{x^\rho\}$ -

¹ In [41] π_∞ is formalised as a proof with two Π_2 -cuts but as far as computing Herbrand sets, the two proofs are essentially identical.

term t of basic type (where x is a fresh symbol of type ρ), whenever $t(r/x) \rightarrow^* u$ for a Σ -term u , then $t(s/x) \rightarrow^* v$ for some Σ -term v such that $u^\circ = v^\circ$.

For instance, if $r \rightarrow s$ via an application of a deterministic production rule then $r \asymp s$, and if $S = S'$ are two representations of the same set of terms then $S \asymp S'$. In general, $r(S/x) \not\asymp \{r(s/x) \mid s \in S\}$ as shown by considering $r = \bar{F}x$ with reduction $\bar{F}x \rightarrow^* \text{max}$.² However, suppose $r = \bar{F}t_1 \cdots t_k x$, S is a set of pairs, \bar{F} is deterministic and $\bar{F}t_1 \cdots t_k \langle x, x' \rangle \rightarrow t$. Then $r(S/x) \asymp \{r(s/x) \mid s \in S\} \asymp \{t(\langle u, v \rangle / (x, x')) \mid \langle u, v \rangle \in S\}$.

Finally, we remark that, generalising Lemma 4.10, for every type ρ with co-domain ϵ and every term $r : \rho$, we have $r \asymp c_\rho$.

Language of π_∞ We now compute the language of \mathcal{H}_{π_∞} focusing on the formula T , i.e. set of terms (after evaluation) derivable from the term $\mathbf{N}_{\pi_\infty}^T \perp c_\epsilon$. The first, and only, production rule applicable to this term is given by the cut rule at the root of the proof:

$$\mathbf{N}_{\pi_\infty}^T \perp c_\epsilon \rightarrow \mathbf{N}_0^T \perp (\mathbf{N}_a^I \perp (\hat{\mathbf{N}}_0^{\bar{I}} \perp c_\epsilon)) c_\epsilon. \tag{2}$$

Analysing derivations directly from this term is complicated. As the right sub-proof at 0 culminates in a \forall_γ inference, the external non-terminal $\mathbf{N}_0^{\bar{I}}$ cannot be reduced until its second argument (the term $\mathbf{N}_a^I \perp (\hat{\mathbf{N}}_0^{\bar{I}} \perp c_\epsilon)$) is reduced to an explicit pair. But the inference at a in the left sub-proof is a contraction, so this immediately introduces non-determinism and duplication of arguments. After resolving the non-determinism and reducing the two continuations of \mathbf{N}_a^I to pairs (say in terms of $\mathbf{N}_d^{I_0} / \mathbf{N}_d^{I_1}$), the external non-terminal can be reduced. The argument $\hat{\mathbf{N}}_0^{\bar{I}} \perp c_\epsilon$ comes into play at this point: the productions for $\mathbf{N}_b^{I^+}$ and \mathbf{N}_c^I increase the nesting of non-terminals which must also be evaluated as pairs in order to proceed beyond $\mathbf{N}_d^{I_0} / \mathbf{N}_d^{I_1}$.

In the following, we compute the language via a top-down approach, analysing derivations starting from relatively simple terms, and building these together to compute the language of more complex interactions between non-terminals. We begin with the most simple derivations available. Recall that $\hat{\tau}_{I_0} = \hat{\tau}_{I_1} = \iota^1$. Concerning non-terminals from the left sub-proof we have the following derivation starting from $\mathbf{N}_d^{I_0} / \mathbf{N}_d^{I_1}$.

$$\begin{aligned} \mathbf{N}_d^{I_i} a \langle r \rangle \langle s \rangle &\rightarrow \mathbf{N}_e^{2+i} ([\hat{\alpha} \leftarrow s] a) \langle r \rangle \langle \rangle \\ &\rightarrow \mathbf{N}_f^{2+i} ([\alpha \leftarrow r][\hat{\alpha} \leftarrow s] a) \langle \rangle \langle \rangle \\ &\rightarrow^* \langle \text{ma} \hat{\alpha} \cdot [\alpha \leftarrow r][\hat{\alpha} \leftarrow s] a \rangle \end{aligned}$$

(Note, as $|\Gamma| = 2$ the $(2+i)$ -th formula at positions e and f is the ancestor of I_i from d .) If r happens to be such that $r \cdot [\hat{\alpha} \leftarrow s] a \asymp r \cdot a$, then since the derivation above follows deterministic reductions only, we deduce

$$\mathbf{N}_d^{I_i} a \langle r \rangle \langle s \rangle \asymp \langle \text{mrs} \cdot a \rangle. \tag{3}$$

Examining the non-terminals from lower in the left sub-proof affords us

$$\begin{aligned} \mathbf{N}_c^I a r s &\rightarrow \langle 0, \hat{\mathbf{N}}_d^{I_0} a s \rangle & \mathbf{N}_c^{I_1} a r s &\rightarrow^* \mathbf{N}_d^{I_1} a (r \langle 0, \hat{\mathbf{N}}_d^{I_0} a s \rangle) s \\ \mathbf{N}_b^{I^+} a r s &\rightarrow \langle 1, \mathbf{N}_c^{I_1} a r \rangle & \mathbf{N}_b^I a r s &\rightarrow \mathbf{N}_c^I a r (s (\mathbf{N}_b^{I^+} a r s)) \\ & & &\rightarrow^* \langle 0, \hat{\mathbf{N}}_d^{I_0} a (s \langle 1, \mathbf{N}_c^{I_1} a r \rangle) \rangle \end{aligned}$$

² Formally, we require an \mathcal{H}_{π_∞} analogue of \bar{F} but this is not difficult to find.

The derivation from $\mathbf{N}_c^{I_1}ars$ can be continued provided that the two arguments of $\mathbf{N}_d^{I_1}$, namely $r\langle 0, \hat{\mathbf{N}}_d^{I_0}as \rangle$ and s , are reducible to pairs. Thus if $r\langle 0, \hat{\mathbf{N}}_d^{I_0}a\langle s \rangle \rangle \rightarrow^* \langle r_0, r'_0 \rangle$ and $r_0 \cdot [\hat{a} \leftarrow s]a \asymp r_0 \cdot a$ then

$$\mathbf{N}_c^{I_1}ar\langle s \rangle \rightarrow^* \mathbf{N}_d^{I_1}a\langle r_0 \rangle\langle s \rangle \asymp \langle mr_0s \cdot a \rangle.$$

Because the reductions governing $\mathbf{N}_c^{I_1}$ are all deterministic, when phrased in terms of equivalences, this becomes

$$\left. \begin{array}{l} r\langle 0, \hat{\mathbf{N}}_d^{I_0}a\langle s \rangle \rangle \asymp \{\langle r_i \rangle \mid i \leq k\} \\ r_i \cdot [\hat{a} \leftarrow s]a \asymp r_i \cdot a \text{ each } i \leq k \end{array} \right\} \text{implies } \mathbf{N}_c^{I_1}ar\langle s \rangle \asymp \{\langle mr_i s \cdot a \rangle \mid i \leq k\} \quad (4)$$

Property (4) will be useful later.

Returning briefly to the derivations from non-terminals \mathbf{N}_b^I and $\mathbf{N}_b^{I^+}$ started earlier, each of these derivations is also deterministic, so therefore

$$\mathbf{N}_a^Iar \asymp \{\mathbf{N}_b^Iarr, \mathbf{N}_b^{I^+}arr\} \asymp \{\langle 0, \hat{\mathbf{N}}_d^{I_0}a(r\langle 1, \mathbf{N}_c^{I_1}ar \rangle) \rangle, \langle 1, \mathbf{N}_c^{I_1}ar \rangle\} \quad (5)$$

which provides the first step in the continuation of the derivation from $\mathbf{N}_{\pi_\infty}^T$. Before extending (2) however we consider some simple derivations arising from the right sub-proof. On this side, the alternation of universal and existential inference rules means that few non-terminals can be adequately analysed in isolation as we did above. Most straightforward are non-terminals $\mathbf{N}_4^{\bar{I}\gamma}$ and $\mathbf{N}_2^{\bar{I}\gamma}$, for which we have

$$\mathbf{N}_4^{\bar{I}\gamma}arst \asymp \langle s\beta \cdot a \rangle \quad \mathbf{N}_2^{\bar{I}\gamma}arst \asymp \langle 0 \cdot a \rangle \asymp \langle 0 \rangle$$

This gives rise to, for example,

$$\mathbf{N}_3^{\bar{I}\gamma}a\langle r_0 \rangle st \asymp \mathbf{N}_4^{\bar{I}\gamma}([\beta \leftarrow r_0]a)\langle \rangle st \asymp \langle sr_0 \cdot a \rangle$$

and hence if $r : (\iota \times (\iota^1 \rightarrow \epsilon)) \rightarrow \iota^1$ is a term such that $r\langle 0 \rangle \asymp \{\langle r_i \rangle \mid i \leq k\}$ then also

$$\begin{aligned} \mathbf{N}_2^{\bar{I}\gamma^+}arst &\asymp \mathbf{N}_3^{\bar{I}\gamma}a(r(\mathbf{N}_2^{\bar{I}\gamma}arst))st \\ &\asymp \mathbf{N}_3^{\bar{I}\gamma}a\{\langle r_i \rangle \mid i \leq k\}st \\ &\asymp \{\langle sr_i \cdot a \rangle \mid i \leq k\} \end{aligned}$$

The equivalences for $\mathbf{N}_2^{\bar{I}\gamma}$ and $\mathbf{N}_2^{\bar{I}\gamma^+}$ combine to yield, given the same r ,

$$\mathbf{N}_1^{\bar{I}\gamma}art \asymp \{\mathbf{N}_2^{\bar{I}\gamma}arrt, \mathbf{N}_2^{\bar{I}\gamma^+}arrt\} \asymp \{\langle 0 \rangle\} \cup \{\langle sr_i \cdot a \rangle \mid i \leq k\}. \quad (6)$$

In particular, choosing $r = \hat{\mathbf{N}}_d^{I_0} \perp \langle s \rangle$, this implies

$$\mathbf{N}_1^{\bar{I}\gamma}a(\hat{\mathbf{N}}_d^{I_0} \perp \langle s \rangle)t \asymp \{\langle 0 \rangle, \langle s(m0s) \cdot a \rangle\} \quad (7)$$

which will be needed later. In addition to (7), it is necessary to analyse the complex term $\mathbf{N}_1^{\bar{I}\gamma}a(\mathbf{N}_c^{I_1} \perp (\hat{\mathbf{N}}_0^{\bar{I}} \perp c_\epsilon))c_\epsilon$. However, here we can use (6) again. If $\delta : \iota$ is a fresh symbol then, applying (7) and (4) (using $r = \mathbf{N}_0^{\bar{I}} \perp c_\epsilon$), we get

$$\mathbf{N}_0^{\bar{I}} \perp \langle 0, \hat{\mathbf{N}}_d^{I_0} \perp \langle \delta \rangle \rangle \mathbf{c}_{\hat{\epsilon}} \asymp \mathbf{N}_1^{\bar{I}\gamma} ([\gamma \leftarrow 0] \perp) (\hat{\mathbf{N}}_d^{I_0} \perp \langle \delta \rangle) \mathbf{c}_{\hat{\epsilon}} \asymp \{ \langle 0 \rangle, \langle \mathbf{s}(\mathbf{m}0\delta) \rangle \} \quad (8)$$

$$\mathbf{N}_c^{I_1} \perp (\hat{\mathbf{N}}_0^{\bar{I}} \perp \mathbf{c}_{\hat{\epsilon}}) \langle \delta \rangle \asymp \{ \langle \mathbf{m}0\delta \rangle, \langle \mathbf{m}(\mathbf{s}(\mathbf{m}0\delta))\delta \rangle \} \quad (9)$$

whence (6) implies

$$\mathbf{N}_1^{\bar{I}\gamma} a (\mathbf{N}_c^{I_1} \perp (\hat{\mathbf{N}}_0^{\bar{I}} \perp \mathbf{c}_{\hat{\epsilon}})) \mathbf{c}_{\hat{\epsilon}} \asymp \{ \langle 0 \rangle, \langle 1 \rangle, \langle 2 \rangle \}. \quad (10)$$

We have still not examined derivations starting from the non-terminals \mathbf{N}_0^T , \mathbf{N}_1^T , and \mathbf{N}_i^T for $i \geq 2$, which will arise in the computation of $L(\pi_\infty)$. The first three non-terminals behave according to

$$\begin{aligned} \mathbf{N}_0^T a \langle r, s \rangle t &\asymp \mathbf{N}_1^T ([\gamma \leftarrow r] a) s t \\ &\asymp \mathbf{N}_2^T ([\gamma \leftarrow r] a) s s t \\ &\asymp \mathbf{N}_3^T ([\gamma \leftarrow r] a) (\mathbf{s}(0)) s t \end{aligned}$$

The remaining behave similarly to the \mathbf{N}_i^A non-terminals analysed earlier, except that it is \mathbf{N}_6^T that provides the only ‘outputs’ in the derivation. In particular,

$$\begin{aligned} \mathbf{N}_6^T a r s &\asymp \langle \beta, \hat{\beta} \rangle \cdot a & \mathbf{N}_5^T a r \langle s_0 \rangle t &\asymp \langle \beta, s_0 \rangle \cdot a \\ \mathbf{N}_3^T a \langle r_0 \rangle s t &\asymp \mathbf{N}_5^T ([\beta \leftarrow r_0] a) \langle \rangle (\mathbf{s} \langle s r_0 \cdot a \rangle) t \end{aligned}$$

Let $\delta : \iota$ be a fresh symbol. Combining the two sets of equations above, if $s : \hat{\iota} \rightarrow \iota^1$ is such that $s \langle \delta \rangle \asymp \{ \langle s_i \rangle \mid i \leq k \}$ and s_i contains neither β or γ for each i , it follows that

$$\mathbf{N}_0^T \perp \langle r, s \rangle t \asymp \{ \langle s_i \cdot [\delta \leftarrow 0] \perp, s_j \cdot [\delta \leftarrow \mathbf{s} s_i] [\delta \leftarrow 0] \perp \mid i, j \leq k \}. \quad (11)$$

