THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Sequent calculi for induction and infinite descent

Citation for published version:

Brotherston, J & Simpson, A 2010, 'Sequent calculi for induction and infinite descent', *Journal of Logic and Computation*, vol. 21, no. 2, pp. 1177-1216. <https://doi.org/10.1093/logcom/exq052>

Digital Object Identifier (DOI):

[10.1093/logcom/exq052](https://doi.org/10.1093/logcom/exq052)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Journal of Logic and Computation

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Sequent calculi for induction and infinite descent

James Brotherston*
Dept. of Computing
Imperial College London

Alex Simpson†
LFCS, School of Informatics
University of Edinburgh

Abstract

This paper formalises and compares two different styles of reasoning with inductively defined predicates, each style being encapsulated by a corresponding sequent calculus proof system.

The first system, LKID, supports traditional proof by induction, with induction rules formulated as rules for introducing inductively defined predicates on the left of sequents. We show LKID to be cut-free complete with respect to a natural class of Henkin models; the eliminability of cut follows as a corollary.

The second system, LKID^ω , uses infinite (non-well-founded) proofs to represent arguments by infinite descent. In this system, the left-introduction rules for inductively defined predicates are simple case-split rules, and an infinitary, global condition on proof trees is required in order to ensure soundness. We show LKID^ω to be cut-free complete with respect to standard models, and again infer the eliminability of cut.

The infinitary system LKID^ω is unsuitable for formal reasoning. However, it has a natural restriction to proofs given by regular trees, i.e. to those proofs representable by finite graphs, which is so suited. We demonstrate that this restricted “cyclic” proof system, CLKID^ω , subsumes LKID, and conjecture that CLKID^ω and LKID are in fact equivalent, i.e., that proof by induction is equivalent to regular proof by infinite descent.

1 Introduction

Many concepts in mathematics are most naturally formulated using inductive definitions. Thus proof support for inductive definitions is an essential component of proof assistants and theorem provers. Often, libraries are provided containing collections of useful induction principles associated with a given set of inductive definitions, see e.g. [28, 15, 33]. In other cases, mechanisms permitting “cyclic” proof arguments are used, with intricate conditions imposed to ensure soundness, see e.g. [44, 32, 14]. These conditions can be broadly construed as versions of the well-known mathematical principle of *infinite descent* originally formalised by Fermat [22]. In this article we develop proof-theoretic foundations for this infinite descent style of inductive reasoning, and compare them with the corresponding (but quite different) foundations for proof by explicit induction.

In the case of classical first-order logic, Gentzen’s sequent calculus LK provides an elegant proof system that is well suited to the goal-directed approach to proof construction employed in many proof assistants. Each logical constant is specified by two types of basic rule, introducing

---

*Research supported by an EPSRC Postdoctoral Fellowship, and by an EPSRC PhD studentship undertaken at LFCS, School of Informatics, University of Edinburgh.

†Research supported by an EPSRC Advanced Research Fellowship (2001–7).

the constant on the left and on the right of sequents respectively. Gentzen’s well-known cut-elimination theorem for LK implies that direct proofs, using these rules alone, are sufficient to derive any valid sequent [12]. In addition to its theoretical elegance, this has implications for proof search, with the locally applicable proof rules thereby constrained by the logical constants appearing in the current goal.

In this paper, we present sequent calculus proof systems that canonically embody two standard approaches to reasoning with inductively defined predicates: (i) explicit rule induction over definitions; and (ii) infinite descent, employing a generalisation of Fermat’s original natural number principle to general inductively defined predicates. In each case, we establish appropriate completeness and cut-eliminability theorems for our proof systems. These theorems constitute the main technical contribution of this article. Aside from their intrinsic technical interest, our results demonstrate our calculi as being canonical ones embodying the two aforementioned styles of inductive reasoning. We hope that this article will help to stimulate wider interest in such systems.

In §3 we present our sequent calculus for induction, LKID, which extends Gentzen’s LK with left- and right-introduction rules for inductively defined predicates. The right-introduction rules for an inductively defined predicate P simply reflect the closure conditions in the definition of P , while the left-introduction rules embody the natural induction principle associated with P . A closely related precursor is Martin-Löf’s natural deduction system for intuitionistic logic with (iterated) inductive definitions [23], in which induction rules are included as elimination rules for inductively defined predicates. As is well known, elimination rules in natural deduction serve the same purpose as left-introduction rules in sequent calculus. Nonetheless, it is only relatively recently that sequent calculus counterparts of Martin-Löf’s system have been explicitly considered, by McDowell, Miller, Momigliano and Tiu [24, 25, 40]. LKID is a natural classical analogue of these intuitionistic systems.

For LKID, we prove soundness and completeness relative to a natural class of “Henkin models” for inductive predicates. In fact, completeness is established for the cut-free fragment of LKID, and so the eliminability of cut follows as an immediate corollary. These results serve to endorse the canonicity of LKID: completeness shows that no proof principles are missing, and cut-eliminability vindicates the formulation of the proof rules. The eliminability of cut in LKID holds in opposition to the not uncommon belief that full cut-elimination results are impossible in the presence of inductive definitions. In fact, the real limitation is not the impossibility of cut-elimination, but rather that the subformula property is not achievable [20], and indeed the subformula property does not hold for cut-free proofs in LKID. In fact, the possibility of obtaining a full cut-elimination result is not surprising if one is familiar with the aforementioned literature on intuitionistic systems with inductive definitions, where similar normalization/cut-elimination results appear [23, 24, 40]. The proofs of normalization/cut-elimination for the intuitionistic systems, in the literature, are all based on Tait’s “computability” method, which does not easily adapt to a two-sided classical sequent calculus like LKID. Compared with such proofs, however, our semantic approach suffers from the weakness of not establishing that any particular cut-elimination strategy terminates. Of course, the use of such semantic methods to establish cut-eliminability is not new. For example, the original proof of Takeuti’s Conjecture (the eliminability of cut in second-order logic) was semantic [38, 13]. However, compared with the semantic proof of Takeuti’s Conjecture, the class of Henkin models we consider seems a natural class of structures, and our completeness result is thus of interest in its own right. We give our completeness proof for LKID in §4.

The remainder of the paper covers in detail our formalisation of infinite descent. For natural numbers, infinite descent exploits the fact that, since there are no infinite strictly decreasing sequences of numbers, any case in a proof that furnishes such a sequence can be ignored as con-

tradictory. This technique can be extended to general inductively defined predicates: any case of a proof which yields an infinite sequence of “unfoldings” of some inductively defined predicate can be similarly dismissed by appealing to the well-foundedness of its chain of approximants. In §5, we formulate a proof system, LKID^ω , in which this principle is implemented. In LKID^ω , the induction rules of LKID are replaced by simple “case-split” rules (which unfold inductively defined predicates on the left of sequents), and proofs are allowed to be infinite (non-well-founded) derivation trees, as opposed to the usual finite derivations. In general, such infinite derivations are not sound, so we impose a *global trace condition* on infinite derivation trees (similar to conditions employed in infinitary μ -calculus proof systems, e.g. [27]) that qualifies such trees as *bona fide* proofs. Informally, this condition states that, for every infinite branch in the tree, some inductively-defined predicate must be unfolded infinitely often along the left-hand side of the sequents on the branch. The precise technical formulation is given in §5. Whereas the soundness proof for LKID, of §3, was relative to a class of Henkin models, the soundness proof for LKID^ω is relative to the subclass of ordinary “standard” models of the inductively defined predicates. This restriction is essential. We show that LKID^ω is complete relative to standard models, and hence strictly stronger than LKID. Once again, our completeness argument establishes completeness for cut-free proofs, and so the eliminability of cut for LKID^ω follows. The proof of completeness is given in §6.

The infinitary system LKID^ω , is unsurprisingly not suitable for practical formal reasoning. In particular, it is impossible to recursively enumerate a complete set of LKID^ω proofs. Nevertheless, LKID^ω does have a natural subsystem that is suitable for formal reasoning, namely the restriction of LKID^ω to *regular* proofs, i.e., to those infinite derivation trees that are representable by a finite (cyclic) graph. We call such proofs “cyclic proofs”. In §7, we formally define the system CLKID^ω of cyclic proofs. The suitability of CLKID^ω for formal reasoning is assured by the fact that the global trace condition is decidable over the graphs generating cyclic proofs. As a result of this, the completeness property of LKID^ω is necessarily lost in the restriction to cyclic proofs: CLKID^ω is unavoidably weaker than LKID^ω .

Since LKID and CLKID^ω are alternative proof systems for formal reasoning about inductive definitions, it is natural to consider the relationship between them. We show that any sequent provable in LKID is also provable in CLKID^ω . Thus cyclic proof subsumes proof by induction for inductively defined predicates. Although we have not been able to establish the converse implication, we conjecture that LKID and CLKID^ω are actually equivalent in power. If one accepts that LKID and CLKID^ω are canonical embodiments of, respectively, proof by induction and regular proof by infinite descent, then the conjecture can be understood as a formal assertion of the equivalence of these two proof styles. We end the paper by stating this conjecture and commenting on the apparent difficulties its proof poses.

Briefly, the structure of the paper is as follows. Section 2 introduces the first-order logic with inductive definitions, FOL_{ID} , that we shall use throughout the paper. Section 3 presents the sequent calculus for induction, LKID, and Section 4 provides its completeness proof. Section 5 presents the infinitary sequent calculus for infinite descent, LKID^ω , and Section 6 gives its completeness proof. Section 7 presents the cyclic subsystem CLKID^ω of LKID^ω . Section 8 concludes. Readers who are not interested in the details of the completeness proofs are advised to skip the technical Sections 4 and 6. Appendix A presents technical details from the soundness proof for LKID.

Earlier incarnations of the results in this paper were outlined in two conference papers [3, 7], and appear in the first author’s PhD thesis [4]. We are grateful to the two anonymous referees for comments that have helped to improve the journal presentation of this work.

2 Syntax and semantics of first-order logic with inductive definitions (FOL_{ID})

In this section we give the syntax and semantics of classical first-order logic with inductively defined predicates, FOL_{ID}. Of the many possible frameworks for inductive definitions, we choose to work with ordinary (mutual) inductive definitions, specified by simple “productions” in the style of Martin-Löf [23]. This choice keeps the logic relatively simple, while encompassing many important examples.

The languages we consider are the standard (countable) first-order languages, except that we designate finitely many of the predicate symbols of the language as *inductive*. A predicate symbol not designated as inductive is called *ordinary*. For the remainder of this paper we consider a fixed language Σ with inductive predicate symbols $P_1, \dots, P_n$. Terms of Σ are defined as usual; we write $t(x_1, \dots, x_m)$ for a term all of whose variables are contained in $\{x_1, \dots, x_m\}$.

The interpretation of the elements of Σ is as usual given by a first-order structure M with domain D ; we write α^M to denote the interpretation of the Σ -symbol α in M . If $t(x_1, \dots, x_m)$ is a term we write $t^M : D^m \rightarrow D$ for the function obtained by replacing each function symbol f in t by its interpretation f^M . Variables are interpreted as elements of D by an environment ρ (which is simply a total function from variables to D); we extend ρ to all terms of Σ in the standard way and write $\rho[x \mapsto d]$ for an environment defined exactly as ρ except that $\rho[x \mapsto d](x) = d$. The formulas of FOL_{ID} are the usual formulas of first-order logic with equality. We then write $M \models_\rho F$ for the standard semantic satisfaction relation for formulas of FOL_{ID}.

Our proof systems will be interpreted relative to only those structures in which inductive predicates have their intended meanings, as specified by *definition sets* for the predicates, adapted from [23].

Definition 2.1 (Inductive definition set). An *inductive definition set* Φ for Σ is a finite set of productions, where a *production* is a pair

$$\langle \{Q_1(\mathbf{u}_1), \dots, Q_h(\mathbf{u}_h), P_{j_1}(\mathbf{t}_1), \dots, P_{j_m}(\mathbf{t}_m)\}, P_i(\mathbf{t}) \rangle$$

of a finite set of atomic formulas together with an atomic formula, where $j_1, \dots, j_m, i \in \{1, \dots, n\}$, $Q_1, \dots, Q_h$ are ordinary predicate symbols, and the bold vector notation abbreviates sequences of terms of appropriate length to match the arities of the predicate symbols. Following [23], we depict such a production

$$\frac{Q_1(\mathbf{u}_1) \dots Q_h(\mathbf{u}_h) P_{j_1}(\mathbf{t}_1) \dots P_{j_m}(\mathbf{t}_m)}{P_i(\mathbf{t})} \quad (1)$$

We call the formulas above the line the *premises* of the production, and the formula below the line the *conclusion*. Often, we shall omit brackets, writing, e.g., $Q_1\mathbf{u}_1$ for $Q_1(\mathbf{u}_1)$. Sometimes it will be convenient to make explicit the variables occurring in a production, and so we shall write formulas as, e.g., $Q_1\mathbf{u}_1(\mathbf{x})$ (where $\mathbf{x}$ is a vector of variables).

Example 2.2. We define the predicates N, E and O via the productions:

$$\frac{}{N0} \quad \frac{Nx}{Nsx} \quad \frac{}{E0} \quad \frac{Ex}{Osx} \quad \frac{Ox}{Esx}$$

In structures in which all “numerals” $s^k 0$ for $k \geq 0$ are interpreted as distinct elements, the predicates N, E and O correspond to the properties of being a natural, even and odd number respectively.

One possible generalisation of Definition 2.1 would be to systems of iterated inductive definitions as considered, e.g., by Martin-Löf [23]. In such schemas, logically complex formulas are allowed to occur in the premises of productions, subject to a suitable stratification of predicate symbols into “levels” which is necessary to ensure monotonicity of the resulting definitions.

From this point onwards we consider an arbitrary fixed inductive definition set Φ for Σ and, when we need to consider an arbitrary production in Φ , will always use the explicit format of (1) above.

The standard interpretation of the inductive predicates (cf. [1]) is obtained as usual by considering prefixed points of a monotone operator constructed from the definition set Φ . It is standard that the least prefixed point of this operator can be constructed in iterative *approximant* stages, indexed by ordinals.

Definition 2.3 (Definition set operator). Let M with domain D be a first-order structure for Σ , and for each $i \in \{1, \dots, n\}$, let k_i be the arity of the inductive predicate symbol P_i . Partition Φ into disjoint subsets $\Phi_1, \dots, \Phi_n \subseteq \Phi$ by:

$$\Phi_i = \{\phi \in \Phi \mid P_i \text{ is the inductive predicate symbol in the conclusion of production } \phi\}$$

Let each rule set Φ_i be indexed by r with $1 \leq r \leq |\Phi_i|$, and for each rule $\Phi_{i,r}$ of the form (1) in Definition 2.1, define a corresponding function $\varphi_{i,r} : \mathcal{P}(D^{k_1}) \times \dots \times \mathcal{P}(D^{k_n}) \rightarrow \mathcal{P}(D^{k_i})$, where $\mathcal{P}(\cdot)$ is powerset, by:

$$\varphi_{i,r}(X_1, \dots, X_n) = \{\rho(\mathbf{t}) \mid Q_1^M(\rho(\mathbf{u}_1)), \dots, Q_h^M(\rho(\mathbf{u}_h)), \\ \rho(\mathbf{t}_1) \in X_{j_1}, \dots, \rho(\mathbf{t}_m) \in X_{j_m}, \rho \text{ an environment}\}$$

where environments ρ are applied pointwise on vectors of terms. (Note that if $\mathbf{t}(\mathbf{x})$ is a vector of terms each with variables in $\mathbf{x}$ then we have $\rho(\mathbf{t}(\mathbf{x})) = \mathbf{t}^M(\rho(\mathbf{x}))$.) Then define the function φ_i for each $i \in \{1, \dots, n\}$ by:

$$\varphi_i(X_1, \dots, X_n) = \bigcup_r \varphi_{i,r}(X_1, \dots, X_n)$$

whence the *definition set operator* for Φ is the operator φ_Φ , with domain and codomain $\mathcal{P}(D^{k_1}) \times \dots \times \mathcal{P}(D^{k_n})$, defined by:

$$\varphi_\Phi(X_1, \dots, X_n) = (\varphi_1(X_1, \dots, X_n), \dots, \varphi_n(X_1, \dots, X_n))$$

Henceforth, we write π_i^n for the i th projection function given by $\pi_i^n(X_1, \dots, X_n) = X_i$, and we extend union and subset inclusion to the corresponding pointwise operations on n -tuples of sets.

Definition 2.4 (Approximants). Let M with domain D be a first-order structure for Σ , and let φ_Φ be the definition set operator for Φ . Define an ordinal-indexed set $(\varphi_\Phi^\alpha \subseteq \mathcal{P}(D^{k_1}) \times \dots \times \mathcal{P}(D^{k_n}))_{\alpha \geq 0}$ by $\varphi_\Phi^\alpha = \bigcup_{\beta < \alpha} \varphi_\Phi(\varphi_\Phi^\beta)$ (note that this implies $\varphi_\Phi^0 = (\emptyset, \dots, \emptyset)$). Then the set $\pi_i^n(\varphi_\Phi^\alpha)$ is called the α^{th} *approximant* of P_i , written as P_i^α .

A *prefixed point* of φ_Φ is a tuple $(X_1, \dots, X_n)$ satisfying $\varphi_\Phi(X_1, \dots, X_n) \subseteq (X_1, \dots, X_n)$. It is a standard result for inductive definitions that the least prefixed point of φ_Φ is given by $\bigcup_\alpha \varphi_\Phi^\alpha$, the union of the approximants of the inductive predicates $(P_1, \dots, P_n)$.

Definition 2.5 (Standard model). A first-order structure M is said to be a *standard model* for (Σ, Φ) if $P_i^M = \bigcup_\alpha P_i^\alpha$ for all $i \in \{1, \dots, n\}$.¹

¹For the form of production considered, we have $\bigcup_\alpha P_i^\alpha = P_i^\omega$, i.e. the closure ordinal of our inductive definitions is at most ω . However, we shall never exploit this fact.

Definition 2.5 fixes a standard interpretation of the inductive predicates. However, we shall also be interested in non-standard *Henkin models* of FOL_{ID} in which the least fixed point of the definition set operator φ_Φ is constructed with respect to a chosen class of sets of tuples over the domain of interpretation. This approach is based on the well-known idea of Henkin for obtaining completeness theorems for higher-order calculi by considering validity with respect to a more general notion of model [17]. Our application in §3 is similar.

Definition 2.6 (Henkin class). Let M with domain D be a structure for Σ . A *Henkin class* for M is a family of sets $\mathcal{H} = \{H_k \subseteq \mathcal{P}(D^k) \mid k \in \mathbb{N}\}$ such that, for each $k \in \mathbb{N}$:

- (H1) $\{(d, d) \mid d \in D\} \in H_2$;
- (H2) if Q is any predicate symbol of arity k then $\{(d_1, \dots, d_k) \mid Q^M(d_1, \dots, d_k)\} \in H_k$;
- (H3) if $R \in H_{k+1}$ and $d \in D$ then $\{(d_1, \dots, d_k) \mid (d_1, \dots, d_k, d) \in R\} \in H_k$;
- (H4) if $R \in H_k$ and $t_1(x_1, \dots, x_m), \dots, t_k(x_1, \dots, x_m)$ are terms then:
 $\{(d_1, \dots, d_m) \mid (t_1^M(d_1, \dots, d_m), \dots, t_k^M(d_1, \dots, d_m)) \in R\} \in H_m$;
- (H5) if $R \in H_k$ then $\overline{R} = D^k \setminus R \in H_k$;
- (H6) if $R_1, R_2 \in H_k$ then $R_1 \cap R_2 \in H_k$;
- (H7) if $R \in H_{k+1}$ then $\{(d_1, \dots, d_k) \mid \exists d. (d_1, \dots, d_k, d) \in R\} \in H_k$.

The following lemma demonstrates that our Henkin classes contain enough sets of tuples to interpret any formula of FOL_{ID} .

Lemma 2.7. *If $\mathcal{H} = \{H_k \mid k \in \mathbb{N}\}$ is a Henkin class for a structure M , ρ is an environment for M , F is a formula of FOL_{ID} and $x_1, \dots, x_k$ are distinct variables, then:*

$$\{(d_1, \dots, d_k) \mid M \models_{\rho[x_1 \mapsto d_1, \dots, x_k \mapsto d_k]} F\} \in H_k$$

Proof. The case when $\{x_1, \dots, x_k\} \subseteq \text{FV}(F)$ is an induction on the structure of the formula F . This case can then be straightforwardly lifted to the case when $x_1, \dots, x_k$ are arbitrary variables. The full details can be found as Proposition 2.3.3 in [4]. $\square$

Definition 2.8 ($\mathcal{H}$ -point). Let M be a structure for Σ and let $\mathcal{H}$ be a Henkin class for M . Also let k_i be the arity of the inductive predicate symbol P_i for each $i \in \{1, \dots, n\}$. Then $(X_1, \dots, X_n) \in \mathcal{P}(D_1^{k_1}) \times \dots \times \mathcal{P}(D_n^{k_n})$ is said to be an $\mathcal{H}$ -point if $X_i \in H_{k_i}$ for each $i \in \{1, \dots, n\}$.

A *prefixed $\mathcal{H}$ -point* of a definition set operator φ_Φ is simply a prefixed point of φ_Φ which is also a $\mathcal{H}$ -point. The next lemma shows that $\mathcal{H}$ -points are closed under definition set operators.