We can now proceed with calculating the language of $\mathbf{N}_{\pi_\infty}^T \perp \mathbf{c}_{\hat{\epsilon}}$. Let $w = \hat{\mathbf{N}}_0^{\bar{I}} \perp \mathbf{c}_{\hat{\epsilon}}$. Following on from (2) and (5) we have

$$\begin{aligned} \mathbf{N}_{\pi_\infty}^T \perp \mathbf{c}_{\hat{\epsilon}} &\asymp \mathbf{N}_0^T \perp (\mathbf{N}_a^I \perp w) \mathbf{c}_{\hat{\epsilon}} \\ &\asymp \left\{ \mathbf{N}_0^T \perp \langle 0, \hat{\mathbf{N}}_d^{I_0} \perp (w \langle 1, \mathbf{N}_c^{I_1} \perp w \rangle) \rangle \mathbf{c}_{\hat{\epsilon}}, \mathbf{N}_0^T \perp \langle 1, \mathbf{N}_c^{I_1} \perp w \rangle \mathbf{c}_{\hat{\epsilon}} \right\} \end{aligned} \quad (12)$$

Thus, we need only compute

$$\mathbf{N}_d^{I_0} \perp \langle \delta \rangle (w \langle 1, \mathbf{N}_c^{I_1} \perp w \rangle) \quad \mathbf{N}_c^{I_1} \perp (\hat{\mathbf{N}}_0^{\bar{I}} \perp \mathbf{c}_{\hat{\epsilon}}) \langle \delta \rangle$$

and apply (11) (assuming that the terms obtained will be free of β and γ). The latter was already established in (9):

$$\mathbf{N}_c^{I_1} \perp (\hat{\mathbf{N}}_0^{\bar{I}} \perp \mathbf{c}_{\hat{\epsilon}}) \langle \delta \rangle \asymp \{ \langle \mathbf{m}0\delta \rangle, \langle \mathbf{m}(\mathbf{s}(\mathbf{m}0\delta))\delta \rangle \}$$

For the former, we have $w \langle 1, \mathbf{N}_c^{I_1} \perp w \rangle \asymp \mathbf{N}_1^{\bar{I}\gamma} ([\gamma \leftarrow 1] \perp) (\mathbf{N}_c^{I_1} \perp w) \mathbf{c}_{\hat{\epsilon}}$, whence (10) implies

$$\begin{aligned} \mathbf{N}_d^{I_0} \perp \langle \delta \rangle (w \langle 1, \mathbf{N}_c^{I_1} \perp w \rangle) &\asymp \mathbf{N}_d^{I_0} \perp \langle \delta \rangle (\mathbf{N}_1^{\bar{I}\gamma} ([\gamma \leftarrow 1] \perp) (\mathbf{N}_c^{I_1} \perp w) \mathbf{c}_{\hat{\epsilon}}) \\ &\asymp \left\{ \mathbf{N}_d^{I_0} \perp \langle \delta \rangle \langle s \rangle \mid s \in \{0, 1, 2\} \right\} \\ &\asymp \{ \langle \mathbf{m}\delta 0 \rangle, \langle \mathbf{m}\delta 1 \rangle, \langle \mathbf{m}\delta 2 \rangle \} \end{aligned}$$

Hence, by (11) and (12), we deduce

$$\begin{aligned} \mathbf{N}_{\pi_\infty}^T \perp \mathbf{c}_\varepsilon \asymp & \left\{ \langle r, s^{[\delta \leftarrow \mathbf{s}r]} \rangle \mid r \in \{\hat{0}, \hat{1}, \hat{2}\}, s \in \{\mathbf{m}\delta 0, \mathbf{m}\delta 1, \mathbf{m}\delta 2\} \right\} \\ & \cup \left\{ \langle r, s^{[\delta \leftarrow \mathbf{s}r]} \rangle \mid r \in \{\hat{0}, \mathbf{m}10\}, s \in \{\mathbf{m}0\delta, \mathbf{m}(\mathbf{s}(\mathbf{m}0\delta))\delta\} \right\} \end{aligned}$$

Under the standard interpretation of the symbols 0, \mathbf{s} and \mathbf{m} (as zero, successor and binary ‘max’) $L(\pi_\infty)$ ascribes to T the set

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

6. Substitution, subsumption and normality

In the previous section we introduced a preorder on terms, $r \preceq s$, specifying that the language induced by s extends the language induced by r . Formally, $r \preceq s$ holds if r and s are of the same type ρ and for every $\mathcal{H} \cup \{x^\rho\}$ -term t of basic type, if $t(r/x) \rightarrow^* r_0$ for a Σ -term r_0 then there exists a Σ -term s_0 such that $t(s/x) \rightarrow^* s_0$ and $r_0^\circ = s_0^\circ$.

This relation can be extended to proofs in a natural way, by defining $\pi' \preceq \pi$ if π and π' have the same end-sequent and $\mathbf{N}_{\pi'}^i \preceq \mathbf{N}_\pi^i$ for each i . For many one-step cut reductions $\pi \rightsquigarrow \pi'$ indeed $\pi' \preceq \pi$ (and even $\pi \preceq \pi'$), from which it immediately follows that $L(\pi') \subseteq L(\pi)$ (resp. $L(\pi') = L(\pi)$). However, there exist reductions $\pi \rightsquigarrow \pi'$ for which $L(\pi') \subseteq L(\pi)$ but $\pi' \not\preceq \pi$. These scenarios all arise in reductions which interact with quantifiers and alter the contexts in which explicit substitutions occur in derivations. In order to prove language preservation, i.e. that $\pi \rightsquigarrow \pi'$ implies $L(\pi') \subseteq L(\pi)$, for these reductions, we will replace the preorder \preceq with a coarser relation which we call *subsumption*, that quantifies not over *all* possible contexts (the ‘ t ’ in the definition of $r \preceq s$ above) but only over contexts of a particular syntactic shape. Such terms we name *normal terms* and will be defined (along with our relation of term subsumption) in Section 6.1 below.

Before embarking on these definitions, it will be convenient to abstract the notion of Herbrand scheme slightly and observe that we can specify a single, *universal*, recursion scheme in which every Herbrand scheme for a regular proof can be viewed as a natural finite sub-scheme.

Definition 6.1 (*Universal Herbrand scheme*). Let Σ be the signature of first-order logic. We let \mathcal{H} denote the infinite recursion scheme comprising:

1. a non-deterministic non-terminal $D_\rho : \rho \rightarrow \rho \rightarrow \rho$ for each basic type ρ with production rules $D_\rho r s \rightarrow r$ and $D_\rho r s \rightarrow s$,
2. all non-terminals \mathbf{N}_π^i , $\mathbf{S}_{\pi,i}$ and \mathbf{c}_ρ from Definition 4.5 with their associated production rules formulated deterministically in terms of the D_ρ non-terminals above,
3. for each non-terminal $\mathbf{N}_\pi^i : \varsigma \rightarrow \tau_0 \rightarrow \dots \tau_m \rightarrow \tau$ from the above with τ prime, a non-terminal $\hat{\mathbf{N}}_\pi^i$ with type and associated production rule

$$\begin{aligned} \hat{\mathbf{N}}_\pi^i : \varsigma \rightarrow \tau_0 \rightarrow \dots \rightarrow \tau_{i-1} \rightarrow \tau_{i+1} \rightarrow \dots \tau_m \rightarrow \tau_i \rightarrow \tau \\ \hat{\mathbf{N}}_\pi^i a x_0 \dots x_m \rightarrow \mathbf{N}_\pi^i a x_0 \dots x_{i-1} x_m x_i \dots x_{m-1} \end{aligned}$$

We refer to \mathcal{H} as the *universal Herbrand scheme*.

Henceforth, a *term* is an \mathcal{H} -term and we write \rightarrow in place of $\rightarrow_{\mathcal{H}}$. Finite sets of \mathcal{H} -terms will represent applications of the non-deterministic non-terminals D_ρ . Specifically, the set $\{s_0^\rho, \dots, s_k^\rho\}$ represents any term

formed by combining all the terms s_0, \dots, s_k (possibly with repetitions) via the non-terminal D_ρ . If S is a finite set of terms of the same type, it follows that $S \rightarrow^* s$ for each $s \in S$.

Notice that there are no start symbols in \mathcal{H} . In this regard we may consider the individual Herbrand scheme \mathcal{H}_π as obtained from \mathcal{H} by specifying an appropriate set of start symbols. The new ‘hat’ non-terminals do not play a role in viewing \mathcal{H} as a universal Herbrand scheme. Rather, they become useful in ‘transferring’ non-terminals lacking their final argument through applications of permutation. For example, the following partial proof (where we assume $u(A) > 0$) gives rise to the production rules on the right:

$$\begin{array}{c} \frac{\pi_1 \vdash A(r/v), B}{\pi_0 \vdash B, A(r/v)} \quad \text{p} \\ \frac{\pi_0 \vdash B, A(r/v)}{\pi \vdash B, \exists v A} \quad \exists_r \end{array} \quad \begin{array}{l} \mathbf{N}_{\pi_0}^i axz \rightarrow \mathbf{N}_{\pi_1}^{1-i} axz \\ \mathbf{N}_{\pi_0}^1 axz \rightarrow \langle r \cdot a, \mathbf{N}_{\pi_0}^1 x \rangle \\ \mathbf{N}_{\pi_0}^0 axz \rightarrow \mathbf{N}_{\pi_0}^0 ax(z(\mathbf{N}_{\pi_0}^1 axz)) \end{array}$$

yielding the derivation $\mathbf{N}_{\pi_0}^0 axz \rightarrow^* \mathbf{N}_{\pi_1}^1 a(z\langle r \cdot a, \mathbf{N}_{\pi_0}^1 ax \rangle)x$. The derivation cannot be extended as it stands because $\mathbf{N}_{\pi_0}^1$ lacks an argument, meaning that it is not formally possible to express the term $\mathbf{N}_{\pi_0}^0 axz$ by reference to the proof π_1 only without instantiating x and z by concrete terms. However, $\mathbf{N}_{\pi_0}^1 ax$ is extensionally equal to the term $\hat{\mathbf{N}}_{\pi_1}^0 ax$, allowing us to equate $\mathbf{N}_{\pi_0}^0 axz$ with the term $\mathbf{N}_{\pi_1}^1 a(z\langle r \cdot a, \hat{\mathbf{N}}_{\pi_1}^0 ax \rangle)x$ for any choice of a, x and z . Equations such as these are useful in the close examination of the cut elimination process carried out in the sections below.

In the previous section a natural subsumption and equivalence relation on terms was introduced given by equating terms that induce the same language in all contexts. In the context of the universal Herbrand scheme \mathcal{H} , this subsumption is given by $r \preceq s$ which holds just if $r, s : \rho$ for some ρ and, for every $\mathcal{H} \cup \{x^\rho\}$ -term t of basic type, whenever $t(r/x) \rightarrow^* r_0$ for a Σ -term r_0 , then $t(s/x) \rightarrow^* s_0$ for some Σ -term s_0 such that $r_0^\circ = s_0^\circ$. The corresponding equivalence relation \asymp is defined by $r \asymp s$ iff $r \preceq s \preceq r$. The following properties of the relations \preceq and \asymp were remarked in the last section.

Lemma 6.2. *Let $r, s : \rho$ be \mathcal{H} -terms of the same type and S a finite set of terms of pair type $\sigma = \sigma_0 \times \dots \times \sigma_l$.*

1. *If $r \rightarrow s$ then $s \preceq r$. If, in addition, the reduction follows from a production rule for a deterministic non-terminal then $r \asymp s$.*
2. *If r and s are representations of the same finite set of \mathcal{H} -terms then $r \asymp s$.*
3. *If $r = \mathbf{F}s_0 \dots s_{k-1}x^\sigma$ for a non-terminal \mathbf{F} with production rule $\mathbf{F}x_0 \dots x_{k-1}\langle x_k, \dots, x_{k+l} \rangle \rightarrow t$ then*

$$\begin{aligned} r(S/x) &\asymp \{r(s/x) \mid s \in S\} \\ &\asymp \{t((s_0 \dots, s_{k+l})/(x_0, \dots, x_{k+l})) \mid \langle s_k, \dots, s_{k+l} \rangle \in S\}. \end{aligned}$$

4. *If the co-domain of ρ is ϵ then $r \asymp c_\rho$.*

Proof. Properties 1–3 are straight-forward, though for 3 we note that only deterministic non-terminals have production rules that invoke pattern-matching. 4 generalises Lemma 4.10 and is proved by induction on ρ and r , noting that $\langle \rangle \cdot a \asymp \langle \rangle$ for any substitution a . \square

6.1. Normal terms and subsumption

In order to focus on the impact of substitutions in \mathcal{H} -terms it is necessary to introduce a notion of free and bound occurrences of Σ -symbols in these terms where, recall, Σ is the signature of first-order logic. The free symbols of a basic Σ -term are simply the Σ -symbols that occur in the term; Σ -terms have no bound symbols. For a basic \mathcal{H} -term t , the free symbols of t are the Σ -symbols occurring in t combined with the

Σ -symbols occurring in any proof π for which a non-terminal \mathbf{N}_π^i or $\hat{\mathbf{N}}_\pi^i$ appears in t ; the *bound* symbols of t are the eigenvariables of the proofs which occur leftmost in t . For non-basic terms, substitutions and substitution stacks are interpreted as contributing to the set of bound symbols, and limiting the set of free symbols in the natural way. Explicitly, for a substitution stack $a : \varsigma$ and \mathcal{H} -term $r : \rho$ we define

$$\begin{aligned}
Bd(r) &= \begin{cases} \emptyset, & \text{if } r \text{ is a } \Sigma\text{-term or } r = \mathbf{c}_\rho, \\ EV(\pi), & \text{if } r = \mathbf{N}_\pi^i \text{ or } r = \hat{\mathbf{N}}_\pi^i \text{ for some } i, \\ Bd(s) \cup Bd(t), & \text{if } r = \langle s, t \rangle, \\ Bd(s) \cup Bd(t), & \text{if } r = \mathbf{D}_\rho st, \\ Bd(s) \cup Bd(a), & \text{if } r = sa \text{ or } r = s \cdot a \text{ for } a : \varsigma, \\ Bd(s), & \text{if } r = st \text{ and } t : \tau \text{ where } \tau \neq \varsigma, \end{cases} \\
Bd(a) &= \begin{cases} \emptyset, & \text{if } a = \perp, \\ \{\alpha\} \cup Bd(b), & \text{if } a = [\alpha \leftarrow s]b, \end{cases} \\
Fr(r) &= \begin{cases} \emptyset, & \text{if } r = \mathbf{c}_\rho \text{ or } r = \mathbf{D}_{\rho'}, \\ \{r\}, & \text{if } r \in \Sigma, \\ Fr(\pi), & \text{if } r = \mathbf{N}_\pi^i \text{ or } r = \hat{\mathbf{N}}_\pi^i \text{ for some } i, \\ Fr(s) \cup Fr(t), & \text{if } r = \langle s, t \rangle, \\ (Fr(s) \setminus Bd(a)) \cup Fr(a), & \text{if } r = sa \text{ or } r = s \cdot a \text{ for } a : \varsigma, \\ Fr(s) \cup Fr(t), & \text{if } r = st \text{ and } t : \tau \text{ where } \tau \neq \varsigma, \end{cases} \\
Fr(a) &= \begin{cases} \emptyset, & \text{if } a = \perp, \\ (Fr(s) \setminus Bd(b)) \cup Fr(b), & \text{if } a = [\alpha \leftarrow s]b, \end{cases}
\end{aligned}$$

$EV(\pi)$ denotes the set of eigenvariables in the proof π , and $Fr(\pi)$ the set of all non-eigenvariable Σ -symbols occurring in the π . Notice that $Bd(\mathbf{N}_\pi^i)$ and $Fr(\mathbf{N}_\pi^j)$ are disjoint sets by definition.