Lemma 2.9. *Let $\mathcal{H}$ be a Henkin class for a Σ -structure M . If $(X_1, \dots, X_n)$ is an $\mathcal{H}$ -point then so is $\varphi_\Phi(X_1, \dots, X_n)$.*

Proof. Since Henkin classes are easily seen to be closed under union, it suffices to show that $\varphi_{i,r}(X_1, \dots, X_n) \in H_{k_i}$, where $\varphi_{i,r}$ is the function corresponding to an arbitrary production $\Phi_{i,r} \in \Phi$ and k_i is the arity of the predicate P_i in the conclusion of the production. This follows from the closure conditions on Henkin classes given in Definition 2.6. The full details appear as Lemma 2.3.6 in [4]. $\square$

Definition 2.10 (Henkin model). Let M be a first-order structure for Σ and $\mathcal{H}$ be a Henkin class for M . The data $(M, \mathcal{H})$ comprises a *Henkin model* for (Σ, Φ) if there exists a least prefixed $\mathcal{H}$ -point $\mu_{\mathcal{H}}.\varphi_{\Phi}$ of φ_{Φ} , and for each $i \in \{1, \dots, n\}$, $P_i^M = \pi_i^n(\mu_{\mathcal{H}}.\varphi_{\Phi})$. (We remark that $\mu_{\mathcal{H}}.\varphi_{\Phi}$, if it exists, is a fixed point of φ_{Φ} .)

As mentioned previously, $\bigcup_{\alpha} \varphi_{\Phi}^{\alpha}$ is the least prefixed point of φ_{Φ} in $\mathcal{P}(D^{k_1}) \times \dots \times \mathcal{P}(D^{k_n})$. Thus a standard model is in particular a Henkin model (with $H_k = \mathcal{P}(D^k)$ for each $k \in \mathbb{N}$).

3 LKID: a proof system for induction in FOL_{ID}

In this section, we formulate a proof system, LKID, formalising the usual notion of proof by induction for FOL_{ID} . We then prove soundness and cut-free completeness of LKID with respect to the Henkin semantics of FOL_{ID} given in the previous section, and infer from these results the eliminability of cut in LKID.

We write sequents of the form $\Gamma \vdash \Delta$ where Γ, Δ are finite sets of formulas, and use the notation $\Gamma[\theta]$ to mean that the substitution θ of terms for free variables is applied to all formulas in Γ .

For first-order logic with equality, we use the (standard) sequent calculus rules, given in Figure 1. By the *principal formula* of a rule instance, we mean the distinguished formula that is introduced by the rule into its conclusion. We remark that we read the comma in sequents as set union, so that contraction rules are unnecessary. For example, we need not delete the principal formula of a rule instance when applying the rule backwards. Somewhat unusually, we include a rule for explicit substitution, and rules for equality (cf. [11]). Although these rules are inessential inclusions in LKID, they will prove useful in our infinitary proof systems for infinite descent in FOL_{ID} , introduced later. To the rules in Figure 1 we add rules for introducing atomic formulas involving inductive predicates on the left and right of sequents.

First, for each production $\Phi_{i,r} \in \Phi$, each necessarily in the format of (1),

$$\frac{Q_1 \mathbf{u}_1(\mathbf{x}) \dots Q_h \mathbf{u}_h(\mathbf{x}) P_{j_1} \mathbf{t}_1(\mathbf{x}) \dots P_{j_m} \mathbf{t}_m(\mathbf{x})}{P_i \mathbf{t}(\mathbf{x})}$$

we include a corresponding sequent calculus right introduction rule for P_i :

$$\frac{\Gamma \vdash Q_1 \mathbf{u}_1(\mathbf{u}), \Delta \quad \dots \quad \Gamma \vdash Q_h \mathbf{u}_h(\mathbf{u}), \Delta \quad \Gamma \vdash P_{j_1} \mathbf{t}_1(\mathbf{u}), \Delta \quad \dots \quad \Gamma \vdash P_{j_m} \mathbf{t}_m(\mathbf{u}), \Delta}{\Gamma \vdash P_i \mathbf{t}(\mathbf{u}), \Delta} (P_i R_r)$$

Here $\mathbf{u}$ is assumed to be a vector of terms of the same length as the vector $\mathbf{x}$ of variables explicitly identified as occurring in the production, and the occurrences of $\mathbf{u}$ in the rule above represent the substitution $[\mathbf{u}/\mathbf{x}]$.

The left-introduction rules for inductively defined predicates manifest themselves as induction rules. In order to formulate these rules correctly and without redundant premises, we first need to define *mutual dependency* between predicates arising from their definitions, as used in [23] (although our formulation slightly improves the notion used there by eliminating redundant premises from some induction rules).

Definition 3.1 (Mutual dependency). Define the binary relation *Prem* on the inductive predicate symbols of Σ as the least relation satisfying: whenever P_i occurs in the conclusion of some production in Φ , and P_j occurs amongst the premises of that production, then $\text{Prem}(P_i, P_j)$ holds. Also define Prem^* to be the reflexive-transitive closure of *Prem*. Then two predicate symbols P_i and P_j are *mutually dependent* if both $\text{Prem}^*(P_i, P_j)$ and $\text{Prem}^*(P_j, P_i)$ hold.

Structural rules:

$$\begin{array}{c}
\frac{}{\Gamma \vdash \Delta} \quad \Gamma \cap \Delta \neq \emptyset \quad (\text{Axiom}) \\
\\
\frac{\Gamma \vdash F, \Delta \quad \Gamma, F \vdash \Delta}{\Gamma \vdash \Delta} \quad (\text{Cut}) \\
\\
\frac{\Gamma' \vdash \Delta' \quad \Gamma' \subseteq \Gamma \quad \Delta' \subseteq \Delta}{\Gamma \vdash \Delta} \quad (\text{Wk}) \\
\\
\frac{\Gamma \vdash \Delta}{\Gamma[\theta] \vdash \Delta[\theta]} \quad (\text{Subst})
\end{array}$$

Logical rules:

$$\begin{array}{c}
\frac{\Gamma \vdash F, \Delta}{\Gamma, \neg F \vdash \Delta} \quad (\neg L) \\
\\
\frac{\Gamma, F \vdash \Delta \quad \Gamma, G \vdash \Delta}{\Gamma, F \vee G \vdash \Delta} \quad (\vee L) \\
\\
\frac{\Gamma, F, G \vdash \Delta}{\Gamma, F \wedge G \vdash \Delta} \quad (\wedge L) \\
\\
\frac{\Gamma \vdash F, \Delta \quad \Gamma, G \vdash \Delta}{\Gamma, F \rightarrow G \vdash \Delta} \quad (\rightarrow L) \\
\\
\frac{\Gamma, F[t/x] \vdash \Delta}{\Gamma, \forall x F \vdash \Delta} \quad (\forall L) \\
\\
\frac{\Gamma, F \vdash \Delta}{\Gamma, \exists x F \vdash \Delta} \quad x \notin FV(\Gamma \cup \Delta) \quad (\exists L) \\
\\
\frac{\Gamma[u/x, t/y] \vdash \Delta[u/x, t/y]}{\Gamma[t/x, u/y], t = u \vdash \Delta[t/x, u/y]} \quad (=L) \\
\\
\frac{\Gamma, F \vdash \Delta}{\Gamma \vdash \neg F, \Delta} \quad (\neg R) \\
\\
\frac{\Gamma \vdash F, G, \Delta}{\Gamma \vdash F \vee G, \Delta} \quad (\vee R) \\
\\
\frac{\Gamma \vdash F, \Delta \quad \Gamma \vdash G, \Delta}{\Gamma \vdash F \wedge G, \Delta} \quad (\wedge R) \\
\\
\frac{\Gamma, F \vdash G, \Delta}{\Gamma \vdash F \rightarrow G, \Delta} \quad (\rightarrow R) \\
\\
\frac{\Gamma \vdash F, \Delta}{\Gamma \vdash \forall x F, \Delta} \quad x \notin FV(\Gamma \cup \Delta) \quad (\forall R) \\
\\
\frac{\Gamma \vdash F[t/x], \Delta}{\Gamma \vdash \exists x F, \Delta} \quad (\exists R) \\
\\
\frac{}{\Gamma \vdash t = t, \Delta} \quad (=R)
\end{array}$$

Figure 1: Sequent calculus proof rules for classical first-order logic with equality.

Now to obtain an instance of the left-introduction rule for any inductive predicate P_j , we first associate with every inductive predicate P_i a tuple $\mathbf{z}_i$ of k_i distinct variables (called *induction variables*), where k_i is the arity of P_i . Furthermore, we associate to every predicate P_i that is mutually dependent with P_j an arbitrary formula (called an *induction hypothesis*) F_i , possibly containing (some of) the induction variables $\mathbf{z}_i$. Next, define the formula G_i for each $i \in \{1, \dots, n\}$ by:

$$G_i = \begin{cases} F_i & \text{if } P_i \text{ and } P_j \text{ are mutually dependent} \\ P_i(\mathbf{z}_i) & \text{otherwise} \end{cases}$$

We write $G_i(\mathbf{t})$, where $\mathbf{t}$ is any tuple of k_i terms, to mean $G_i[\mathbf{t}/\mathbf{z}_i]$ (and similarly for F_i). Then an instance of the induction rule for P_j has the following schema:

$$\frac{\text{minor premises} \quad \Gamma, F_j(\mathbf{u}) \vdash \Delta}{\Gamma, P_j(\mathbf{u}) \vdash \Delta} (\text{Ind } P_j)$$

where the premise $\Gamma, F_j(\mathbf{u}) \vdash \Delta$ is called the *major premise* of the rule instance, and for each production of Φ having in its conclusion a predicate P_i that is mutually dependent with P_j , say:

$$\frac{Q_1 \mathbf{u}_1(\mathbf{x}) \dots Q_h \mathbf{u}_h(\mathbf{x}) P_{j_1} \mathbf{t}_1(\mathbf{x}) \dots P_{j_m} \mathbf{t}_m(\mathbf{x})}{P_i \mathbf{t}(\mathbf{x})}$$

there is a corresponding *minor premise*:

$$\Gamma, Q_1 \mathbf{u}_1(\mathbf{y}), \dots, Q_h \mathbf{u}_h(\mathbf{y}), G_{j_1} \mathbf{t}_1(\mathbf{y}), \dots, G_{j_m} \mathbf{t}_m(\mathbf{y}) \vdash F_i \mathbf{t}(\mathbf{y}), \Delta$$

where $\mathbf{y}$ is a vector of distinct variables of the same length as the vector $\mathbf{x}$ of variables explicitly identified in the production, and $y \notin FV(\Gamma \cup \Delta \cup \{P_j \mathbf{u}\})$ for all $y \in \mathbf{y}$ ($FV(\cdot)$ being the usual free variable function on sets of formulas).

The induction rule for a predicate P_j can be seen to embody the natural principle of rule induction over the productions defining P_j .

Example 3.2. The induction rule for the “natural number” predicate N defined in Example 2.2 is:

$$\frac{\Gamma \vdash F0, \Delta \quad \Gamma, Fx \vdash Fsx, \Delta \quad \Gamma, Ft \vdash \Delta}{\Gamma, Nt \vdash \Delta} (\text{Ind } N)$$

where x is fresh and F is the induction hypothesis associated with the predicate N . This is one way of writing the usual induction scheme for N in sequent calculus style.

Example 3.3. The induction rule for the “even number” predicate E defined in Example 2.2 is:

$$\frac{\Gamma \vdash F_E 0, \Delta \quad \Gamma, F_E x \vdash F_O sx, \Delta \quad \Gamma, F_O x \vdash F_E sx, \Delta \quad \Gamma, F_E t \vdash \Delta}{\Gamma, E t \vdash \Delta} (\text{Ind } E)$$

where x is fresh and F_E and F_O are the induction hypotheses associated with the (mutually dependent) predicates E and O respectively.

Definition 3.4 (Henkin validity / Validity). Let $(M, \mathcal{H})$ be a Henkin model for (Σ, Φ) . A sequent $\Gamma \vdash \Delta$ is said to be *true in* $(M, \mathcal{H})$ if, for all environments ρ , whenever $M \models_\rho J$ for all $J \in \Gamma$ then $M \models_\rho K$ for some $K \in \Delta$. A sequent is said to be *Henkin valid* if it is true in all Henkin models. A sequent is said to be *valid* if it is true in all standard models.

Since every standard model is a Henkin model, it is clear that Henkin validity implies validity.

By a *derivation tree*, we mean a tree of sequents in which each parent sequent is obtained as the conclusion of an inference rule with its children as premises. As usual, a *proof* in LKID is a finite derivation tree all of whose branches end in an axiom (i.e. a proof rule with no premises).

Proposition 3.5 (Henkin soundness of LKID). *If there is an LKID proof of $\Gamma \vdash \Delta$ then $\Gamma \vdash \Delta$ is Henkin valid.*

Proof. Soundness follows as usual from the local soundness of each proof rule. The proofs of local soundness are straightforward in all cases except that of the induction rules. For this case, although the proof goes through roughly as expected, dealing correctly with possible mutual dependency between predicates is delicate, and we therefore include the details as Appendix A of this paper. Full details of all cases appear as Lemma 3.2.2 in [4]. $\square$

We say that a sequent $\Gamma \vdash \Delta$ is *cut-free provable* iff there is an LKID proof of $\Gamma \vdash \Delta$ that does not contain any instances of the cut, weakening or substitution rules. Our main result about LKID is the following.

Theorem 3.6 (Cut-free Henkin completeness of LKID). *If $\Gamma \vdash \Delta$ is Henkin valid, then it is cut-free provable in LKID.*

The detailed proof of Theorem 3.6 is postponed to §4.

Corollary 3.7 (Eliminability of cut for LKID). *If $\Gamma \vdash \Delta$ is provable in LKID then it is cut-free provable.*

Proof. If $\Gamma \vdash \Delta$ is provable in LKID, it is Henkin valid by soundness (Proposition 3.5), and hence cut-free provable in LKID by Theorem 3.6. $\square$

Although cut is eliminable, LKID does not enjoy the subformula property because of the induction rules. This is an unavoidable phenomenon, and corresponds to the well-known need for generalising induction hypotheses in inductive arguments (an issue which causes serious trouble for theorem provers [8]). Nevertheless, cut-eliminability for LKID is potentially a useful property for constraining proof search; see [24] for related discussion in the intuitionistic case.

There are two natural questions arising from our completeness and cut-eliminability results for LKID. The first is whether LKID might be complete with respect to standard validity, rather than Henkin validity. The second is whether a syntactic proof of cut-elimination would be feasible, as opposed to our semantic proof. As we show below, it is possible to encode true arithmetic as sequents which are valid with respect to standard models (of a suitably chosen arithmetical language and inductive definition of the natural numbers). It thus follows from Gödel's incompleteness theorem that LKID is incomplete with respect to standard models, as indeed is any effective proof system for inductive definitions. We also show that the eliminability of cut in LKID implies the consistency of Peano Arithmetic, so there can be no straightforward combinatorial proof of Corollary 3.7.

Definition 3.8 (Peano Arithmetic / True Arithmetic). Let Σ_{PA} be the first-order language consisting of the constant symbol 0, unary function symbol s , and binary function symbols $\cdot$ and $+$. Then *Peano Arithmetic* (PA) is the theory in the language Σ_{PA} axiomatized by the usual six *Peano axioms* (PA1)–(PA6), plus the induction schema; see, e.g., [13]. *True Arithmetic* is given by the theory of the first-order structure $\mathcal{N}$ for Σ_{PA} whose domain is the natural numbers $\mathbb{N}$ and in which 0, s , $+$ and $\cdot$ have their standard arithmetical interpretations.

Definition 3.9 (TA-model). Let Σ'_{PA} be the language obtained by extending Σ_{PA} with a unary inductive predicate symbol N , and let Φ_N be the inductive definition set consisting of the “natural number” productions for N defined in Example 2.2. A TA-model is a standard model for (Σ'_{PA}, Φ_N) which satisfies the first six Peano axioms (PA1)–(PA6).

Note that in any TA-model M we have $N^M \cong \mathbb{N}$ (to simplify notation, we henceforth assume equality), since $N^M = \{(s^M)^n 0^M \mid n \in \mathbb{N}\}$ by the fact that M is a standard model of (Σ'_{PA}, Φ_N) , whence the first two Peano axioms ensure that all the elements of N^M are distinct.

Definition 3.10. Define the function $\ulcorner - \urcorner$ from Σ_{PA} -formulas to Σ'_{PA} -formulas by relativising the quantifiers so that they range over the interpretation of the inductive predicate N , explicitly:

$$\begin{aligned} \ulcorner t = u \urcorner &= t = u \\ \ulcorner \neg F \urcorner &= \neg \ulcorner F \urcorner \\ \ulcorner F_1 * F_2 \urcorner &= \ulcorner F_1 \urcorner * \ulcorner F_2 \urcorner \quad (* \in \{\wedge, \vee, \rightarrow\}) \\ \ulcorner \forall x F \urcorner &= \forall x (Nx \rightarrow \ulcorner F \urcorner) \\ \ulcorner \exists x F \urcorner &= \exists x (Nx \wedge \ulcorner F \urcorner) \end{aligned}$$

Lemma 3.11. Let M be a TA-model with domain D , and let ρ be an environment for $\mathcal{N}$ (and thus also for M since $\mathbb{N} \subseteq D$ by our observation above). Then, for any Σ_{PA} -formula F , we have $M \models_\rho \ulcorner F \urcorner$ iff $\mathcal{N} \models_\rho F$.

Proof. A straightforward structural induction on F . □

Lemma 3.12. A Σ_{PA} -formula F with $FV(F) \subseteq \{x_1, \dots, x_k\}$ is a statement of true arithmetic iff the sequent:

$$(PA1), \dots, (PA6), Nx_1, \dots, Nx_k \vdash \ulcorner F \urcorner$$

is valid with respect to standard models of (Σ'_{PA}, Φ_N) .

Proof. By definition, F is in true arithmetic iff $\mathcal{N} \models_\rho F$ for all $\mathcal{N}$ -environments ρ . Thus, by Lemma 3.11, F is in true arithmetic iff $M \models_\rho \ulcorner F \urcorner$ for all TA-models M and M -environments ρ such that $\rho(x_i) \in \mathbb{N}$ for all $1 \leq i \leq k$. Since $N^M = \mathbb{N}$ in TA-models, the latter holds iff the sequent $Nx_1, \dots, Nx_k \vdash \ulcorner F \urcorner$ is valid with respect to TA-models. But since (PA1), ..., (PA6) are closed first-order formulas, this is the case exactly if $(PA1), \dots, (PA6), Nx_1, \dots, Nx_k \vdash \ulcorner F \urcorner$ is valid for standard models of (Σ'_{PA}, Φ_N) . □

Theorem 3.13 (Standard incompleteness of LKID). *There are sequents that are valid, but unprovable in LKID.*

Proof. We first note that, since standard models are Henkin models, LKID is sound for standard models by Proposition 3.5, i.e., any provable sequent is valid. Assuming for contradiction that LKID is complete for standard models, we then have that a sequent is LKID-provable iff it is valid.

Now, fixing the language Σ'_{PA} and definition set Φ_N , for any formula F in true arithmetic with $FV(F) \subseteq \{x_1, \dots, x_k\}$ the sequent:

$$(PA1), \dots, (PA6), Nx_1, \dots, Nx_k \vdash \ulcorner F \urcorner$$

is valid with respect to standard models of (Σ'_{PA}, Φ_N) by Lemma 3.12 and hence provable by completeness. We can clearly recursively enumerate the LKID proofs, so we can construct a recursive enumeration of the sequents that are valid with respect to standard models of (Σ'_{PA}, Φ_N) . Thus, by discarding any sequents not of the form above, we can obtain a recursive enumeration of true arithmetic, which is known to be impossible. We conclude by contradiction that LKID cannot be complete with respect to standard models. □

Our next result shows that our embedding of true arithmetic in FOL_{ID} also gives an embedding of Peano arithmetic in LKID. To prove the result, we instantiate “provability in PA” as meaning provability (with respect to the signature Σ_{PA}) in the sequent calculus LK_e for first-order logic with equality given by Figure 1, extended with axiom sequents $\vdash A$, for each axiom A from the Peano axioms (PA1)–(PA6), and for each instance A of the induction schema.

Lemma 3.14. $\Gamma \vdash \Delta$ is provable in PA, where $FV(\Gamma \cup \Delta) \subseteq \{x_1, \dots, x_k\}$, if and only if the sequent $Nx_1, \dots, Nx_k, \ulcorner \Gamma \urcorner \vdash \ulcorner \Delta \urcorner$ is provable in $\text{LKID} + (\text{PA1})\text{--}(\text{PA6})$, where the language for LKID is Σ'_{PA} and the inductive definition set is Φ_N .

Proof. ($\Rightarrow$) By induction on the height of the PA derivation of $\Gamma \vdash \Delta$. For the PA axioms, we note that for each $i \in \{1, \dots, 6\}$ the sequent $(\text{PA}i) \vdash \ulcorner (\text{PA}i) \urcorner$ is derivable in LKID, and that each instance of the induction schema is provable in LKID using the induction rule (Ind N). The remaining cases follow straightforwardly by the induction hypothesis; the cases ($\exists$ R) and ($\forall$ L) require the use of an auxiliary lemma stating that $Nx_1, \dots, Nx_k \vdash Nt$ is provable for any Σ'_{PA} -term t whose variables are contained in $\{x_1, \dots, x_k\}$.