Definition 6.3 (*Normal terms*). A *normal term* is an \mathcal{H} -term r satisfying:

1. if a is substitution stack which is a sub-term of r then $Bd(a) \cap Fr(a) = \emptyset$,
2. if st is an application which is a sub-term of r then $Bd(s) \cap Fr(t) = \emptyset$,
3. if $s \cdot a$ is a substitution occurring as a sub-term of r then s is of simple type.

As mentioned at the beginning of this section, the aim of the above definition is to provide a class of terms for which we can examine a more refined subsumption relation on \mathcal{H} -terms that captures both language inclusion and equality for a wide range of cut reduction rules. The subsumption relation that achieves this is essentially the restriction of \preceq that only quantifies over normal contexts.

Definition 6.4 (*Subsumption*). Given normal \mathcal{H} -terms $r, s : \rho$ of the same type, s *subsumes* r , in symbols $r \sqsubset s$, just if, for every $\mathcal{H} \cup \{x^\rho\}$ -term t of basic type such that $t(r/x)$ and $t(s/x)$ are both normal, whenever $t(r/x) \rightarrow^* u$ for a Σ -term u then $t(s/x) \rightarrow^* v$ for some Σ -term v satisfying $u^\circ = v^\circ$. Define $r \sim s$ if $r \sqsubset s$ and $s \sqsubset r$.

Clearly, for normal terms r and s , $r \preceq s$ implies $r \sqsubset s$, and $r \succ s$ implies $r \sim s$. Hence, if π, π' are two regular proofs of a Σ_1 sequent Γ and $\mathbf{S}_{\pi',i} \sqsubset \mathbf{S}_{\pi,i}$ for every $i < |\Gamma|$ then $L(\pi') \subseteq L(\pi)$. However, what we require is the more general property that if for every \mathcal{H} -derivation $\mathbf{S}_{\pi',i} \rightarrow^* u$ of a Σ -term there exists

\mathcal{H} -terms r , s and t such that $S_{\pi',i} \rightarrow^* t(r/x) \rightarrow^* u$, $S_{\pi,i} \rightarrow^* t(s/x)$ and $r \sqsubset s$, then we may conclude $L(\pi') \subseteq L(\pi)$. This result holds trivially for \preceq in place of \sqsubset . For it to work for subsumption, the terms r , s and t must all be normal, i.e., we require

Lemma 6.5. *If $r \rightarrow s$ and r is normal then s is normal. In particular, if $S_{\pi,i} \rightarrow^* s$ then s is a normal term.*

Lemma 6.5 is not difficult to establish but requires some technical observations concerning the preservation of free and bound symbols through \mathcal{H} -derivations. The proof is given in Section 6.2 below.

To accompany Lemma 6.5 it is necessary to know that every maximal derivation from a start symbol terminates in a Σ -term. By the previous lemma, it would suffice to show an arbitrary normal term of basic type is either reducible or a Σ -term but this claim is easily seen to be false. In place of the general statement we have the next two lemmas.

Lemma 6.6 (Finite basis lemma). *For every normal \mathcal{H} -term r of pair type there exists terms $\langle s_0, t_0 \rangle, \dots, \langle s_k, t_k \rangle$ such that $r \sim \{ \langle s_i, t_i \rangle \mid i \leq l \}$.*

Lemma 6.7. *Suppose $A \in \Sigma_1$. If $r : \tau_A$ is a normal term and not a Σ -term then $r \rightarrow s$ for some term s .*

There are two scenarios in which explicit substitutions can block derivations. In the first, there is a term r of pair type which cannot be reduced to an explicit pair because it has the form, say, $s \cdot a$. Even if r itself is a Σ -term, i.e., does not contain any non-terminals, it may appear as the argument of a non-terminal whose reduction depends on pattern-matching r against an explicit pair. The second scenario is if there exists a sub-term of the form $Fr_1 \cdots r_k \cdot a$ where F is a non-terminal of arity greater than k . Normality rules out both scenarios. This is the main idea behind Lemmas 6.6 and 6.7.

6.2. Proofs of Lemmas 6.5–6.7

We now prove the three lemmas stated above. The arguments rely on a number of technical details concerning derivations of normal terms. Following, we examine the interaction of subsumption and explicit substitutions which will prove important for establishing language preservation for the case quantifier inferences. At this point the reader may wish to proceed directly to Section 7 and refer back the technical results as needed.

We begin with Lemma 6.5 which requires two technical observations on free and bound symbols in normal terms. Their effect is to reduce the problem of proving the lemma to checking that each production rule of the universal Herbrand scheme preserve normality.

Lemma 6.8. *If $r(s/x)$ is a normal term, t is a normal term of the same type as s , $Fr(t) \subseteq Fr(s)$ and $Bd(t) \subseteq Bd(s)$ then $r(t/x)$ is normal.*

Proof. By definition. \square

Lemma 6.9. *If $Fx_0 \cdots x_{k-1} \langle x_k, \dots, x_{k+l} \rangle \rightarrow t$ is a production rule of \mathcal{H} , and r_0, \dots, r_{k+l} are such that $s = Fr_0 \cdots r_{k-1} \langle r_k, \dots, r_{k+l} \rangle$ is normal, then $Fr(t(\vec{r}/\vec{x})) \subseteq Fr(s)$ and $Bd(t(\vec{r}/\vec{x})) \subseteq Bd(s)$.*

Proof. We examine two particular cases, namely the quantifier rules, and leave the remaining for the reader to check. Consider an instance of the production rule for \forall_α for a single eigenvariable:

$$N_\pi^i ar_0 \cdots r_{k-1} \langle s, r_k \rangle \rightarrow N_{\pi_0}^i ([\alpha \leftarrow s]a)r_0 \cdots r_k$$

where r_0, \dots, r_k and s are terms of suitable type and number, and $a : \varsigma$ is a substitution stack. Let m and n abbreviate the left- and righthand term in the above rule respectively. Assume m is a normal term.

$$\begin{aligned} Bd(m) &= Bd(\mathbf{N}_{\pi}^i a) \\ &= EV(\pi) \cup Bd(a) \\ &= EV(\pi_0) \cup \{\alpha\} \cup Bd(a) \\ &= Bd(n) \end{aligned}$$

Concerning free symbols, we have

$$\begin{aligned} Fr(m) &= (Fr(\pi) \setminus Bd(a)) \cup Fr(a) \cup Fr(r_0, \dots, r_k, s) \\ Fr(n) &= (Fr(\pi_0) \setminus Bd([\alpha \leftarrow s]a)) \cup Fr([\alpha \leftarrow s]a) \cup Fr(r_0, \dots, r_k) \\ &= (Fr(\pi_0) \setminus (\{\alpha\} \cup Bd(a))) \cup (Fr(s) \setminus Bd(a)) \cup Fr(a) \cup Fr(r_0, \dots, r_k) \end{aligned}$$

where $Fr(u_0, \dots, u_l) = \bigcup_{i \leq l} Fr(u_i)$. By normality of m , $Fr(s) \setminus Bd(a) = Fr(s)$ and as $Fr(\pi) = Fr(\pi_0) \setminus \{\alpha\}$, so $Fr(n) \subseteq Fr(m)$.

For production rules resulting from the inference rule \exists_s , suppose

$$\mathbf{N}_{\pi}^i ar_0 \cdots r_k \rightarrow \mathbf{N}_{\pi_0}^i ar_0 \cdots r_{k-1} (r_k \langle s \cdot a, \mathbf{N}_{\pi_0}^k ar_0 \cdots r_{k-1} \rangle) \quad (13)$$

for suitable terms r_0, \dots, r_k and a . Recall that s is the Σ -term instantiating the existential quantifier in the active formula of π_0 . By our regularity condition on proofs, $Fr(s) \subseteq Fr(\pi)$, so, letting m and n denote the left and right side of the reduction in (13), we have

$$\begin{aligned} Fr(n) &= Fr(\mathbf{N}_{\pi_0}^i a) \cup Fr(r_0, \dots, r_k) \cup Fr(s \cdot a) \cup Fr(\mathbf{N}_{\pi_0}^k a) \\ &\subseteq Fr(\mathbf{N}_{\pi}^i a) \cup Fr(r_0, \dots, r_k) \\ &= Fr(m). \quad \square \end{aligned}$$

We can now prove Lemma 6.5.

Proof of Lemma 6.5. By the previous two lemmas it suffices to show that every production rule of \mathcal{H} locally preserves normality. As in the proof of Lemma 6.9, we offer the argument for the important cases of the two quantifier rules and leave the remaining cases to the reader. Let the derivation

$$\mathbf{N}_{\pi}^i ar_1 \cdots r_k \langle s, r_{k+1} \rangle \rightarrow \mathbf{N}_{\pi_0}^i ([\alpha \leftarrow s]a) r_1 \cdots r_k r_{k+1}$$

arise from an inference \forall_{α} . Let m and n denote respectively the left and righthand term of the above equation. Assume m is normal. In particular,

$$(EV(\pi) \cup Bd(a)) \cap (Fr(a) \cup Fr(s)) = \emptyset \quad (14)$$

We first show that for every application $s't'$ occurring in n , $Bd(s') \cap Fr(t') = \emptyset$. This is evident if $s't'$ is a sub-term of a, r_1, \dots, r_k, s or t . Moreover, it holds for the case $s' = \mathbf{N}_{\pi_0}^i$ and $t' = [\alpha \leftarrow s]a$ because

$$\begin{aligned} Bd(s') \cap Fr(t') &= EV(\pi_0) \cap ((Fr(s) \setminus Bd(a)) \cup Fr(a)) \\ &\subseteq EV(\pi) \cap (Fr(s) \cup Fr(a)) \\ &= \emptyset \end{aligned}$$

and for $s' = N_{\pi_0}^i([\alpha \leftarrow s]a)r_1 \cdots r_j$, $t' = r_{j+1}$ ($j \leq k$) because $Bd(N_{\pi_0}^i([\alpha \leftarrow s]a)) \subseteq Bd(N_{\pi}^i a)$. The other requirement to check for normality is that the sets $Bd([\alpha \leftarrow s]a)$ and $Fr([\alpha \leftarrow s]a)$ are disjoint, but this follows from (14), given that $Fr([\alpha \leftarrow s]a) \subseteq Fr(s) \cup Fr(a)$ and $Bd([\alpha \leftarrow s]a) \subseteq EV(\pi) \cup Bd(a)$. Hence n is normal.

The second production rule is the one arising from the inference \exists_s :

$$N_{\pi}^i ar_0 \cdots r_k \rightarrow N_{\pi_0}^i ar_0 \cdots r_{k-1} (r_k \langle s \cdot a, N_{\pi_0}^k ar_0 \cdots r_{k-1} \rangle)$$

Suppose $N_{\pi}^i ar_0 \cdots r_k$ is a normal term so, in particular, $Bd(N_{\pi}^i a)$ is disjoint from $Fr(r_i)$ for each $i \leq k$. In this case it suffices to show

$$Bd(N_{\pi_0}^i ar_0 \cdots r_{k-1}) \cap Fr(r_k \langle s \cdot a, N_{\pi_0}^k ar_0 \cdots r_{k-1} \rangle) = \emptyset$$

i.e., that

$$Bd(N_{\pi_0}^i a) \cap \left(\bigcup_{i \leq k} Fr(r_i) \cup Fr(s \cdot a) \cup Fr(N_{\pi_0}^k a) \right) = \emptyset,$$

as all other cases follow immediately from normality of $N_{\pi}^i ar_0 \cdots r_k$. But by regularity of π , $Fr(s) \subseteq Fr(\pi)$, so we have

$$\begin{aligned} Bd(N_{\pi_0}^i a) &= EV(\pi) \cup Bd(a) \\ Fr(s \cdot a) &\subseteq (Fr(\pi) \setminus Bd(a)) \cup Fr(a) \\ Fr(N_{\pi_0}^k a) &= (Fr(\pi) \setminus Bd(a)) \cup Fr(a) \end{aligned}$$

and, as $Bd(a)$ is disjoint from $Fr(a)$ and $EV(\pi)$ is disjoint from $Fr(\pi) \cup Fr(a)$, we are done. \square

We now turn to the task of proving Lemma 6.7 which follows from the next two lemmas. The first characterises the syntactic form of normal \mathcal{H} -terms and will be useful in the subsequent analysis of derivations in Herbrand schemes.

Lemma 6.10. *If $r : \rho$ is a normal \mathcal{H} -term and ρ is a basic type whose co-domain is a pair $\sigma \times \tau$ in which σ is simple and τ is not simple, then either $r = \langle s, t \rangle$ for a Σ -term s and \mathcal{H} -term t , or $r = Fr_1 \cdots r_k$ for some non-terminal F and terms r_1, \dots, r_k .*

Proof. By induction on r . Let $r : \rho = \rho_1 \rightarrow \cdots \rightarrow \rho_l \rightarrow \sigma \times \tau$ be a normal \mathcal{H} -term satisfying the hypothesis of the lemma. Since ρ is not a simple type, r is not a Σ -symbol nor of the form $s \cdot a$ for a substitution stack a (by definition of normal terms). This leaves three cases: i) $l = 0$ and $r = \langle s, t \rangle$ for $s : \sigma$ and $t : \tau$; ii) $r = st$ for $s : \tau' \rightarrow \rho$ and $t : \tau'$; or iii) r is a non-terminal of \mathcal{H} . If (i), as σ is simple, s is a Σ -term by Lemma 4.8 and we are done. In case (ii), suppose $r = st$ is an application and $s : \sigma' = \tau' \rightarrow \rho$ and $t : \tau'$. If σ' is not basic then Lemma 4.9 implies $s = F$ for some π and i , whence $r = Ft$. On the other hand, if σ' is basic the induction hypothesis applies and $s = Fr_1 \cdots r_k$ for terms r_1, \dots, r_k , and so similarly for r . So we are done. \square

Lemma 6.11. *If $r : \iota \times \rho$ is a normal \mathcal{H} -term of pair type but not a pair then $r \rightarrow s$ for some \mathcal{H} -term s .*

Proof. Assume to the contrary that $r : \iota \times \rho$ is an \mathcal{H} -term which is not a pair and that there is no s such that $r \rightarrow s$. Without loss of generality assume r is minimal in length. By Lemma 6.10, $r = Fr_1, \dots, r_k$ for

some non-terminal F and terms r_1, \dots, r_k . It follows that $F \neq c_\sigma$ for any σ as otherwise $r \rightarrow \langle c, c_\rho \rangle$. Also $F \neq D_{\iota \times \rho}$ (as then $r \rightarrow r_1$) and $F \neq \hat{N}_\pi^i$ for any π and i . So $F = N_\pi^i$ for some π and i . The fact that r is not reducible means that the production rule for N_π^i requires pattern-matching on the final argument. But then $r_k : \iota \times \sigma$ for some σ , is not a pair and is not reducible, contradicting minimality of r . \square

We can now prove the two remaining lemmas.

Proof of Lemma 6.6. Let $r : \rho$ be a \mathcal{H} -term where $\rho = \sigma \times \tau$. Without loss of generality, we may assume $r \neq D_\rho r_0 r_1$ for any r_0 and r_1 . If r has the form $\langle s, t \rangle$ then trivially $r \sim \{\langle s, t \rangle\}$ and if $r = c_\rho$ then $r \sim \{\langle c_\sigma, c_\tau \rangle\}$. Otherwise, Lemma 6.10 implies that $r \sim N_\pi^i a r_0 \cdots r_k$ for some $\pi, i, a, r_0, \dots, r_k$. An induction on π determines terms s_0, \dots, s_l and t_0, \dots, t_l such that $r \sim \{\langle s_j, t_j \rangle \mid j \leq l\}$. Note that Lemma 6.11 implies there is no issue with pattern-matching stopping derivations from fully writing out. \square

Proof of Lemma 6.7. Let A be a Σ_1 formula and suppose $r : \tau_A$ is an irreducible normal term that is not a Σ -term. Without loss of generality, we assume every sub-term of r satisfies the statement of the lemma. Considering the types of non-terminals that form the \mathcal{H} -terms we deduce ρ is a pair type. By Lemma 6.6 we may assume r is a pair, say $r = \langle s, t \rangle$. As s is of simple type, it is a Σ -term. Hence t is an irreducible normal term which is not a Σ -term and, by the assumption on A , is of type τ_B for some $B \in \Sigma_1$. \square

6.3. Substitution and normality

Here we present some results concerning the interaction of subsumption with explicit substitutions which are needed for analysing the cut reduction and permutation rules for quantifiers.

Lemma 6.12. *If $t(r/x)$ and $t(s/x)$ are normal terms and $r \sqsubset s$ then $t(r/x) \sqsubset t(s/x)$.*

Proof. Direct consequence of the definition. \square

Lemma 6.13. *If $ru \sqsubset su$ for every term u then $r \sqsubset s$.*

Proof. For every derivation $t(r/x) \rightarrow^* r_0$ of a Σ -term, there are terms t', u_0, \dots, u_k such that $t(r/x) \rightarrow^* t'((ru_i)_{i \leq k}/\vec{x}) \rightarrow^* r_0$ and $t(s/x) \rightarrow^* t'((su_i)_{i \leq k}/\vec{x})$. Since normality is preserved through derivations, we are done. \square

Lemma 6.14. *Let r be a basic \mathcal{H} -term and a be a substitution stack over Σ . Then,*

1. $Fr(r(s/\alpha) \cdot a) \subseteq Fr(r \cdot ([\alpha \leftarrow s]a))$,
2. $Bd(r(s/\alpha) \cdot a) \subseteq Bd(r \cdot ([\alpha \leftarrow s]a))$ provided s is basic,
3. if $Fr(r) \cap Bd(a) = \emptyset$ then $r^a = r$,
4. if $\alpha \notin Bd(a) \cup Fr(a)$ then $r(s/\alpha)^a = r^a(s^a/\alpha)$ provided s is basic.

Proof. By induction on r and a . \square

Lemma 6.15. *Let $r : \rho$ and $a : \varsigma$ be Σ -terms, $s : \sigma$ a basic Σ -term, $\alpha^\rho \in \Sigma$ and π a regular proof. Under the assumption that $N_\pi^i([\alpha \leftarrow r]a)$ is normal the following hold.*

1. $s \cdot ([\alpha \leftarrow r]a) \sim s(r/\alpha) \cdot a$,
2. $r \sim r^\circ$,
3. $[\alpha \leftarrow r]a \sim [\alpha \leftarrow r \cdot a]a$,

4. If $EV(\pi) \cap Fr(r) = \emptyset$ then $\mathbf{N}_\pi^i([\alpha \leftarrow r]a) \sim \mathbf{N}_{\pi(r^\circ/\alpha)}^i a$,
5. If $\alpha \notin Fr(\pi)$ then $\mathbf{N}_\pi^i a \sim \mathbf{N}_\pi^i([\alpha \leftarrow r]a)$.

Proof. 1 is proved via induction on the basic term s . That r and a are Σ -terms is necessary for showing $s(r/\alpha) \cdot a \sqsubset s \cdot ([\alpha \leftarrow r]a)$. 2 follows from 1 by induction on r . Regarding 3, Lemma 6.14(1, 2) imply $Fr(a^\circ) \subseteq Fr(a)$ and $Bd(a^\circ) \subseteq Bd(a)$, so $\alpha \notin Fr(a^\circ)$ by normality. Hence, if t is a basic term then

$$\begin{aligned} t \cdot [\alpha \leftarrow r]a &\sim t(r^\circ/\alpha)^{a^\circ} \sim t^{a^\circ}((r^\circ)^{a^\circ}/\alpha) \\ &\sim t^{a^\circ}(((r^\circ)^{a^\circ})^{a^\circ}/\alpha) \\ &\sim t((r^\circ)^{a^\circ}/\alpha)^{a^\circ} \\ &\sim t \cdot [\alpha \leftarrow r \cdot a]a. \end{aligned}$$

The first and last equivalence are applications of 2; the second and fourth equivalence are consequences of Lemma 6.14(4); and the third equivalence uses Lemma 6.14(3) and the fact that $Fr(a^\circ) \cap Bd(a^\circ) = \emptyset$. Via 2 the above holds for t an arbitrary Σ -term, and from there generalises to deduce $[\alpha \leftarrow r]a \sim [\alpha \leftarrow r \cdot a]a$.

4 is derived by induction on π . By 2 we may assume r is a basic term, i.e., $r = r^\circ$. In the base case, where π is an axiom, the equivalence is trivial as $\mathbf{N}_\pi^i([\alpha \leftarrow r]a)r_1r_2 \sim \mathbf{N}_{\pi(r/\alpha)}^i ar_1r_2$ for any choice of r_1 and r_2 of appropriate type. The induction step is straightforward except in the case of quantifier rules. If π ends in the inference

$$\frac{\exists_{\vec{s}} \pi_0 \vdash \Gamma, A(\vec{s}/\vec{v})}{\pi \vdash \Gamma, \exists \vec{v}A}$$

where $\vec{s} = (s_j)_{j \leq k}$ then we have, if $i = |\Gamma|$, $b = [\alpha \leftarrow r]a$ and $r_1, \dots, r_{|\Gamma|}$ and t are suitable normal terms,

$$\begin{aligned} \mathbf{N}_\pi^{|\Gamma|} br_1 \cdots r_{|\Gamma|} t &\sim \langle s_0 \cdot b, \dots, s_k \cdot b, \mathbf{N}_{\pi_0}^{|\Gamma|} br_1 \cdots r_{|\Gamma|} \rangle \\ &\sim \langle s_0(r/\alpha) \cdot a, \dots, s_k(r/\alpha) \cdot a, \mathbf{N}_{\pi_0(r/\alpha)}^{|\Gamma|} ar_1 \cdots r_{|\Gamma|} \rangle \\ &\sim \mathbf{N}_{\pi(r/\alpha)}^{|\Gamma|} ar_1 \cdots r_{|\Gamma|} t \end{aligned}$$

where the second equivalence due to the induction hypothesis for $\mathbf{N}_{\pi_0}^{|\Gamma|}([\alpha \leftarrow r]a)$. The case $i < |\Gamma|$ is similar. For applications of the \forall inferences, we consider the inference

$$\frac{\forall_{\vec{\beta}} \pi_0 \vdash \Gamma, A(\vec{\beta}/\vec{v})}{\pi \vdash \Gamma, \forall \vec{v}A}$$

for $\vec{\beta} = (\beta_j^i)_{j \leq k}$. By normality of $\mathbf{N}_\pi^i([\alpha \leftarrow r]a)$ and the assumption that $Fr(r) \cap EV(\pi) = \emptyset$, it follows that $\beta_j \notin Fr(a) \cup Fr(r)$ for each j . Let $r_1, \dots, r_{|\Gamma|}$, s_0, \dots, s_k and t be such that $n := \mathbf{N}_\pi^i([\alpha \leftarrow r]a)r_1 \cdots r_{|\Gamma|}(s_0, \dots, s_k, t)$ is well-typed and normal. In particular, $\alpha, \beta_0, \dots, \beta_k \notin Fr(\langle s_0, \dots, s_k, t \rangle)$. Moreover, as s_j has simple type for each $j \leq k$, Lemma 4.8 implies that \vec{s} is a sequence of Σ -terms. Then assuming $\alpha \notin \{\beta_j \mid j \leq k\}$, and writing $[\vec{\beta} \leftarrow \vec{s}]b$ in place of $[\beta_0 \leftarrow s_0] \cdots [\beta_k \leftarrow s_k]b$, we have

$$\begin{aligned} n &\sim \mathbf{N}_{\pi_0}^i([\vec{\beta} \leftarrow \vec{s}][\alpha \leftarrow r]a)r_1 \cdots r_{|\Gamma|} t \\ &\sim \mathbf{N}_{\pi_0(\vec{s}/\vec{\beta})(r/\alpha)}^i ar_1 \cdots r_{|\Gamma|} t \\ &\sim \mathbf{N}_{\pi_0(r/\alpha)(\vec{s}/\vec{\beta})}^i ar_1 \cdots r_{|\Gamma|} t \end{aligned}$$

$$\begin{aligned} &\sim \mathbf{N}_{\pi_0}^i([\vec{\beta} \leftarrow \vec{s}]a)r_1 \cdots r_{|\Gamma|}t \\ &\sim \mathbf{N}_{\pi(r/\alpha)}^i ar_1 \cdots r_{|\Gamma|} \langle s_0, \dots, s_k, t \rangle. \end{aligned}$$

The third equivalence holds since $\alpha \notin Fr(\langle s_0, \dots, s_k, t \rangle)$ and $\beta_j \notin Fr(r)$ for any j . If $\alpha = \beta_j$ then $\pi^{(r/\alpha)} = \pi$ and, using again that $\alpha \notin Fr(\langle s_0, \dots, s_k \rangle)$, we have

$$\begin{aligned} n &\sim \mathbf{N}_{\pi_0}^i([\vec{\beta} \leftarrow \vec{s}][\alpha \leftarrow r]a)r_1 \cdots r_{|\Gamma|}t \\ &\sim \mathbf{N}_{\pi_0}^i([\vec{s}/\vec{\beta}](r/\alpha)) ar_1 \cdots r_{|\Gamma|}t \\ &\sim \mathbf{N}_{\pi_0}^i([\vec{s}/\vec{\beta}]) ar_1 \cdots r_{|\Gamma|}t \\ &\sim \mathbf{N}_{\pi_0}^i([\vec{\beta} \leftarrow \vec{s}]a)r_1 \cdots r_{|\Gamma|}t \\ &\sim \mathbf{N}_{\pi}^i ar_1 \cdots r_{|\Gamma|} \langle s_0, \dots, s_k, t \rangle \\ &\sim \mathbf{N}_{\pi(r/\alpha)}^i ar_1 \cdots r_{|\Gamma|} \langle s_0, \dots, s_k, t \rangle. \end{aligned}$$

Note, 5 is a special case of 4. \square

7. Language preservation for Gentzen-style cut elimination

Recall the relation $\pi \rightsquigarrow \pi'$ which expresses that π' is obtained from π by the application of a reduction rule in Figs. 2 and 3 to a sub-proof of π . In the present section we determine in which cases \rightsquigarrow supports: (i) *language inclusion*: $\pi \rightsquigarrow \pi'$ implies $L(\pi') \subseteq L(\pi)$; and (ii) *language equality*: $\pi \rightsquigarrow \pi'$ implies $L(\pi') = L(\pi)$. Establishing language inclusion for the cut reduction steps will suffice to derive the main theorem; language equality offers a finer study of the Herbrand content of proofs since if π_0 and π_1 can be connected by a sequence of forward and backward language preserving reductions then $L(\pi_0) = L(\pi_1)$.

Let π and π' be regular proofs of some sequent Γ . We say that π *subsumes* π' , in symbols $\pi' \sqsubset \pi$, if $\mathbf{N}_{\pi'}^i \sqsubset \mathbf{N}_{\pi}^i$ for every $i < |\Gamma|$. If π and π' each subsumes the other then π and π' are *equivalent*, in symbols $\pi \sim \pi'$.