($\Leftarrow$) By induction on the height of the derivation of $Nx_1, \dots, Nx_k, \ulcorner \Gamma \urcorner \vdash \ulcorner \Delta \urcorner$ in $\text{LKID} + (\text{PA1})\text{--}(\text{PA6})$. We note that $\Gamma \vdash \Delta$ is essentially obtained from this sequent by reading all formulas of the form Nt as $\top$ (easily definable in FOL_{ID}). Most of the rule cases are then straightforward, with applications of (Ind N) being translated as uses of the Peano induction axiom. $\square$

We remark that, by Lemma 3.14 and our completeness result for LKID, PA-derivability corresponds to Henkin validity over $(\Sigma'_{\text{PA}}, \Phi_N)$ in the same way that membership of true arithmetic corresponds to standard validity over $(\Sigma'_{\text{PA}}, \Phi_N)$ (cf. Lemma 3.12). It should also be possible to prove this via a direct argument.

Theorem 3.15. *Eliminability of cut in LKID implies consistency of PA.*

Proof. Suppose PA is inconsistent, i.e., there is a proof of the empty sequent $\vdash$ in PA. By Lemma 3.14, $\vdash$ is then provable in $\text{LKID} + (\text{PA1})\text{--}(\text{PA6})$. Since $(\text{PA1})\text{--}(\text{PA6})$ are closed first-order formulas, it follows that the sequent $(\text{PA1}), \dots, (\text{PA6}) \vdash$ is provable in LKID, and thus cut-free provable by cut-eliminability in LKID (Corollary 3.7). But every rule of LKID, except (Cut), having an inductive predicate in one of its premises also has an inductive predicate in its conclusion. Therefore, since the sequent $(\text{PA1}), \dots, (\text{PA6}) \vdash$ contains no inductive predicates, there are no instances of the rules for inductive predicates occurring anywhere in its cut-free proof. We thus have a cut-free proof of $(\text{PA1}), \dots, (\text{PA6}) \vdash$ in the system LK_e , and so $\vdash$ is derivable in $\text{LK}_e + (\text{PA1})\text{--}(\text{PA6})$, i.e., the axioms $(\text{PA1})\text{--}(\text{PA6})$ are inconsistent. But this system can be proved consistent by elementary means (see e.g. [13]). Hence we have the required contradiction and conclude PA is consistent. $\square$

Corollary 3.16. *The eliminability of cut in LKID (Corollary 3.7) is not provable in PA.*

Proof. The eliminability of cut in LKID implies the consistency of PA by Theorem 3.15, the proof of which is evidently itself formalisable in PA. Therefore, were the eliminability of cut in LKID provable in PA, then the consistency of PA would be provable in PA, contradicting Gödel’s second incompleteness theorem. $\square$

4 Proof of Henkin completeness of LKID (Theorem 3.6)

In this section, we present the proof of completeness of LKID with respect to Henkin models (Theorem 3.6). The proof is an extension of the direct style of completeness proof for Gentzen’s

LK as given in e.g. [9]. Briefly, supposing that $\Gamma \vdash \Delta$ is not cut-free provable in LKID, we use a uniform proof-search procedure to construct a sequence of underivable sequents $\Gamma_i \vdash \Delta_i$, which can together be used to build a syntactic countermodel to the original sequent. The required modifications to the standard argument in our case concern the rules for equality and inductively defined predicates, and also the need to construct a Henkin class over the model.

Definition 4.1 (Schedule). An *LKID-schedule element* for Σ is defined as any of the following, where k_i is the arity of the inductive predicate symbol P_i for each $i \in \{1, \dots, n\}$:

- a formula of the form $\neg F$, $F_1 \wedge F_2$, $F_1 \vee F_2$, or $F_1 \rightarrow F_2$;
- a pair of the form $\langle \forall x F, t \rangle$ or $\langle \exists x F, t \rangle$ where $\forall x F$ and $\exists x F$ are formulas and t is a Σ -term;
- a tuple of the form $\langle P_i \mathbf{t}, \mathbf{z}_1, F_1, \dots, \mathbf{z}_n, F_n \rangle$ where P_i is an inductive predicate symbol, $\mathbf{t}$ is a sequence of k_i terms of Σ and, for each $j \in \{1, \dots, n\}$, $\mathbf{z}_j$ is a sequence of k_j distinct variables and F_j is a formula.

An *LKID-schedule* for Σ is then an enumeration $(E_i)_{i \geq 0}$ of schedule elements of Σ such that every schedule element of Σ appears infinitely often in the enumeration.

Note that an LKID-schedule for Σ exists since we assume our languages countable. Our next definition proceeds assuming a given LKID-schedule $(E_i)_{i \geq 0}$. Also, we henceforth assume a fixed sequent $\Gamma \vdash \Delta$ that is not cut-free provable.

Definition 4.2 (Limit sequent). We define an infinite sequence $(\Gamma_i \vdash \Delta_i)_{i \geq 0}$ of sequents such that each $\Gamma_i \vdash \Delta_i$ is not cut-free provable. We set $\Gamma_0 \vdash \Delta_0 = \Gamma \vdash \Delta$, so this is trivially the case for $i = 0$. Now we assume inductively that we have constructed $(\Gamma_j \vdash \Delta_j)_{0 \leq i \leq j}$, and show how to construct $S = \Gamma_{j+1} \vdash \Delta_{j+1}$.

First note that no formula can be in both Γ_j and Δ_j , otherwise $\Gamma_j \vdash \Delta_j$ would be an axiom instance and thus cut-free provable. We proceed by case distinction on E_j , the j th element in the schedule. Let F be the main formula of E_j (i.e. the one occurring leftmost, if E_j is a tuple). If $F \notin \Gamma_j \cup \Delta_j$, or if F is of the form $P_i \mathbf{t}$ and $F \in \Delta_j$, then we define $S = \Gamma_j \vdash \Delta_j$. Otherwise, we consider the derivation obtained by, if F is non-atomic, applying the sequent rule $(-L)$ or $(-R)$ as appropriate with principal formula F , where $-$ is the main connective of F , or, if F is of the form $P_i \mathbf{t}$ and occurs in Γ_j , applying the rule $(\text{Ind } P_i)$ with F as principal formula. Where applying the rule requires us to perform an instantiation, we use the extra information appearing alongside F in the schedule element E_j . As $\Gamma_j \vdash \Delta_j$ is not cut-free provable, it then follows that one of the premises of this last rule application is not cut-free provable, and we pick S to be any such premise. We show some sample cases: the other cases are similar.

- Case $E_j = F_1 \wedge F_2$. If $F_1 \wedge F_2 \in \Gamma_j$ then by the rule application:

$$\frac{\Gamma_j, F_1, F_2 \vdash \Delta_j}{\Gamma_j, F_1 \wedge F_2 \vdash \Delta_j} (\wedge L)$$

it is clear that $\Gamma_j, F_1, F_2 \vdash \Delta_j$ is not cut-free provable, since otherwise $\Gamma_j \vdash \Delta_j$ is cut-free provable, contradicting the inductive hypothesis. (Note that we have $\Gamma_j = (\Gamma_j, F_1 \wedge F_2)$ since Γ_j is a set and comma is set union.) We thus define $S = \Gamma_j, F_1, F_2 \vdash \Delta_j$. Otherwise, if $F_1 \wedge F_2 \in \Delta_j$ then by the rule application:

$$\frac{\Gamma_j \vdash F_1, \Delta_j \quad \Gamma_j \vdash F_2, \Delta_j}{\Gamma_j \vdash F_1 \wedge F_2, \Delta_j} (\wedge R)$$

it is clear that one of $\Gamma_j \vdash \Delta_j, F_1$ or $\Gamma_j \vdash \Delta_j, F_2$, is not cut-free provable. We define S to be $\Gamma_j \vdash \Delta_j, F_1$ if $\Gamma_j \vdash \Delta_j, F_1$ is not cut-free provable and $\Gamma_j \vdash \Delta_j, F_2$ otherwise.

- Case $E_j = \langle \exists xF, t \rangle$. If $\exists xF \in \Gamma_j$ then by the rule application:

$$\frac{\Gamma_j, F[z/x] \vdash \Delta_j}{\Gamma_j, \exists xF \vdash \Delta_j} (\exists L)$$

where $z \notin FV(\Gamma_j \cup \Delta_j)$, it is clear that $\Gamma_j, F[z/x] \vdash \Delta_j$ cannot be cut-free provable and we thus define $S = \Gamma_j, F[z/x] \vdash \Delta_j$. On the other hand, if $\exists xF \in \Delta_j$ then by the inference:

$$\frac{\Gamma_j \vdash F[t/x], \Delta_j}{\Gamma_j \vdash \exists xF, \Delta_j} (\exists R)$$

it is clear that $\Gamma_j \vdash F[t/x], \Delta_j$ cannot be cut-free provable and we thus define $S = \Gamma_j \vdash F[t/x], \Delta_j$.

- Case $E_j = \langle P_i \mathbf{t}, \mathbf{z}_1, F_1, \dots, \mathbf{z}_n, F_n \rangle$. If $P_i \mathbf{t} \notin \Gamma_j$ we just set $S = \Gamma_j \vdash \Delta_j$. If $P_i \mathbf{t} \in \Gamma_j$ we have the inference:

$$\frac{\text{minor premises } \Gamma_j, F_i \mathbf{t} \vdash \Delta_j}{\Gamma_j, P_i \mathbf{t} \vdash \Delta_j} (\text{Ind } P_i)$$

where the minor premises are obtained by using (some of) the tuples of variables $\mathbf{z}_1, \dots, \mathbf{z}_n$ as the induction variables and (some of) the formulas $F_1, \dots, F_n$ as induction hypotheses in the instance of $(\text{Ind } P_i)$. Again, we pick S to be any of the premises of the instance that is not cut-free provable (clearly there is at least one).

By construction, we have $\Gamma_j \subseteq \Gamma_{j+1}$ and $\Delta_j \subseteq \Delta_{j+1}$ for all $j \geq 0$. Let $\Gamma_\omega = \bigcup_{j \geq 0} \Gamma_j$ and $\Delta_\omega = \bigcup_{j \geq 0} \Delta_j$. Then the *limit sequent* for $\Gamma \vdash \Delta$ is defined to be $\Gamma_\omega \vdash \Delta_\omega$. Strictly speaking $\Gamma_\omega \vdash \Delta_\omega$ need not be a sequent since the sets Γ_ω and Δ_ω may be infinite. When we say that such an infinite “sequent” is cut-free provable, we mean that some finite subsequent of the infinite “sequent” is cut-free provable. Clearly, $\Gamma_\omega \vdash \Delta_\omega$ is not cut-free provable.

Definition 4.3. Define the relation $\sim$ to be the smallest congruence relation on terms of Σ that satisfies: $t_1 \sim t_2$ whenever $(t_1 = t_2) \in \Gamma_\omega$. We write $[t]$ for the equivalence class of t with respect to $\sim$, i.e. $[t] = \{u \mid t \sim u\}$. If $\mathbf{t} = (t_1, \dots, t_k)$ then we shall write $[\mathbf{t}]$ to mean $([t_1], \dots, [t_k])$.