Lemma 7.1. *Suppose π and π' are proofs of the same Σ_1 sequent. If $\pi' \sqsubset \pi$ then $L(\pi') \subseteq L(\pi)$.*

Proof. By definition. \square

Herbrand schemes have the property that their languages are invariant under many basic proof transformations. The first example we give concerns the operation of substitution in proofs:

Lemma 7.2. *Suppose π and π' are proofs with the same end-sequent such that π' is the result of replacing a sub-proof π_0 of π by π'_0 . If $\pi_0 \sqsubset \pi'_0$ then $\pi \sqsubset \pi'$.*

Proof. Let π, π_0, π' and π'_0 be as in the statement. We assume π and π' have the same end-sequent, say Γ . Given a subproof $\hat{\pi}$ of π which is not a proper subproof of π_0 , let $\hat{\pi}'$ denote the corresponding subproof of π' . Observe that if $\hat{\pi}$ is a subproof of π but not a proper subproof of π_0 then the non-terminals $\mathbf{N}_{\hat{\pi}}^j$ and $\mathbf{N}_{\hat{\pi}'}^j$ are of the same type for each j . Fix $i < |\Gamma|$ and a normal term $t_0 = t(\mathbf{N}_{\hat{\pi}}^i/x)$. Suppose $t_0 \rightarrow t_1 \rightarrow \cdots \rightarrow t_k = r$ is a derivation in \mathcal{H} of a Σ -term r . By Lemma 6.5, t_i is normal for every $i \leq k$ and, without loss of generality, we may assume t does not feature any non-terminals labelled by proofs with π_0 as a sub-proof. Throughout this derivation, recursively replace each occurrence of a non-terminal $\mathbf{N}_{\hat{\pi}}^j$ for which $\hat{\pi}$ is not a proper subproof of π_0 by the non-terminal $\mathbf{N}_{\hat{\pi}'}^j$. Arguing by induction on k , using $\pi_0 \sqsubset \pi'_0$, we deduce $t(\mathbf{N}_{\hat{\pi}}^i/x) \rightarrow^* s$ for some Σ -term s with $s^\circ = r^\circ$. \square

We begin our analysis of cut elimination by observing three common scenarios in which language computations can be simplified.

Lemma 7.3. *Let $\pi \vdash A_1, \dots, A_m, B, C_1, \dots, C_n$ and let $a : \varsigma$, $r_i : \hat{\tau}_{A_i}$, $s : \hat{\tau}_B$ and $t_j : \hat{\tau}_{C_j}$ be terms for $1 \leq i \leq m$ and $1 \leq j \leq n$ such that $N_\pi^m a \vec{r} s \vec{t}$ is normal. Let $\rho = \hat{\tau}_B$.*

1. *If B is prenex Σ_1 then $N_\pi^m a \vec{r} s \vec{t} \sim N_\pi^m a \vec{r} c_\rho \vec{t}$.*
2. *If $e(B) > 0$ and there are no applications of contraction to B in π then $N_\pi^m a \vec{r} s \vec{t} \sim N_\pi^m a \vec{r} c_\rho \vec{t}$.*
3. *If the final inference in π is an application of ρ with immediate sub-proof $\pi' \vdash A_1, \dots, A_m, C_1, B, C_2, \dots, C_n$ then for each $j \in [0, m) \cup [m + 2, m + n]$,*

$$N_\pi^j a r_1 \cdots r_m s t_1 \cdots t_n \sim N_{\pi'}^j a r_1 \cdots r_m t_1 s t_2 \cdots t_n.$$

Proof. 1. Since B is Σ_1 , $\hat{\tau}_B = \tau_B \rightarrow \epsilon$, whereby Lemma 6.2 completes the proof. 2 and 3 are proved by induction on π . \square

We now turn our attention to the analysis of the subsumption relation with respect to the cut reduction and permutation steps of Figs. 2 and 3. Only the most interesting cases will be covered in detail: the cut and quantifier permutation, and contraction and quantifier reduction. As before, we leave instances of the permutation inference implicit and make use of Lemma 7.3(3) without reference. Recall the characterisation of the cut inference from Remark 4.6:

$$\text{cut} \frac{\pi_0 \vdash \Gamma, A \quad \pi_1 \vdash \Delta, \bar{A}}{\pi \vdash \Gamma, \Delta} \quad N_\pi^i a \vec{x} \vec{y} \sim \begin{cases} N_{\pi_0}^i a \vec{x} (N_{\pi_1}^n a \vec{y} (N_{\pi_0}^m a \vec{x})), & \text{if } u(A) > 0, \\ N_{\pi_0}^i a \vec{x} (N_{\pi_1}^n a \vec{y}), & \text{if } e(A) > 0, \\ N_{\pi_0}^i a \vec{x} \langle \rangle, & \text{if } A \text{ is q.f.}, \end{cases}$$

where $i < |\Gamma|$ and \vec{x}, \vec{y} and a are terms of suitable type.

7.1. Cut permutation

Suppose $\pi \rightsquigarrow \pi'$ are the two proofs

$$\text{cut} \frac{\text{cut} \frac{\pi_0}{\Gamma, A, B} \quad \text{cut} \frac{\pi_1}{\Delta, \bar{A}}}{\hat{\pi} \vdash \Gamma, B, \Delta} \quad \pi_2}{\Lambda, \bar{B}} \quad \rightsquigarrow \quad \text{cut} \frac{\text{cut} \frac{\pi_0}{\Gamma, A, B} \quad \pi_2}{\hat{\pi}' \vdash \Gamma, A, \Lambda} \quad \pi_1}{\Delta, \bar{A}} \quad \pi'$$

Due to the asymmetry in the production rules for cut, it is necessary to split the analysis of this reduction into two cases, depending on whether or not A and B are both universally quantified. Provided at least one of the two formulæ is existentially quantified or quantifier free, the two proofs above are equivalent and their languages are equal. This is proved in Lemma 7.4. If both A and B are universally quantified we do not expect equivalence to hold in general. However, if there are no contractions to the formula \bar{A} in π_1 or the formula \bar{B} in π_2 , the proofs π and π' are equivalent. This is relevant to the cut reduction strategies employed in Theorem 1.1 and is treated in Lemma 7.7.

Lemma 7.4. *For $\pi \rightsquigarrow \pi'$ as above, if at least one of $u(A)$ and $u(B)$ is zero then $\pi \sim \pi'$.*

Proof. If one of A or B is quantifier-free the argument is straightforward following the production rules for cut. This leaves the following three cases to consider: $u(\bar{A}), u(\bar{B}) > 0$, $u(A), u(\bar{B}) > 0$ and $u(\bar{A}), u(\bar{B}) > 0$. We consider only the first case as the second is symmetric and the third follows a simpler argument. Thus assume $e(A), u(B) > 0$.

Let $\vec{r}, \vec{s}, \vec{t}$ be sequences of terms of length $m = |\Gamma|$, $n = |\Delta|$ and $o = |\Lambda|$ respectively, and let $a : \varsigma$ be an arbitrary substitution stack. By the production rules for cut we have, for each $i \leq m$, $j < n$ and $k < o$, and each term w, w' of suitable type,

$$\begin{aligned} N_{\hat{\pi}}^i a\vec{r}w\vec{s} &\sim \begin{cases} N_{\pi_0}^i a\vec{r}(N_{\pi_1}^n a\vec{s})w, & \text{if } i < m, \\ N_{\pi_0}^{m+1} a\vec{r}(N_{\pi_1}^n a\vec{s})w, & \text{if } i = m, \end{cases} \\ N_{\hat{\pi}}^{m+1+j} a\vec{r}w\vec{s} &\sim N_{\pi_1}^j a\vec{s}(N_{\pi_0}^m a\vec{r}(N_{\pi_1}^n a\vec{s})w) \\ N_{\hat{\pi}}^i a\vec{r}w'\vec{t} &\sim N_{\pi_0}^i a\vec{r}w'(N_{\pi_2}^o a\vec{t}(N_{\pi_0}^{m+1} a\vec{r}w')) \\ N_{\hat{\pi}'}^{m+1+k} a\vec{r}w'\vec{t} &\sim N_{\pi_2}^k a\vec{t}(N_{\pi_0}^{m+1} a\vec{r}w') \end{aligned}$$

In particular,

$$\hat{N}_{\hat{\pi}}^m a\vec{r}\vec{s} \sim N_{\pi_0}^{m+1} a\vec{r}(N_{\pi_1}^n a\vec{s}) \quad (15)$$

and so, for $i \leq m$,

$$N_{\hat{\pi}'}^i a\vec{r}(N_{\pi_1}^n a\vec{s})\vec{t} \sim N_{\pi_0}^i a\vec{r}(N_{\pi_1}^n a\vec{s})(N_{\pi_2}^o a\vec{t}(\hat{N}_{\hat{\pi}}^m a\vec{r}\vec{s})). \quad (16)$$

We prove $N_{\hat{\pi}}^i a\vec{r}\vec{s}\vec{t} \sim N_{\hat{\pi}'}^i a\vec{r}\vec{s}\vec{t}$ for every $i < m + n + o$, from which Lemma 6.13 implies $N_{\hat{\pi}}^i \sim N_{\hat{\pi}'}^i$. For $i < m$ we have

$$\begin{aligned} N_{\hat{\pi}}^i a\vec{r}\vec{s}\vec{t} &\sim N_{\hat{\pi}}^i a\vec{r}(N_{\pi_2}^o a\vec{t}(\hat{N}_{\hat{\pi}}^m a\vec{r}\vec{s}))\vec{s} \sim N_{\pi_0}^i a\vec{r}(N_{\pi_1}^n a\vec{s})(N_{\pi_2}^o a\vec{t}(\hat{N}_{\hat{\pi}}^m a\vec{r}\vec{s})) \\ &\sim N_{\hat{\pi}'}^i a\vec{r}(N_{\pi_1}^n a\vec{s})\vec{t} \\ &\sim N_{\hat{\pi}'}^i a\vec{r}\vec{s}\vec{t} \end{aligned}$$

by applying (16). For $j < n$,

$$\begin{aligned} N_{\hat{\pi}}^{m+1+j} a\vec{r}\vec{s}\vec{t} &\sim N_{\hat{\pi}}^{m+1+j} a\vec{r}(N_{\pi_2}^o a\vec{t}(\hat{N}_{\hat{\pi}}^m a\vec{r}\vec{s}))\vec{s} \\ &\sim N_{\pi_1}^j a\vec{s}(N_{\pi_0}^m a\vec{r}(N_{\pi_1}^n a\vec{s})(N_{\pi_2}^o a\vec{t}(\hat{N}_{\hat{\pi}}^m a\vec{r}\vec{s}))) \\ &\sim N_{\pi_1}^j a\vec{s}(N_{\hat{\pi}}^m a\vec{r}(N_{\pi_1}^n a\vec{s})\vec{t}) \\ &\sim N_{\hat{\pi}'}^{m+1+j} a\vec{r}\vec{s}\vec{t} \end{aligned}$$

again applying (16). For $k < o$, using (15):

$$\begin{aligned} N_{\hat{\pi}}^{m+n+k} a\vec{r}\vec{s}\vec{t} &\sim N_{\pi_2}^k a\vec{t}(\hat{N}_{\hat{\pi}}^m a\vec{r}\vec{s}) \sim N_{\pi_2}^k a\vec{t}(N_{\pi_0}^{m+1} a\vec{r}(N_{\pi_1}^n a\vec{s})) \\ &\sim N_{\hat{\pi}'}^{m+1+k} a\vec{r}(N_{\pi_1}^n a\vec{s})\vec{t} \\ &\sim N_{\hat{\pi}'}^{m+n+k} a\vec{r}\vec{s}\vec{t}. \quad \square \end{aligned}$$

As noted above, in the case $u(A)$ and $u(B)$ are both positive, language equality holds only in particular circumstances. A sufficient condition for this is given by the next lemma.

Lemma 7.5. *Let $\pi \rightsquigarrow \pi'$ be as above and assume $u(A), u(B) > 0$. Let $\rho = \hat{\tau}_{\bar{A}}$, $\sigma = \hat{\tau}_{\bar{B}}$, $R = \hat{N}_{\pi}^m a\vec{r}\vec{t}$ and $S = \hat{N}_{\pi}^m a\vec{r}\vec{s}$. If $R \sim \hat{N}_{\pi_0}^m a\vec{r}(\mathbf{N}_{\pi_2}^o a\vec{t}S)$ and $S \sim \hat{N}_{\pi_0}^{m+1} a\vec{r}(\mathbf{N}_{\pi_1}^n a\vec{s}R)$ then $\pi \sim \pi'$.*

Proof. Recall that R and S have type ρ and σ respectively. Then for $i < m$,

$$\begin{aligned} \mathbf{N}_{\pi}^i a\vec{r}\vec{s}\vec{t} &\sim \mathbf{N}_{\pi}^i a\vec{r}(\mathbf{N}_{\pi_2}^o a\vec{t}S)\vec{s} \\ &\sim \mathbf{N}_{\pi_0}^i a\vec{r}(\mathbf{N}_{\pi_1}^n a\vec{s}(\hat{N}_{\pi_0}^m a\vec{r}(\mathbf{N}_{\pi_2}^o a\vec{t}S)))(\mathbf{N}_{\pi_2}^o a\vec{t}S) \\ &\sim \mathbf{N}_{\pi_0}^i a\vec{r}(\mathbf{N}_{\pi_1}^n a\vec{s}R)(\mathbf{N}_{\pi_2}^o a\vec{t}(\hat{N}_{\pi_0}^{m+1} a\vec{t}(\mathbf{N}_{\pi_1}^n a\vec{s}R))) \\ &\sim \mathbf{N}_{\pi'}^i a\vec{r}(\mathbf{N}_{\pi_1}^n a\vec{s}R)\vec{t} \\ &\sim \mathbf{N}_{\pi'}^i a\vec{r}\vec{s}\vec{t} \end{aligned}$$

The other cases, namely $m \leq i < n + o$ follow similar reasoning. \square

Lemma 7.6. *If $A, B \in \Pi_1 \cup \Sigma_1$ then $\pi \sim \pi'$.*

Proof. Assume \bar{A} and \bar{B} are both Σ_1 formulæ (if not, apply Lemma 7.4). Let R and S be as in Lemma 7.5. We have, by Lemma 7.3,

$$\begin{aligned} R w &\sim \mathbf{N}_{\pi_0}^m a\vec{r}w(\mathbf{N}_{\pi_2}^n a\vec{t}(\hat{N}_{\pi_0}^{m+1} a\vec{r}w)) & S w' &\sim \mathbf{N}_{\pi_0}^{m+1} a\vec{r}(\mathbf{N}_{\pi_1}^n a\vec{s}(\hat{N}_{\pi_0}^m a\vec{r}w'))w' \\ &\sim \mathbf{N}_{\pi_0}^m a\vec{r}w(\mathbf{N}_{\pi_2}^n a\vec{t}S) & &\sim \mathbf{N}_{\pi_0}^{m+1} a\vec{r}(\mathbf{N}_{\pi_1}^n a\vec{s}R)w' \end{aligned}$$

and hence

$$R \sim \hat{N}_{\pi_0}^m a\vec{r}(\mathbf{N}_{\pi_2}^o a\vec{t}S) \quad S \sim \hat{N}_{\pi_0}^{m+1} a\vec{r}(\mathbf{N}_{\pi_1}^n a\vec{s}R).$$

The previous lemma then implies $\pi \sim \pi'$. \square

Lemma 7.7. *For the same π and π' , if there are no contractions to either the formula \bar{A} in the sub-proof π_1 or the formula \bar{B} in the sub-proof π_2 then $\pi \sim \pi'$.*

Proof. Suppose there are no contractions to \bar{B} in π_2 and let R and S be as above. By Lemma 7.3, $\mathbf{N}_{\pi_2}^k a\vec{t}u \sim \mathbf{N}_{\pi_2}^k a\vec{t}v$ for any two terms $u, v : \hat{\tau}_{\bar{B}}$. Hence, in particular,

$$R \sim \hat{N}_{\pi_0}^m a\vec{r}(\mathbf{N}_{\pi_2}^o a\vec{t}S) \quad \mathbf{N}_{\pi_2}^o a\vec{t}S \sim \hat{N}_{\pi_0}^{m+1} a\vec{r}(\mathbf{N}_{\pi_1}^n a\vec{s}R)$$

which suffice, by the proof of Lemma 7.5, to show $\pi \sim \pi'$. \square

7.2. Contraction reduction

Consider the two proofs

$$\begin{array}{c} \begin{array}{c} \text{---} \\ \diagup \quad \diagdown \\ \pi_0 \\ \text{---} \\ \Gamma, A, A \\ \text{---} \\ \text{c} \\ \hat{\pi} \vdash \Gamma, A \\ \text{---} \\ \text{cut} \\ \pi \vdash \Gamma, \Delta \end{array} \quad \begin{array}{c} \text{---} \\ \diagup \quad \diagdown \\ \pi_1 \\ \text{---} \\ \Delta, \bar{A} \end{array} \\ \rightsquigarrow \\ \begin{array}{c} \begin{array}{c} \text{---} \\ \diagup \quad \diagdown \\ \pi_0 \\ \text{---} \\ \Gamma, A, A \\ \text{---} \\ \text{cut} \\ \hat{\pi}' \vdash \Gamma, A, \Delta \\ \text{---} \\ \text{cut} \\ \Gamma, \Delta, \Delta \\ \text{---} \\ \text{c}^* \\ \pi' \vdash \Gamma, \Delta \end{array} \quad \begin{array}{c} \text{---} \\ \diagup \quad \diagdown \\ \pi_1 \\ \text{---} \\ \Delta, \bar{A} \\ \text{---} \\ \text{cut} \\ \Delta, \bar{A} \end{array} \end{array}$$

where π_1^* denotes a copy of π_1 with fresh eigenvariables. Observe that $\pi_1 \sim \pi_1^*$.

Although the reduction above does not in general induce language inclusion, for the two scenarios required in Theorem 1.1, namely either $u(A) = 0$ or there are no applications of contraction are applied to the formula \bar{A} in the sub-proof π_1 , we have $\pi' \sqsubset \pi$. The following two lemmas deal with these two cases.

Lemma 7.8. *If $u(A) = 0$ then $\pi' \sqsubset \pi$.*

Proof. If A is quantifier-free then $\pi \sim \pi'$ is easily established by following the reduction rules for cut. So assume $u(A) = 0 < u(\bar{A})$. Let $m = |\Gamma|$ and $n = |\Delta|$, and fix $i < m$ and $j < n$. Let $r = \mathbf{N}_{\pi_1}^n a\vec{y}$ and $r_* = \mathbf{N}_{\pi_1^*}^n a\vec{y}$. Unravelling the production rules for the two proofs yield

$$\begin{aligned} \mathbf{N}_{\pi}^i a\vec{x}\vec{y} &\sim \mathbf{N}_{\pi_0}^i a\vec{x}rr & \mathbf{N}_{\pi}^{m+j} a\vec{x}\vec{y} &\sqsubset \{\mathbf{N}_{\pi_1}^j a\vec{y}(\mathbf{N}_{\pi_0}^{m+1} a\vec{x}rr), \mathbf{N}_{\pi_1}^j a\vec{y}(\mathbf{N}_{\pi_0}^m a\vec{x}rr)\} \\ \mathbf{N}_{\pi'}^i a\vec{x}\vec{y} &\sim \mathbf{N}_{\pi_0}^i a\vec{x}r_*r & \mathbf{N}_{\pi'}^{m+j} a\vec{x}\vec{y} &\sim \{\mathbf{N}_{\pi_1}^j a\vec{y}(\mathbf{N}_{\pi_0}^{m+1} a\vec{x}r_*r), \mathbf{N}_{\pi_1^*}^j a\vec{y}(\mathbf{N}_{\pi_0}^m a\vec{x}r_*r)\} \end{aligned}$$

Since $\pi_1 \sim \pi_1^*$, Lemma 6.12 implies $\pi' \sqsubset \pi$. \square

Lemma 7.9. *If $u(A) > 0$ and there are no contractions on the formula \bar{A} in π_1 , then $\pi' \sim \pi$.*

Proof. Suppose $u(A) > 0$. Let $\tau = \hat{\tau}_{\bar{A}}$. Lemma 7.3 implies $\mathbf{N}_{\pi_1}^j a\vec{y}s \sim \mathbf{N}_{\pi_1}^j a\vec{y}c_{\tau}$ for every $s : \tau$. Concerning derivations from π , this yields the following equivalences for $i < |\Gamma|$ and $j < |\Delta|$.

$$\begin{aligned} \mathbf{N}_{\pi}^i a\vec{x}\vec{y} &\sim \mathbf{N}_{\pi}^i a\vec{x}(\mathbf{N}_{\pi_1}^n a\vec{y}(\hat{\mathbf{N}}_{\pi}^m a\vec{x})) & \mathbf{N}_{\pi}^{m+j} a\vec{x}\vec{y} &\sim \mathbf{N}_{\pi_1}^j a\vec{y}(\hat{\mathbf{N}}_{\pi}^m a\vec{x}) \\ &\sim \mathbf{N}_{\pi}^i a\vec{x}(\mathbf{N}_{\pi_1}^n a\vec{y}c_{\tau}) & &\sim \mathbf{N}_{\pi_1}^j a\vec{y}c_{\tau} \\ &\sim \mathbf{N}_{\pi_0}^i a\vec{x}(\mathbf{N}_{\pi_1}^n a\vec{y}c_{\tau})(\mathbf{N}_{\pi_1}^n a\vec{y}c_{\tau}) \end{aligned}$$

Starting from π' we obtain

$$\begin{aligned} \mathbf{N}_{\pi'}^i a\vec{x}\vec{y} &\sim \mathbf{N}_{\pi'}^i a\vec{x}(\mathbf{N}_{\pi_1}^n a\vec{y}(\hat{\mathbf{N}}_{\pi'}^m a\vec{x}))\vec{y} & \mathbf{N}_{\pi'}^{m+j} a\vec{x}\vec{y} &\sim \{\mathbf{N}_{\pi'}^{m+1+j} a\vec{x}(\mathbf{N}_{\pi_1^*}^n a\vec{y}c_{\tau})\vec{y}, \mathbf{N}_{\pi_1}^j a\vec{y}c_{\tau}\} \\ &\sim \mathbf{N}_{\pi'}^i a\vec{x}(\mathbf{N}_{\pi_1}^n a\vec{y}c_{\tau})\vec{y} & &\sim \{\mathbf{N}_{\pi_1}^j a\vec{y}c_{\tau}, \mathbf{N}_{\pi_1}^j a\vec{y}c_{\tau}\} \\ &\sim \mathbf{N}_{\pi_0}^i a\vec{x}(\mathbf{N}_{\pi_1}^n a\vec{y}c_{\tau})(\mathbf{N}_{\pi_1}^n a\vec{y}c_{\tau}) \end{aligned}$$

So $\pi' \sim \pi$. \square

7.3. Quantifier permutation

Concerning permuting quantifier rules with cut, consider the following two proofs.

$$\begin{array}{ccc} \begin{array}{c} \triangle \pi_0 \\ \Gamma, A(\vec{\alpha}/\vec{v}), B \\ \forall_{\vec{\alpha}} \frac{}{\hat{\pi} \vdash \Gamma, \forall \vec{v} A, B} \\ \triangle \pi_1 \\ \Delta, \bar{B} \\ \text{cut} \frac{}{\pi \vdash \Gamma, \forall \vec{v} A, \Delta} \end{array} & \rightsquigarrow & \begin{array}{c} \triangle \pi_0 \quad \triangle \pi_1 \\ \Gamma, A(\vec{\alpha}/\vec{v}), B \quad \Delta, \bar{B} \\ \text{cut} \frac{}{\hat{\pi}' \vdash \Gamma, A(\vec{\alpha}/\vec{v}), \Delta} \\ \forall \frac{}{\pi' \vdash \Gamma, \forall \vec{v} A, \Delta} \end{array} \end{array} \quad (17)$$

Let $\vec{\alpha} = (\alpha_j)_{j \leq p}$ and $\vec{v} = (v_j)_{j \leq p}$. Regularity ensures that $u(A) = 0$. In the following, if $\vec{u} = (u_j)_{j \leq p}$ is a sequence of terms of type ι and $u_{p+1} : \hat{\tau}_A$, we write $\vec{u} * u_{p+1}$ to abbreviate the sequence term $\langle u_0, \dots, u_{p+1} \rangle : \hat{\tau}_{\forall \vec{v} A}$.

Like with the case of permuting cuts, an application of the quantifier permutation reduction does not preserve equivalence of proofs in all cases. For the main theorem it suffice to prove only $\pi' \sqsubset \pi$. This is taken up in Lemma 7.11 below. First, however, we show that if B is not universally quantified then indeed $\pi \sim \pi'$.

Lemma 7.10. *For π and π' above, if $u(B) = 0$ then $\pi' \sim \pi$.*

Proof. Suppose $u(B) = 0$ and B is not quantifier-free. The other cases involve much similar arguments. Fix \vec{r} and \vec{s} sequences of normal terms of length $m = |\Gamma|$ and $n = |\Delta|$ respectively, and normal terms t and $\vec{u} \star u'$ of type $\hat{\tau}_{\forall \vec{v}A}$. By regularity of π and Lemma 6.15,

$$\mathbf{N}_{\pi_1}^j([\vec{\alpha} \leftarrow \vec{u}]a)\vec{s} \sim \mathbf{N}_{\pi_1}^j a\vec{s}$$

for each $j \leq n$. Concerning π the following equivalences therefore appear for $i \leq m$, $j < n$ and $k \leq m+1$,

$$\begin{aligned} \mathbf{N}_{\pi}^i a\vec{r}t\vec{s} &\sim \mathbf{N}_{\pi}^i a\vec{r}t(\mathbf{N}_{\pi_1}^n a\vec{s}) & \mathbf{N}_{\pi}^k a\vec{r}(\vec{u} \star u') &\sim \mathbf{N}_{\pi_0}^k([\vec{\alpha} \leftarrow \vec{u}]a)\vec{r}u' \\ \mathbf{N}_{\pi}^{m+1+j} a\vec{r}t\vec{s} &\sim \mathbf{N}_{\pi_1}^j a\vec{s}(\mathbf{N}_{\pi}^{m+1} a\vec{r}t(\mathbf{N}_{\pi_1}^n a\vec{s})) \end{aligned}$$

So, if t is a normal term and $t \sim \{\vec{u}_0 \star u'_0, \dots, \vec{u}_l \star u'_l\}$ is given by Lemma 6.6 then for $i \leq m$ and $j < n$,

$$\begin{aligned} \mathbf{N}_{\pi}^i a\vec{r}t\vec{s} &\sim \{\mathbf{N}_{\pi_0}^i([\vec{\alpha} \leftarrow \vec{u}_k]a)\vec{r}u'_k(\mathbf{N}_{\pi_1}^n a\vec{s}) \mid k \leq l\} \\ \mathbf{N}_{\pi}^{m+1+j} a\vec{r}t\vec{s} &\sim \{\mathbf{N}_{\pi_1}^j a\vec{s}(\mathbf{N}_{\pi_0}^{m+1}([\vec{\alpha} \leftarrow \vec{u}_k]a)\vec{r}u'_k(\mathbf{N}_{\pi_1}^n a\vec{s})) \mid k \leq l\} \end{aligned}$$

Examining π' , we observe

$$\begin{aligned} \mathbf{N}_{\pi'}^i a\vec{r}t\vec{s} &\sim \{\mathbf{N}_{\pi'}^i([\vec{\alpha} \leftarrow \vec{u}_k]a)\vec{r}u'_k\vec{s} \mid k \leq l\} \\ &\sim \{\mathbf{N}_{\pi_0}^i([\vec{\alpha} \leftarrow \vec{u}_k]a)\vec{r}u'_k(\mathbf{N}_{\pi_1}^n a\vec{s}) \mid k \leq l\} \\ \mathbf{N}_{\pi'}^{n+j} a\vec{r}t\vec{s} &\sim \{\mathbf{N}_{\pi_1}^j a\vec{s}(\mathbf{N}_{\pi_0}^{m+1}([\vec{\alpha} \leftarrow \vec{u}_k]a)\vec{r}u'_k(\mathbf{N}_{\pi_1}^n a\vec{s})) \mid k \leq l\} \end{aligned}$$

Hence $\mathbf{N}_{\pi'}^i \sim \mathbf{N}_{\pi}^i$ for every $i \leq m+n$ and so $\pi' \sim \pi$. \square