Lemma 4.4. *If $t \sim u$ then, for any formula F , it holds that $\Gamma_\omega \vdash F[t/x]$ is cut-free provable if and only if $\Gamma_\omega \vdash F[u/x]$ is cut-free provable.*

Proof. By rule induction on the conditions defining $t \sim u$. □

Definition 4.5 (Counter-interpretation). Define a first-order structure M_ω for Σ by:

- the domain of M_ω is $\text{Terms}(\Sigma) / \sim =_{\text{def}} \{[t] \mid t \text{ a } \Sigma\text{-term}\}$, the set of $\sim$ -equivalence classes of Σ -terms;
- for any function symbol f in Σ of arity k , $f^{M_\omega}([t_1], \dots, [t_k]) = [f(t_1, \dots, t_k)]$;
- for any ordinary predicate symbol Q in Σ of arity k , Q^{M_ω} is defined by:

$$Q^{M_\omega}([t_1], \dots, [t_k]) \Leftrightarrow \exists u_1, \dots, u_k. t_1 \sim u_1, \dots, t_k \sim u_k \text{ and } Q(u_1, \dots, u_k) \in \Gamma_\omega$$

- for any inductive predicate symbol P_i in Σ of arity k_i , $P_i^{M_\omega}$ is defined by:

$$(P_1^{M_\omega}, \dots, P_n^{M_\omega}) = \text{least}(X_1, \dots, X_n). (\varphi_\Phi(X_1, \dots, X_n) \subseteq (X_1, \dots, X_n) \text{ and } (\forall i \in \{1, \dots, n\}. P_i \mathbf{t} \in \Gamma_\omega \Rightarrow [\mathbf{t}] \in X_i))$$

i.e. $(P_1^{M_\omega}, \dots, P_n^{M_\omega})$ is the least prefixed point of φ_Φ whose i th component contains $[t]$ whenever $P_i t \in \Gamma_\omega$. Note that φ_Φ acts on the structure M_ω ; this is not a circular definition, since the definition of φ_Φ (cf. Defn. 2.3) only requires the interpretation given to the constants, function symbols and ordinary predicates of Σ by M_ω , which we have already defined. To see that the least prefixed point $(P_1^{M_\omega}, \dots, P_n^{M_\omega})$ actually exists, first note that the set $(Terms(\Sigma)/\sim)^n$ is a prefixed point of φ_Φ whose i th component trivially contains $[t]$ whenever $P_i t \in \Gamma_\omega$. Then observe that, given any two such prefixed points of φ_Φ , their intersection is also a prefixed point of φ_Φ , of smaller or equal size, and possessing the same property. Thus $(P_1^{M_\omega}, \dots, P_n^{M_\omega})$ is given by the intersection of all such prefixed points of φ_Φ .

Also, we define an environment ρ_ω for M_ω by $\rho_\omega(x) = [x]$ for all variables x . Then (M_ω, ρ_ω) is called the *counter-interpretation* for $\Gamma_\omega \vdash \Delta_\omega$.

Lemma 4.6. *For any inductive predicate P_i , if $M_\omega \models_{\rho_\omega} P_i t$ then $\Gamma_\omega \vdash P_i t$ is cut-free provable.*

Proof. It can easily be established that $\rho_\omega(t) = [t]$, whence we immediately have $M_\omega \models_{\rho_\omega} P_i t \Leftrightarrow [t] \in P_i^{M_\omega}$. Define an n -tuple of sets $(X_1, \dots, X_n)$ by:

$$X_i = \{[t] \mid \Gamma_\omega \vdash P_i t \text{ cut-free provable}\} \quad (i \in \{1, \dots, n\})$$

It is thus immediate that if $P_i t \in \Gamma_\omega$ then $[t] \in X_i$. As $(P_1^{M_\omega}, \dots, P_n^{M_\omega})$ is the least prefixed point of φ_Φ satisfying this condition (cf. Defn 4.5), we show that $(X_1, \dots, X_n)$ is a prefixed point of φ_Φ . It then follows that if $[t] \in P_i^{M_\omega}$ then $[t] \in X_i$, and so by definition of X_i and Lemma 4.4, $\Gamma_\omega \vdash P_i t$ is cut-free provable as required.

To see that $(X_1, \dots, X_n)$ is indeed a prefixed point of φ_Φ , it suffices to show the inclusion $\varphi_{i,r}(X_1, \dots, X_n) \subseteq X_i$ for an arbitrary production $\Phi_{i,r} \in \Phi$. That is, we must show for $\Phi_{i,r}$ of the form:

$$\frac{Q_1 \mathbf{u}_1(\mathbf{x}) \dots Q_h \mathbf{u}_h(\mathbf{x}) P_{j_1} \mathbf{t}_1(\mathbf{x}) \dots P_{j_m} \mathbf{t}_m(\mathbf{x})}{P_i \mathbf{t}(\mathbf{x})}$$

that we have the inclusion:

$$\{[t(\mathbf{x})] \mid Q_1^{M_\omega}[\mathbf{u}_1(\mathbf{x})], \dots, Q_h^{M_\omega}[\mathbf{u}_h(\mathbf{x})], [t_1(\mathbf{x})] \in X_{j_1}, \dots, [t_m(\mathbf{x})] \in X_{j_m}\} \subseteq X_i$$

By the definition of M_ω and of $(X_1, \dots, X_n)$, and making use of Lemma 4.4, this amounts to showing the following implication:

$$\begin{aligned} \Gamma_\omega \vdash Q_1 \mathbf{u}_1(\mathbf{x}), \dots, \Gamma_\omega \vdash Q_h \mathbf{u}_h(\mathbf{x}), \Gamma_\omega \vdash P_{j_1} \mathbf{t}_1(\mathbf{x}), \dots, \Gamma_\omega \vdash P_{j_m} \mathbf{t}_m(\mathbf{x}) \text{ all cut-free provable} \\ \implies \Gamma_\omega \vdash P_i \mathbf{t}(\mathbf{x}) \text{ cut-free provable} \end{aligned}$$

which follows from the fact that cut-free provability from Γ_ω is closed under the right-introduction rule $(P_i R_r)$. $\square$

Lemma 4.7. *If $F \in \Gamma_\omega$ then $M_\omega \models_{\rho_\omega} F$, and if $F \in \Delta_\omega$ then $M_\omega \not\models_{\rho_\omega} F$.*

Proof. By structural induction on the formula F . All the cases apart from those for equality and inductive predicates follow by the analogous cases in the standard first-order completeness argument (see e.g. [9]). We show the non-standard cases and a representative selection of the standard ones.

Case $F = P_i t$, where P_i is an inductive predicate symbol of Σ . If $P_i t \in \Gamma_\omega$ then $[t] \in P_i^{M_\omega}$ by definition of M_ω , i.e. $M_\omega \models_{\rho_\omega} P_i t$ as required. On the other hand, if $P_i t \in \Delta_\omega$ then we must

have $M_\omega \not\models_{\rho_\omega} P_i \mathbf{t}$, for otherwise $\Gamma_\omega \vdash P_i \mathbf{t}$ would be cut-free provable by Lemma 4.6, and so $\Gamma_\omega \vdash \Delta_\omega$ would be cut-free provable, a contradiction.

Case $F = (t_1 = t_2)$. If $t_1 = t_2 \in \Gamma_\omega$ then we have $t_1 \sim t_2$ and thus $[t_1] = [t_2]$, i.e. $\rho_\omega(t_1) = \rho_\omega(t_2)$ and so $M_\omega \models_{\rho_\omega} t_1 = t_2$.

Now suppose $t_1 = t_2 \in \Delta_\omega$, and suppose for contradiction that $M_\omega \models_{\rho_\omega} t_1 = t_2$, i.e. $\rho_\omega(t_1) = \rho_\omega(t_2)$, whence $[t_1] = [t_2]$ and so $t_1 \sim t_2$. Now, observe that $\Gamma_\omega \vdash t_1 = t_1$ is cut-free provable via an application of (=R). Hence by Lemma 4.4, $\Gamma_\omega \vdash t_1 = t_2$ is also cut-free provable. But since $t_1 = t_2 \in \Delta_\omega$, we would then have a cut-free proof of $\Gamma_\omega \vdash \Delta_\omega$, which is a contradiction. Hence $M_\omega \not\models_{\rho_\omega} t_1 = t_2$.

Case $F = F_1 \wedge F_2$. If $F_1 \wedge F_2 \in \Gamma_\omega$ then by construction of $\Gamma_\omega \vdash \Delta_\omega$, we have $F_1 \in \Gamma_\omega$ and $F_2 \in \Gamma_\omega$. By induction hypothesis $M_\omega \models_{\rho_\omega} F_1$ and $M_\omega \models_{\rho_\omega} F_2$, i.e. $M_\omega \models_{\rho_\omega} F_1 \wedge F_2$ as required. If on the other hand $F_1 \wedge F_2 \in \Delta_\omega$ then by construction we have $F_1 \in \Delta_\omega$ or $F_2 \in \Delta_\omega$. In the former case we have by induction hypothesis $M_\omega \not\models_{\rho_\omega} F_1$ whence it is clear that $M_\omega \not\models_{\rho_\omega} F_1 \wedge F_2$; the other case is similar.

Case $F = \exists x F'$. If $\exists x F' \in \Gamma_\omega$ then by construction of $\Gamma_\omega \vdash \Delta_\omega$, we have $F'[z/x] \in \Gamma_\omega$ for some variable z , whence $M_\omega \models_{\rho_\omega} F'[z/x]$ by induction hypothesis and so $M_\omega \models_{\rho_\omega[x \mapsto \rho_\omega(z)]} F'$, i.e. $M_\omega \models_{\rho_\omega} \exists x F'$.