Lemma 7.11. *For π and π' as in (17), $\pi' \sqsubset \pi$.*

Proof. Fix \vec{r} and \vec{s} sequences of terms of length $m = |\Gamma|$ and $n = |\Delta|$ respectively. Let $\vec{u} \star u' : \hat{\tau}_{\forall \vec{v}A}$. Suppose $u(B) > 0$ and $i \leq m$. The other cases have been considered earlier or involve similar but simpler arguments. As was observed earlier,

$$\mathbf{N}_{\pi_1}^j([\vec{\alpha} \leftarrow \vec{u}]a)\vec{s} \sim \mathbf{N}_{\pi_1}^j a\vec{s}$$

for each $j \leq n$. With respect to π' the following equivalences therefore appear.

$$\begin{aligned} \mathbf{N}_{\pi'}^i a\vec{r}(\vec{u} \star u')\vec{s} &\sim \mathbf{N}_{\pi'}^i([\vec{\alpha} \leftarrow \vec{u}]a)\vec{r}u'\vec{s} \\ &\sim \mathbf{N}_{\pi_0}^i([\vec{\alpha} \leftarrow \vec{u}]a)\vec{r}u'(\mathbf{N}_{\pi_1}^n a\vec{s}(\mathbf{N}_{\pi_0}^{m+1}([\vec{\alpha} \leftarrow \vec{u}]a)\vec{r}u')) \end{aligned}$$

whereas the rules for π yield, for arbitrary $t : \hat{\tau}_{\forall \vec{v}A}$,

$$\mathbf{N}_{\pi}^i a\vec{r}t\vec{s} \sim \mathbf{N}_{\pi}^i a\vec{r}t(\mathbf{N}_{\pi_1}^n a\vec{s}(\mathbf{N}_{\pi}^{m+1} a\vec{r}t)) \quad \mathbf{N}_{\pi}^k a\vec{r}(\vec{u} \star u') \sim \mathbf{N}_{\pi_0}^k([\vec{\alpha} \leftarrow \vec{u}]a)\vec{r}u'$$

If t is a normal term and $t \sim \{\vec{u}_0 \star u'_0, \dots, \vec{u}_l \star u'_l\}$ is given by Lemma 6.6 then for each $i \leq m$,

$$\begin{aligned} \mathbf{N}_{\pi}^i a \vec{r} t &\sim \{\mathbf{N}_{\pi_0}^i([\vec{\alpha} \leftarrow \vec{u}_k] a) \vec{r} u'_k \mid k \leq l\} \\ \mathbf{N}_{\pi}^i a \vec{r} t \vec{s} &\sim \{\mathbf{N}_{\pi_0}^i([\vec{\alpha} \leftarrow \vec{u}_k] a) \vec{r} u'_k (\mathbf{N}_{\pi_1}^n a \vec{s} (\mathbf{N}_{\pi_0}^{m+1}([\vec{\alpha} \leftarrow \vec{u}_j] a) \vec{r} u'_j)) \mid k, j \leq l\} \end{aligned} \quad (18)$$

whereas, due to pattern-matching in the production rule for π' ,

$$\mathbf{N}_{\pi'}^i a \vec{r} t \vec{s} \sim \{\mathbf{N}_{\pi_0}^i([\vec{\alpha} \leftarrow \vec{u}_k] a) \vec{r} u'_k (\mathbf{N}_{\pi_1}^n a \vec{s} (\mathbf{N}_{\pi_0}^{m+1}([\vec{\alpha} \leftarrow \vec{u}_k] a) \vec{r} u'_k)) \mid k \leq l\} \quad (19)$$

Hence $\mathbf{N}_{\pi'}^i \sqsubset \mathbf{N}_{\pi}^i$ and $\pi' \sqsubset \pi$. \square

The contrast between equations (18) and (19) demonstrates why $\pi \sqsubset \pi'$ need not hold in general.

7.4. Quantifier reduction

Consider the reduction

$$\begin{array}{c} \begin{array}{ccc} \begin{array}{c} \pi_0 \\ \Gamma, A(\vec{\alpha}/\vec{v}) \\ \forall_{\vec{\alpha}} \frac{\pi_0 \vdash \Gamma, \forall \vec{v} A}{\pi_0 \vdash \Gamma, \forall \vec{v} A} \\ \text{cut} \end{array} & \begin{array}{c} \pi_1 \\ \Delta, \vec{A}(\vec{s}/\vec{v}) \\ \exists_{\vec{s}} \frac{\pi_1 \vdash \Delta, \exists \vec{v} \vec{A}}{\pi_1 \vdash \Delta, \exists \vec{v} \vec{A}} \end{array} & \rightsquigarrow \\ \hline \pi \vdash \Gamma, \Delta \end{array} & \rightsquigarrow & \begin{array}{ccc} \begin{array}{c} \pi_0^{(\vec{s}/\vec{\alpha})} \\ \Gamma, A(\vec{s}/\vec{v}) \\ \text{cut} \end{array} & \begin{array}{c} \pi_1 \\ \Delta, \vec{A}(\vec{s}/\vec{v}) \end{array} & \\ \hline \pi' \vdash \Gamma, \Delta \end{array} \end{array}$$

Lemma 7.12. *If $\pi \rightsquigarrow \pi'$ is the reduction above then $\pi \sim \pi'$.*

Proof. Let $m = |\Gamma|$, $n = |\Delta|$, $\vec{\alpha} = (\alpha_i)_{i \leq p}$ and $\vec{s} = (s_i)_{i \leq p}$. Recall that $\vec{s} \cdot a = (s_i \cdot a)_{i \leq p}$. Note that regularity of π implies $u(A) = 0$. This leaves two cases to consider: A is quantifier-free or $e(A) > 0$. Suppose the latter, so the cut in π' remains a quantified cut (the case A is q.f. follows an analogous argument). The following equivalences arise, where $i < m$ and $j < n$.

$$\begin{aligned} \mathbf{N}_{\pi}^i a \vec{r} t &\sim \mathbf{N}_{\pi_0}^i a \vec{r} (\mathbf{N}_{\pi_1}^n a \vec{t} (\mathbf{N}_{\pi_0}^m a \vec{r})) & \mathbf{N}_{\pi}^{m+j} a \vec{r} t &\sim \mathbf{N}_{\pi_1}^j a \vec{t} (\mathbf{N}_{\pi_0}^m a \vec{r}) \\ &\sim \mathbf{N}_{\pi_0}^i a \vec{r} (\vec{s} \cdot a \star \mathbf{N}_{\pi_1}^n a \vec{t}) & &\sim \mathbf{N}_{\pi_1}^j a \vec{t} (\mathbf{N}_{\pi_0}^m a \vec{r} (\vec{s} \cdot a \star \mathbf{N}_{\pi_1}^n a \vec{t})) \\ &\sim \mathbf{N}_{\pi_0}^i ([\vec{\alpha} \leftarrow \vec{s} \cdot a] a) \vec{r} (\mathbf{N}_{\pi_1}^n a \vec{t}) & &\sim \mathbf{N}_{\pi_1}^j a \vec{t} (\mathbf{N}_{\pi_0}^m ([\vec{\alpha} \leftarrow \vec{s} \cdot a] a) \vec{r} (\mathbf{N}_{\pi_1}^n a \vec{t})) \\ &\sim \mathbf{N}_{\pi_0^{(\vec{s}/\vec{\alpha})}}^i a \vec{r} (\mathbf{N}_{\pi_1}^n a \vec{t}) & &\sim \mathbf{N}_{\pi_0^{(\vec{s}/\vec{\alpha})}}^j a \vec{t} (\mathbf{N}_{\pi_1}^n a \vec{r} (\mathbf{N}_{\pi_0}^m a \vec{t})) \\ &\sim \mathbf{N}_{\pi'}^i a \vec{r} t & &\sim \mathbf{N}_{\pi'}^{m+j} a \vec{r} t. \end{aligned}$$

The penultimate equivalence in each column is given by Lemma 6.15. \square

7.5. Remaining reductions

The remaining rules are all straightforward to analyse and all induce language equality with the exception of weakening reduction for which only language inclusion holds in general.

7.6. Proof of main theorem

We can now prove Theorem 1.1. Let $\pi \vdash \exists \vec{v} F_{qf}$ be a regular proof and $\pi = \pi_0 \rightsquigarrow \pi_1 \rightsquigarrow \dots \rightsquigarrow \pi_n$ be a reduction of π to a quasi cut-free proof π_n such that for each $i < n$, the reduction $\pi_i \rightsquigarrow \pi_{i+1}$ applies a cut

reduction or permutation rule from Figs. 2 or 3 to a sub-proof of π_i with the restriction that a rule reducing the strong quantifier side of a cut is applied only if no other reduction of this cut is possible. By Lemma 7.2 and the analysis in the previous section, $L(\pi_{i+1}) \subseteq L(\pi_i)$ for each $i < n$. This together with Lemma 4.12 establishes part (iii) of the theorem. The existence of a reduction of the form above is well-known: see, e.g. [40], hence (i). Acyclicity of \mathcal{H}_π is shown in Lemma 4.7, the bound on the order of \mathcal{H}_π is given by Lemma 4.14, and the language bound in (ii) follows from Theorem 4.15.

8. Discussion

This work contributes to the structural analysis of first-order proofs with respect to their Herbrand content. To a first-order classical proof $\pi \vdash F$ of a Σ_1 formula we associate a recursion scheme \mathcal{H} with a finite language that constitutes a Herbrand set for F . More generally, the language of \mathcal{H} covers the Herbrand set implicit in any quasi cut-free proof obtained from π by a sequence of reductions fulfilling the following two restrictions.

1. A contraction on a universally quantified formula is reduced only when no other reduction rule is applicable to this cut;
2. If two cuts are permuted in the following form then either there are no contractions on the formula \bar{B} in the relevant subproof, or one of A and B is not universally quantified.

$$\text{cut} \frac{\text{cut} \frac{\Gamma, A, B \quad \Delta, \bar{A}}{\Gamma, \Delta, B} \quad \Lambda, \bar{B}}{\Gamma, \Delta, \Lambda} \rightsquigarrow \text{cut} \frac{\text{cut} \frac{\Gamma, A, B \quad \Lambda, \bar{B}}{\Gamma, \Lambda, A} \quad \Delta, \bar{A}}{\Gamma, \Delta, \Lambda}$$

The size of the Herbrand set is bounded by $2^{4|\pi|_n^3}$ where $|\pi|$ is the number of inferences in π and n is the maximal quantifier rank of a cut in π . Comparing with related work, the bound on the cardinality of the Herbrand expansion obtained by Gerhardy and Kohlenbach [17, Corollary 15] is $2^{3\|t\|_{\text{dg}(\phi)+1}}$ where ϕ is a proof in Shoenfield’s calculus [38], t is the realiser extracted from ϕ , and $\|t\|$ is the number of symbols in t . The degree $\text{dg}(\phi)$ is the maximal \neg -depth of a cut formula in ϕ . The \neg -depth of a formula is defined precisely in the discussion on pp. 17–25 of [16] as the maximal number of nested negations over quantifier-free sub-formulae (that may contain an arbitrary number of negations). This is sufficient for describing the height of the tower of exponentials since, in Shoenfield’s system, $\exists x$ is considered an abbreviation of $\neg\forall x\neg$. Thus (the translation of) a $\Pi_n \cup \Sigma_n$ formula has \neg -depth at most n . Presumably it is possible to give a polynomial translation from the sequent calculus into Shoenfield’s system which preserves the maximal \neg -depth of cut formulae (but, to the knowledge of the authors, this has not been done in the literature) and, moreover, to bound $\|t\|$ polynomially in terms of the number of inferences of ϕ . Under these assumptions, the bound of Gerhardy and Kohlenbach would yield the upper bound $2^{p(|\pi|_n)}$ on the cardinality of a Herbrand expansion for some polynomial p and any sequent calculus proof π with $\Pi_n \cup \Sigma_n$ -cuts. This would be one exponent less than our own.

Closely related is the bound obtained by Buss in [12]. The proof of Theorem 9 of [12] shows, given a proof π where all cut formulae are contained in $\Pi_n \cup \Sigma_n$, that there is a cut-free proof whose number of inferences is no greater than $2^{|\pi|_{n+2}}$. As an immediate corollary this also yields the upper bound of $2^{|\pi|_{n+2}}$ on the cardinality of the Herbrand expansion. If one is interested in the cardinality of the Herbrand expansion, Buss’s bound and our Theorem 1.1 give the same number of iterations of the exponential function, but Gerhardy and Kohlenbach’s would give one less. If one is interested in the number of inferences in the cut-free proof, Buss’s bound is one exponential better than ours but has the same number of exponentials as the one that could be obtained from Gerhardy and Kohlenbach’s since the number of inferences is at most exponential in the

cardinality of the Herbrand expansion (considering the symbolic complexity of the end-sequent is constant). That being said, the bounds we obtain apply to any cut-free proof (and Herbrand expansion) that can be reached by the class of reductions pertaining to 1 and 2 above. In particular, it places no restriction on which cut is to be reduced at any given step, and therefore accommodates a variety of strategies, including top-most and maximal cut-complexity. Whether this freedom of strategies necessitates the larger bound is not entirely clear, and requires further investigation.

Below we highlight finer features of our representation of Herbrand's theorem and some potential applications.

8.1. Sequent versus trace grammars

In this paper, the grammar associated to a proof is 'sequent based' in the following sense. Consider an inference of the form

$$\text{r} \frac{G_0, \dots, G_n}{F_0, \dots, F_m}$$

The production rules corresponding to r can be seen as transforming a sequence of inputs (x_0, \dots, x_m) for the formulæ F_0, \dots, F_m to a sequence of terms (t_0, \dots, t_n) which are used as inputs for G_0, \dots, G_n in the inference rule immediately above r . The production rules affect the whole sequence of inputs regardless of which formula is active. This is in contrast with the 'trace'-based grammars of, e.g., [21,1,2] where an inference of the form

$$\text{r} \frac{\Gamma, G}{\Gamma, F}$$

is associated a production rule that updates an input for F to an input for G , entirely ignoring presence of formulæ in Γ . In the latter type of grammars the derivations can be viewed as traces that climb up and also down the proof tree mimicking the traces revealed through Gentzen-style cut-elimination. These grammars are generally cyclic and it is necessary to place equality constraints (the 'rigidity' conditions of [21,1]) on derivations to ensure finite languages. For proofs that contain cuts with complexity greater than Π_2/Σ_2 the trace-based analysis quickly becomes infeasible. In contrast, the sequent-based approach generates an acyclic term grammar that not only ensures a finite language but allows one to obtain upper bounds on language size by standard language-theoretic arguments.