Now suppose $\exists x F' \in \Delta_\omega$, and observe that by construction of $\Gamma_\omega \vdash \Delta_\omega$ there is then an $i \geq 0$ such that $\exists x F' \in \Gamma_j$ for all $j \geq i$. Now consider an arbitrary term t of Σ and note that the element $(\exists x F', t)$ appears infinitely often on the schedule $(E_i)_{i \geq 0}$ according to which $\Gamma_\omega \vdash \Delta_\omega$ is constructed. So there is a $j \geq i$ such that $E_j = \langle \exists x F', t \rangle$ and thus we have $F'[t/x] \in \Delta_\omega$. As t was chosen arbitrarily, it follows that for every term t , $F'[t/x] \in \Delta_\omega$. So by induction hypothesis $M_\omega \not\models_{\rho_\omega} F'[t/x]$ for every term t . Suppose for contradiction that $M_\omega \models_{\rho_\omega} \exists x F'$. Then for some $t \in \text{Terms}(\Sigma)$, we would have $M_\omega \models_{\rho_\omega[x \mapsto [t]]} F'$, i.e. $M_\omega \models_{\rho_\omega[x \mapsto \rho_\omega(t)]} F'$, and so it follows that $M_\omega \models_{\rho_\omega} F'[t/x]$, which contradicts our induction hypotheses. Hence $M_\omega \not\models_{\rho_\omega} \exists x F'$. $\square$

Definition 4.8 (Henkin counter-class). Define $\mathcal{H}_\omega = \{H_k \mid k \in \mathbb{N}\}$ by:

$$H_k = \{ \{([t_1], \dots, [t_k]) \mid M_\omega \models_{\rho_\omega} F[t_1/x_1, \dots, t_k/x_k]\} \mid \\ F \text{ a formula and } x_1, \dots, x_k \text{ distinct variables} \}$$

$\mathcal{H}_\omega$ is said to be the *Henkin counter-class* for $\Gamma_\omega \vdash \Delta_\omega$.

Lemma 4.9. *The Henkin counter-class for $\Gamma_\omega \vdash \Delta_\omega$ is indeed a Henkin class for M_ω .*

Proof. One easily verifies that the closure conditions defining Henkin classes (cf. Definition 2.6) hold for $\mathcal{H}_\omega$. A similar verification can be found as Proposition 2.3.4 in [4]. $\square$

Lemma 4.10. *$(M_\omega, \mathcal{H}_\omega)$ is a Henkin model for (Σ, Φ) .*

Proof. First, note that $\mathcal{H}_\omega$ is a Henkin class for M_ω by Lemma 4.9. We must prove that $(P_1^{M_\omega}, \dots, P_n^{M_\omega})$ is the least prefixed $\mathcal{H}_\omega$ -point of φ_Φ . It is easily seen to be an $\mathcal{H}_\omega$ -point and is a prefixed point of φ_Φ by definition. To see it is the least such point, it suffices to show that an arbitrary prefixed $\mathcal{H}_\omega$ point $(X_1, \dots, X_n)$ satisfies the condition in the definition of $(P_1^{M_\omega}, \dots, P_n^{M_\omega})$, i.e. that $P_j \mathbf{t} \in \Gamma_\omega$ implies $[\mathbf{t}] \in X_j$ for each $j \in \{1, \dots, n\}$.

To see this, first observe that as $(X_1, \dots, X_n)$ is an $\mathcal{H}_\omega$ -point, for each $i \in \{1, \dots, n\}$ there is a tuple of variables $\mathbf{z}_i$ and a formula F_i such that $X_i = \{[\mathbf{t}] \mid M_\omega \models_{\rho_\omega} F_i[\mathbf{t}/\mathbf{z}_i]\}$. Now assume that $P_j \mathbf{t} \in \Gamma_\omega$; it follows that there is a point in the construction of $\Gamma_\omega \vdash \Delta_\omega$ at which the rule (Ind P_j) is applied with principal formula $P_j \mathbf{t}$, induction variables $\mathbf{z}_1, \dots, \mathbf{z}_n$ and induction

hypotheses $F_1, \dots, F_n$. One of the premises of this rule application is thus a subsequent of $\Gamma_\omega \vdash \Delta_\omega$. If it is the major premise, then we have $F_j \mathbf{t} \in \Gamma_\omega$ and so $M_\omega \models_{\rho_\omega} F_j \mathbf{t}$ by Lemma 4.7, i.e. $[\mathbf{t}] \in X_j$ as required. Otherwise, it is a minor premise which, again by Lemma 4.7, must be false in M_ω . But this contradicts the assumed closure of $(X_1, \dots, X_n)$ under φ_Φ . The full details appear as Lemma 3.3.12 of [4]. $\square$

We now complete the proof of Theorem 3.6. Suppose that $\Gamma \vdash \Delta$ is not cut-free provable in LKID. Letting (M_ω, ρ_ω) be the counter-interpretation and $\mathcal{H}_\omega$ be the Henkin counter-class for the limit sequent $\Gamma_\omega \vdash \Delta_\omega$ constructed from $\Gamma \vdash \Delta$, we have by Lemma 4.10 that $(M_\omega, \mathcal{H}_\omega)$ is a Henkin model for (Σ, Φ) . By Lemma 4.7, every finite subsequent of $\Gamma_\omega \vdash \Delta_\omega$ is false in $(M_\omega, \mathcal{H}_\omega)$, including $\Gamma \vdash \Delta$, so $\Gamma \vdash \Delta$ is not Henkin valid.

5 LKID^ω: a proof system for infinite descent in FOL_{ID}

In this section, we formulate an infinitary proof system, LKID^ω, formalising a version of proof by infinite descent in FOL_{ID}. As in the case of LKID, we prove soundness and cut-free completeness of LKID^ω, and thereby infer the eliminability of cut in LKID^ω. However, for LKID^ω, the soundness and completeness results are relative to the more natural class of standard models of FOL_{ID}, rather than the wider class of Henkin models.

The proof rules of the system LKID^ω are the rules of LKID described in §3, except that for each inductive predicate P_i of Σ , the induction rule (Ind P_i) of LKID is replaced by the *case-split rule*:

$$\frac{\text{case distinctions}}{\Gamma, P_i \mathbf{u} \vdash \Delta} \text{ (Case } P_i \text{)}$$

where for each production having predicate P_i in its conclusion, say:

$$\frac{Q_1 \mathbf{u}_1(\mathbf{x}) \dots Q_h \mathbf{u}_h(\mathbf{x}) P_{j_1} \mathbf{t}_1(\mathbf{x}) \dots P_{j_m} \mathbf{t}_m(\mathbf{x})}{P_i \mathbf{t}(\mathbf{x})}$$

there is a corresponding case distinction:

$$\Gamma, \mathbf{u} = \mathbf{t}(\mathbf{y}), Q_1 \mathbf{u}_1(\mathbf{y}), \dots, Q_h \mathbf{u}_h(\mathbf{y}), P_{j_1} \mathbf{t}_1(\mathbf{y}), \dots, P_{j_m} \mathbf{t}_m(\mathbf{y}) \vdash \Delta$$

where $\mathbf{y}$ is a vector of distinct variables of the same length as $\mathbf{x}$, and $y \notin FV(\Gamma \cup \Delta \cup \{P_i \mathbf{u}\})$ for all $y \in \mathbf{y}$. The formulas $P_{j_1} \mathbf{t}_1(\mathbf{y}), \dots, P_{j_m} \mathbf{t}_m(\mathbf{y})$ occurring in a case distinction are said to be *case-descendants* of the principal formula $P_i \mathbf{u}$.

Example 5.1. The case-split rule for N from Example 2.2 is:

$$\frac{\Gamma, t = 0 \vdash \Delta \quad \Gamma, t = sx, Nx \vdash \Delta}{\Gamma, Nt \vdash \Delta} x \notin FV(\Gamma \cup \Delta \cup \{Nt\}) \text{ (Case } N \text{)}$$

The formula Nx occurring in the right-hand premise is the only case-descendant of the formula Nt occurring in the conclusion.

Example 5.2. The rule for E from Example 2.2 is:

$$\frac{\Gamma, t = 0 \vdash \Delta \quad \Gamma, t = sx, Ox \vdash \Delta}{\Gamma, Et \vdash \Delta} x \notin FV(\Gamma \cup \Delta \cup \{Et\}) \text{ (Case } E \text{)}$$

The formula Ox occurring in the right-hand premise is the only case-descendant of the formula Et occurring in the conclusion.

Our proof system LKID^ω will be based upon infinite derivation trees. For convenience, we distinguish between “leaves” and “buds” in derivation trees. By a *leaf* we mean an axiom, i.e. the conclusion of a 0-premise inference rule. By a *bud* we mean any sequent occurrence in the tree that is not the conclusion of a proof rule.

Definition 5.3 (LKID^ω pre-proof). An LKID^ω *pre-proof* of a sequent $\Gamma \vdash \Delta$ is a (possibly infinite) derivation tree $\mathcal{D}$, constructed according to the proof rules of LKID^ω , such that $\Gamma \vdash \Delta$ is the root of $\mathcal{D}$ and $\mathcal{D}$ has no buds.

We remark that LKID^ω pre-proofs are not sound in general. For example, there is a pre-proof of the invalid sequent $A \wedge B \vdash$ consisting of infinitely many applications of $(\wedge\text{L})$, and there are infinitely many pre-proofs of any invalid sequent consisting of nothing but spurious applications of (Cut) . We therefore impose a global condition on pre-proofs which ensures their soundness by requiring that all infinite paths correspond to well-founded arguments.

As is standard, we define a (finite or infinite) *path* in a derivation tree to be a sequence $(S_i)_{0 \leq i < \alpha}$, for some $\alpha \in \mathbb{N} \cup \{\infty\}$, of sequent occurrences in the tree such that S_{i+1} is a child of S_i for all $i + 1 < \alpha$.

Definition 5.4 (Trace). Let $\mathcal{D}$ be an LKID^ω pre-proof and let $(\Gamma_i \vdash \Delta_i)_{i \geq 0}$ be a path in $\mathcal{D}$. A *trace following* $(\Gamma_i \vdash \Delta_i)_{i \geq 0}$ is a sequence $(\tau_i)_{i \geq 0}$ such that, for all i :

- $\tau_i = P_{j_i} \mathbf{t}_i \in \Gamma_i$, where $j_i \in \{1, \dots, n\}$;
- if $\Gamma_i \vdash \Delta_i$ is the conclusion of (Subst) then $\tau_i = \tau_{i+1}[\theta]$, where θ is the substitution associated with the rule instance;
- if $\Gamma_i \vdash \Delta_i$ is the conclusion of $(=\text{L})$ with principal formula $t = u$ then there is a formula F and variables x, y such that $\tau_i = F[t/x, u/y]$ and $\tau_{i+1} = F[u/x, t/y]$;
- if $\Gamma_i \vdash \Delta_i$ is the conclusion of a case-split rule then either (a) $\tau_{i+1} = \tau_i$, or (b) τ_i is the principal formula of the rule instance and τ_{i+1} is a case-descendant of τ_i . In the latter case, i is said to be a *progress point* of the trace;
- if $\Gamma_i \vdash \Delta_i$ is the conclusion of any other rule then $\tau_{i+1} = \tau_i$.

An *infinitely progressing trace* is a trace having infinitely many progress points.

Definition 5.5 (LKID^ω proof). An LKID^ω pre-proof $\mathcal{D}$ is an LKID^ω *proof* if it satisfies the following *global trace condition*: for every infinite path $(\Gamma_i \vdash \Delta_i)_{i \geq 0}$ in $\mathcal{D}$, there is an infinitely progressing trace following some tail of the path, $(\Gamma_i \vdash \Delta_i)_{i \geq k}$, for some $k \geq 0$.