8.2. Providing a minimal grammar

Part of the motivation behind this study is to ultimately invert the cut-elimination procedure and find an algorithmic method for introducing cuts into cut-free proofs. The idea has been successfully carried out for Π_1/Σ_1 -cut introduction and more recently for the introduction of a single Π_2/Σ_2 -cut (see [24,23,22,27]). The general method proceeds as follows. Given a cut-free proof π , one computes a concise representation of π as a term grammar (such as a regular tree grammar whose language contains the Herbrand set induced by π). The grammar is then viewed as a proof with cut, in which the cut-formulæ are yet to be determined. Finding the cut-formulæ involves solving a unification problem induced by the grammar. Key to successfully carrying out this procedure is identifying natural classes of formal grammars that describe the instantiation structure of a proof with cut. Higher order recursion schemes are a promising candidate to lift the method of cut-introduction above the Π_2 level.

8.3. First-order logic in finite types

A natural extension to consider is first-order logic in finite types, namely many-sorted predicate logic with a sort of individuals for each simple type and well-typed application as a term forming operation. On the sequent calculus side, we add new quantifier inferences for each type:

$$\forall_{\alpha}^{\sigma} \frac{\Gamma, A(\alpha^{\sigma}/v^{\sigma})}{\Gamma, \forall v^{\sigma} A} \qquad \exists_r^{\sigma} \frac{\Gamma, A(r^{\sigma}/v^{\sigma})}{\Gamma, \exists v^{\sigma} A}$$

At the level of types, the definition of τ_A and $\hat{\tau}_A$ is extended to incorporate higher-type quantification: for example, $\tau_{\forall v^{\sigma} F} = \tau_F$ and $\tau_{\exists v^{\sigma} F} = \sigma \times \hat{\tau}_{\bar{F}}$. The production rules corresponding to the new quantifier inferences will be as before, though the move to higher-type means substitution stacks may contain symbols of non-ground types. With appropriate modifications to the notion of normal terms, we expect the analogous language preservation lemmas to hold.

8.4. Lifting the prenex restriction

Our representation of first-order proofs as recursion schemes forces an asymmetric interpretation of formulæ to types that does not easily generalise to non-prenex cuts. Specifically, the type of an existentially quantified formula is, except in the case of Σ_1 , an order higher than the dual universally quantified formula. This disparity is due to the production rules for cut which treat the cut formula from one premise as a function which receives as its input ‘witnesses’ for the dual (cut) formula in the other premise. If the same representation is to be applied to non-prenex cuts we should expect the types associated to, say, a conjunction to be an order higher than those assigned to the dual disjunction. Motivated by the existing connection to the functional interpretation (Theorem 4.4), the duality of the logical connectives in accounts of Shoenfield’s functional interpretation (such as is examined in [39]) may well be relevant. Nevertheless, the primary desideratum is that the associated production rules respect the local operations of reductive cut elimination, so the types of any additional connectives should respect the ‘proof-theoretic’ semantics imbued by cut-elimination over other computational interpretations.

8.5. Functional interpretation for sequent calculus

Shoenfield’s functional interpretation [38] maps every first-order formula A to a Π_2 formula in finite types $A^S = \forall \vec{x} \exists \vec{y} A_S(\vec{x}, \vec{y})$ where A_S is quantifier-free. Gerhardy and Kohlenbach [17] show how a sequence of terms \vec{t} can be extracted from a proof of A such that $A_S(\vec{x}, \vec{t})$ is derivable in a quantifier-free predicate logic for simple types. As we saw in Theorem 4.4, modulo a calculus for basic types we may assume A^S has the specific form $\forall x^{\hat{\tau}_A} \exists y^{\tau_A} A_S(x, y)$, whence the approach of [17] extracts, from a proof $\pi \vdash A$, a realiser $t_{\pi} : \hat{\tau}_A \rightarrow \tau_A$ and a proof of $A_S(x, t_{\pi} x)$. The Herbrand schemes of this article operate similarly. Given a proof $\pi \vdash A$ with singleton end-sequent, the term $N_{\pi}^0 \perp$ has the type of realisers for A , namely $\hat{\tau}_A \rightarrow \tau_A$. Moreover, it is easily shown that the induced quantifier-free formula $A_S(c_{\hat{\tau}_A}, N_{\pi}^0 \perp c_{\hat{\tau}_A})$ can be derived in a quantifier-free first-order logic extended by an equational calculus for the production rules of the Herbrand scheme \mathcal{H}_{π} (recall $c_{\hat{\tau}_A}$ is a canonical constant inhabiting the type $\hat{\tau}_A$).

At this stage it is unclear how close our approach is to the functional interpretation. For instance, an obvious question is whether there exists a class of classical sequent calculus proofs for which Herbrand schemes behave, in a suitable sense, identically to the functional interpretation. Although we do not know the answer to this question, even the superficial connection between the two formalisms that has come to light points to an unexpected correlation between witness extraction and classical cut elimination.

Declaration of competing interest

Declarations of interest: none.

Acknowledgements

This research was supported through the programme “Research in Pairs” by the Mathematisches Forschungsinstitut Oberwolfach in 2016, the Vienna Science Technology Fund (WWTF) [VRG12-004], the Austrian Science Fund (FWF) [P25160], the Swedish Research Council (VR) [2016-03502 & 2017-05111], and by the Knut and Alice Wallenberg Foundation (KAW) [2015.0179].

The authors also wish to thank Luke Ong for valuable discussions on higher order recursion schemes and pattern matching which have greatly improved the presentation of the results. We also extend our gratitude to the anonymous referee for their insightful comments and suggestions, and for drawing our attention to potential deeper connections between our approach via higher order recursion and the functional interpretation.

References

- [1] Bahareh Afshari, Stefan Hetzl, Graham E. Leigh, Herbrand disjunctions, cut elimination and context-free tree grammars, in: Thorsten Altenkirch (Ed.), 13th International Conference on Typed Lambda Calculi and Applications (TLCA 2015), in: Leibniz International Proceedings in Informatics (LIPIcs), vol. 38, Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 2015, pp. 1–16.
- [2] Bahareh Afshari, Stefan Hetzl, Graham E. Leigh, Herbrand confluence for first-order proofs with Π_2 -cuts, in: Proof: Concepts of Proof in Mathematics, Philosophy, and Computer Science, in: Ontos Mathematical Logic, vol. 6, De Gruyter, 2016, pp. 4–50.
- [3] Bahareh Afshari, Stefan Hetzl, Graham E. Leigh, Herbrand’s theorem revisited, *Proc. Appl. Math. Mech.* 16 (1) (2016) 905–906.
- [4] Bahareh Afshari, Stefan Hetzl, Graham E. Leigh, On the Herbrand content of LK, in: Ulrich Kohlenbach, Steffen van Bakel, Stefano Berardi (Eds.), Proceedings of Sixth International Workshop on Classical Logic and Computation (CL&C 2016), in: EPTCS, vol. 213, 2016, pp. 1–10.
- [5] Federico Aschieri, Stefan Hetzl, Daniel Weller, Expansion trees with cut, *Math. Struct. Comput. Sci.* 29 (8) (2019) 1009–1029.
- [6] Jeremy Avigad, Solomon Feferman, Gödel’s functional (“dialectica”) interpretation, in: Sam Buss (Ed.), *The Handbook of Proof Theory*, North-Holland, 1999, pp. 337–405.
- [7] Matthias Baaz, Stefan Hetzl, On the non-confluence of cut-elimination, *J. Symb. Log.* 76 (1) (2011) 313–340.
- [8] Matthias Baaz, Stefan Hetzl, Alexander Leitsch, Clemens Richter, Hendrik Spohr, Cut-elimination: experiments with CERES, in: Franz Baader, Andrei Voronkov (Eds.), *Logic for Programming, Artificial Intelligence, and Reasoning (LPAR 2004)*, in: Lecture Notes in Computer Science, vol. 3452, Springer, 2005, pp. 481–495.
- [9] Matthias Baaz, Alexander Leitsch, Cut-elimination and redundancy-elimination by resolution, *J. Symb. Comput.* 29 (2) (2000) 149–176.
- [10] Franco Barbanera, Stefano Berardi, Massimo Schivalocchi, “Classical” programming-with-proofs in λ_{PA}^{Sym} : an analysis of non-confluence, in: Martín Abadi, Takayasu Ito (Eds.), *Theoretical Aspects of Computer Software*, in: Lecture Notes in Computer Science, vol. 1281, Springer, Berlin, Heidelberg, 1997, pp. 365–390.
- [11] Arnold Beckmann, Exact bounds for lengths of reductions in typed lambda-calculus, *J. Symb. Log.* 66 (3) (2001) 1277–1285.
- [12] Sam Buss, Cut elimination in situ, in: R. Kahle, M. Rathjen (Eds.), *Gentzen’s Centenary: The Quest for Consistency*, Springer, 2015, pp. 245–277.
- [13] Thierry Coquand, A semantics of evidence for classical arithmetic, *J. Symb. Log.* 60 (1) (1995) 325–337.
- [14] Fernando Ferreira, Gilda Ferreira, A herbrandized functional interpretation of classical first-order logic, *Arch. Math. Log.* 56 (5) (Aug 2017) 523–539.
- [15] Gerhard Gentzen, Untersuchungen über das logische Schließen II, *Math. Z.* 39 (3) (1935) 405–431.
- [16] Philipp Gerhardy, Applications of proof interpretations, PhD thesis, University of Aarhus, 2006.
- [17] Philipp Gerhardy, Ulrich Kohlenbach, Extracting Herbrand disjunctions by functional interpretation, *Arch. Math. Log.* 44 (2005) 633–644.
- [18] Kurt Gödel, Über eine noch nicht benützte Erweiterung des finiten Standpunktes, *Dialectica* 12 (1958) 280–287.
- [19] Willem Heijltjes, Classical proof forestry, *Ann. Pure Appl. Log.* 161 (11) (2010) 1346–1366.
- [20] Hugo Herbelin, Séquents qu’on calcule, PhD thesis, Université Paris 7, 1995.
- [21] Stefan Hetzl, Applying tree languages in proof theory, in: Adrian-Horia Dediú, Carlos Martín-Vide (Eds.), *Language and Automata Theory and Applications (LATA 2012)*, in: Lecture Notes in Computer Science, vol. 7183, Springer, 2012, pp. 301–312.

- [22] Stefan Hetzl, Alexander Leitsch, Giselle Reis, Janos Tapolczai, Daniel Weller, Introducing quantified cuts in logic with equality, in: Stéphane Demri, Deepak Kapur, Christoph Weidenbach (Eds.), 7th International Joint Conference on Automated Reasoning (IJCAR 2014), in: *Lecture Notes in Computer Science*, vol. 8562, Springer, 2014, pp. 240–254.
- [23] Stefan Hetzl, Alexander Leitsch, Giselle Reis, Daniel Weller, Algorithmic introduction of quantified cuts, *Theor. Comput. Sci.* 549 (2014) 1–16.
- [24] Stefan Hetzl, Alexander Leitsch, Daniel Weller, Towards algorithmic cut-introduction, in: Nikolaj Bjørner, Andrei Voronkov (Eds.), *Logic for Programming, Artificial Intelligence and Reasoning (LPAR 2012)*, in: *Lecture Notes in Computer Science*, vol. 7180, Springer, 2012, pp. 228–242.
- [25] David Hilbert, Paul Bernays, *Grundlagen der Mathematik II*, Springer, 1939.
- [26] Naoki Kobayashi, Types and higher-order recursion schemes for verification of higher-order programs, in: POPL, ACM, 2009, pp. 416–428.
- [27] Alexander Leitsch, Michael Peter Lettmann, The problem of Π_2 -cut-introduction, *Theor. Comput. Sci.* 706 (2018) 83–116.
- [28] Richard McKinley, Proof nets for Herbrand’s theorem, *ACM Trans. Comput. Log.* 14 (1) (2013) 5:1–5:31.
- [29] Dale Miller, A compact representation of proofs, *Stud. Log.* 46 (4) (1987) 347–370.
- [30] Georg Moser, Richard Zach, The epsilon calculus and Herbrand complexity, *Stud. Log.* 82 (1) (2006) 133–155.
- [31] C.-H. Luke Ong, On model-checking trees generated by higher-order recursion schemes, in: LICS, IEEE Computer Society, 2006, pp. 81–90.
- [32] C.-H. Luke Ong, Steven J. Ramsay, Verifying higher-order functional programs with pattern-matching algebraic data types, in: *Proceedings of the 38th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL 2011)*, ACM, 2011, pp. 587–598.
- [33] C.-H. Luke Ong, Recursion schemes, collapsible pushdown automata and higher-order model checking, in: Adrian-Horia Dediu, Carlos Martín-Vide, Bianca Truthe (Eds.), *Language and Automata Theory and Applications (LATA 2013)*, Springer, 2013, pp. 13–41.
- [34] David Michael Ritchie Park, Fixpoint induction and proofs of program properties, in: B. Meltzer, D. Michie (Eds.), *Machine Intelligence*, vol. 5, 1970.
- [35] Diana Ratiu, Trifon Trifonov, Exploring the computational content of the infinite pigeonhole principle, *J. Log. Comput.* 22 (2) (2012) 329–350.
- [36] Helmut Schwichtenberg, Complexity of normalization in the pure typed lambda-calculus, in: Anne S. Troelstra, D. van Dalen (Eds.), *The L.E.J. Brouwer Centenary Symposium*, North-Holland, 1982, pp. 453–457.
- [37] Monika Seisenberger, On the constructive content of proofs, PhD thesis, Ludwig-Maximilians-Universität München, 2003.
- [38] Joseph R. Shoenfield, *Mathematical Logic*, 2nd edition, AK Peters, 2001.
- [39] Thomas Streicher, Ulrich Kohlenbach, Shoenfield is Gödel after Krivine, *Math. Log. Q.* 53 (2) (2007) 176–179.
- [40] Anne S. Troelstra, Helmut Schwichtenberg, *Basic Proof Theory*, Cambridge University Press, Cambridge, 1996.
- [41] Christian Urban, *Classical logic and computation*, PhD thesis, University of Cambridge, 2000.