Example 5.6. Let N, E and O be the predicates given in Example 2.2. Figure 2 gives the initial part of an LKID^ω pre-proof of the sequent $Nx_0 \vdash Ex_0, Ox_0$, with the underlined formulas showing a trace following the right hand branch of the pre-proof. This trace progresses because its second element Nx_1 is a case-descendant of its first element Nx_0 . One can easily see that by continuing the expansion of this derivation, we obtain an infinite tree with exactly one infinite branch. Furthermore, one can similarly continue the trace along this branch to obtain an infinitely progressing trace: $(Nx_0, Nx_1, \dots, Nx_1, Nx_2, \dots)$. The pre-proof thereby obtained is thus indeed an LKID^ω proof.

The following lemma is a consequence of the local soundness of the proof rules and the fact that a trace “tracks” case-descendants of an inductive predicate along a path in a derivation tree.

$$\begin{array}{c}
\text{(etc.)} \\
\vdots \\
\frac{}{Nx_1 \vdash Ex_1, Ox_1} \text{(Case } N) \\
\frac{}{Nx_1 \vdash Ox_1, Osx_1} (OR_1) \\
\frac{}{Nx_1 \vdash Esx_1, Osx_1} (ER_2) \\
\frac{}{\vdash E0, O0} (ER_1) \\
\frac{}{x_0 = 0 \vdash Ex_0, Ox_0} (=L) \quad \frac{}{x_0 = sx_1, \underline{Nx_1} \vdash Ex_0, Ox_0} (=L) \\
\hline
\underline{Nx_0} \vdash Ex_0, Ox_0 \quad \text{(Case } N)
\end{array}$$

Figure 2: Portion of an $LKID^\omega$ proof, with a progressing trace denoted by the underlined formulas.

Lemma 5.7. *Let $\mathcal{D}$ be an $LKID^\omega$ pre-proof of $\Gamma_0 \vdash \Delta_0$, and let M be a standard model such that $\Gamma_0 \vdash \Delta_0$ is false in M under the environment ρ_0 . Then there is an infinite path $(\Gamma_i \vdash \Delta_i)_{i \geq 0}$ in $\mathcal{D}$ and an infinite sequence $(\rho_i)_{i \geq 0}$ of environments such that:*

1. *for all $i \geq 0$, $\Gamma_i \vdash \Delta_i$ is false in M under ρ_i ;*
2. *if there is a trace $(\tau_i = P_{j_i} \mathbf{t}_i)_{i \geq n}$ following some tail $(\Gamma_i \vdash \Delta_i)_{i \geq n}$ of $(\Gamma_i \vdash \Delta_i)_{i \geq 0}$, then the sequence $(\alpha_i)_{i \geq n}$ of ordinals defined by:*

$$\alpha_i = \text{least } \alpha \text{ such that } \rho_i(\mathbf{t}_i) \in P_{j_i}^\alpha$$

is non-increasing, i.e. $\alpha_{j+1} \leq \alpha_j$ for all $j \geq n$. Furthermore, if j is a progress point of $(\tau_i)_{i \geq n}$ then we have $\alpha_{j+1} < \alpha_j$.

Proof. We write $\Gamma \not\models_\rho \Delta$ to mean that the sequent $\Gamma \vdash \Delta$ is false in the model M under the environment ρ .

First note that the ordinal sequence $(\alpha_i)_{i \geq n}$ defined in property 2 of the lemma is well-defined, for, by the definition of trace (Definition 5.4), we have $\tau_i = P_{j_i} \mathbf{t}_i \in \Gamma_i$ for each $i \geq n$, and since $\Gamma_i \not\models_{\rho_i} \Delta_i$ for all i by property 1 of the lemma we must have $M \models_{\rho_i} P_{j_i} \mathbf{t}_i$, i.e. $\rho_i(\mathbf{t}_i) \in \bigcup_\alpha P_{j_i}^\alpha$, for each $i \geq n$. Now $\rho_i(\mathbf{t}_i) \in \bigcup_\alpha P_{j_i}^\alpha$ iff $\rho_i(\mathbf{t}_i) \in P_{j_i}^\alpha$ for some ordinal α , and there is a least such α by the well-ordered property of the ordinals, so α_i is defined for each $i \geq n$.

The two properties required by the lemma are trivially true of the 1-element sequences $(\Gamma_0 \vdash \Delta_0)$ and (ρ_0) . We assume we have sequences $(\Gamma_i \vdash \Delta_i)_{0 \leq i \leq k}$ and $(\rho_i)_{0 \leq i \leq k}$ satisfying the two properties of the lemma and inductively show how to construct $\Gamma_{k+1} \vdash \Delta_{k+1}$ and ρ_{k+1} . We always choose $\Gamma_{k+1} \vdash \Delta_{k+1}$ to be a premise of the rule instance in $\mathcal{D}$ of which $\Gamma_k \vdash \Delta_k$ is the conclusion, so that $(\Gamma_i \vdash \Delta_i)_{i \geq 0}$ is an infinite path in $\mathcal{D}$ as required. To establish that property 2 holds of the constructed sequence, it suffices to assume the existence of an arbitrary trace (τ_k, τ_{k+1}) following the edge $(\Gamma_k \vdash \Delta_k, \Gamma_{k+1} \vdash \Delta_{k+1})$, and show that $\alpha_{k+1} \leq \alpha_k$, with the inequality holding strictly if k is a progress point of the trace. It is clear that this construction can be iterated infinitely often, thus yielding the required infinite sequences.

We note that since $\mathcal{D}$ is an $LKID^\omega$ derivation tree, the sequent $\Gamma_k \vdash \Delta_k$ is the conclusion of an instance of some $LKID^\omega$ proof rule, which clearly cannot be a rule with no premises, as the conclusion of every such rule is easily seen to be a valid sequent but $\Gamma_k \not\models_{\rho_k} \Delta_k$ by the induction hypothesis. We therefore distinguish a case for each of the remaining proof rules. In all cases the fact that $\Gamma_{k+1} \not\models_{\rho_{k+1}} \Delta_{k+1}$ follows immediately from the local soundness of the proof rule in question. Furthermore, ρ_{k+1} can always be constructed in a manner consistent with the

requirement for the second property. We only examine the main interesting case here, which occurs when $\Gamma_k \vdash \Delta_k$ is the conclusion $\Gamma, P_i \mathbf{u} \vdash \Delta$ of an application of rule (Case P_i). (The remaining cases are treated in Lemma 4.2.4 of [4].) As $\Gamma, P_i \mathbf{u} \not\models_{\rho_k} \Delta$ by induction hypothesis, we have $M \models_{\rho_k} P_i \mathbf{u}$, i.e. $\rho_k(\mathbf{u}) \in \bigcup_{\alpha} P_i^{\alpha}$. Let α' be the least ordinal α such that $\rho_k(\mathbf{u}) \in P_i^{\alpha} = \pi_i^n(\varphi_{\Phi}^{\alpha})$. By Definition 2.4 we thus have $\rho_k(\mathbf{u}) \in \pi_i^n(\bigcup_{\beta < \alpha'} \varphi_{\Phi}(\varphi_{\Phi}^{\beta}))$. By construction of φ_{Φ} (cf. Definition 2.3), there is then a $\beta < \alpha'$ and a production $\Phi_{i,r} \in \Phi$ such that $\rho_k(\mathbf{u}) \in \varphi_{i,r}(\varphi_{\Phi}^{\beta})$. Now $\Phi_{i,r}$ is a production with P_i in its conclusion,

$$\frac{Q_1 \mathbf{u}_1(\mathbf{x}) \dots Q_h \mathbf{u}_h(\mathbf{x}) P_{j_1} \mathbf{t}_1(\mathbf{x}) \dots P_{j_m} \mathbf{t}_m(\mathbf{x})}{P_i \mathbf{t}(\mathbf{x})}$$

Letting l be the length of the vector $\mathbf{x}$ appearing in the rule, by definition of $\varphi_{i,r}$ we have:

$$\rho_k(\mathbf{u}) \in \{\mathbf{t}^M(\mathbf{d}) \mid Q_1^M \mathbf{u}_1^M(\mathbf{d}), \dots, Q_h^M \mathbf{u}_h^M(\mathbf{d}), \\ \mathbf{t}_1^M(\mathbf{d}) \in \pi_{j_1}^n(\varphi_{\Phi}^{\beta}), \dots, \mathbf{t}_m^M(\mathbf{d}) \in \pi_{j_m}^n(\varphi_{\Phi}^{\beta}), \mathbf{d} \in D^l\}$$

$$\text{i.e. } \exists \mathbf{d} \in D^l. \rho_k(\mathbf{u}) = \mathbf{t}^M(\mathbf{d}) \text{ and } Q_1^M \mathbf{u}_1^M(\mathbf{d}), \dots, Q_h^M \mathbf{u}_h^M(\mathbf{d}), \\ \mathbf{t}_1^M(\mathbf{d}) \in P_{j_1}^{\beta}, \dots, \mathbf{t}_m^M(\mathbf{d}) \in P_{j_m}^{\beta} (*)$$

Now define $\Gamma_{k+1} \vdash \Delta_{k+1}$ to be the (case distinction) premise corresponding to $\Phi_{i,r}$:

$$\Gamma, \mathbf{u} = \mathbf{t}(\mathbf{y}), Q_1 \mathbf{u}_1(\mathbf{y}), \dots, Q_h \mathbf{u}_h(\mathbf{y}), P_{j_1} \mathbf{t}_1(\mathbf{y}), \dots, P_{j_m} \mathbf{t}_m(\mathbf{y}) \vdash \Delta$$

where $y \notin FV(\Gamma \cup \Delta \cup \{P_i \mathbf{u}\})$ for all $y \in \mathbf{y}$, and define $\rho_{k+1} = \rho_k[\mathbf{y} \mapsto \mathbf{d}]$.

For property 1, we need to show $\Gamma_{k+1} \not\models_{\rho_{k+1}} \Delta_{k+1}$. It is clear that we have $M \models_{\rho_{k+1}} J$ for all $J \in \Gamma$ and $M \not\models_{\rho_{k+1}} K$ for all $K \in \Delta$ by the induction hypothesis, since ρ_{k+1} agrees with ρ_k on all variables free in $\Gamma \cup \Delta$. Also, we have $\rho_{k+1}(\mathbf{u}) = \rho_k(\mathbf{u})$ since $y \notin FV(P_i \mathbf{u})$ for all $y \in \mathbf{y}$. Now $\rho_{k+1}(\mathbf{t}(\mathbf{y})) = \mathbf{t}^M(\rho_{k+1}(\mathbf{y})) = \mathbf{t}^M(\mathbf{d}) = \rho_k(\mathbf{u})$ by the definition of ρ_{k+1} and the statement (*) above, so we have $\rho_{k+1}(\mathbf{u}) = \rho_{k+1}(\mathbf{t}(\mathbf{y}))$, i.e. $M \models_{\rho_{k+1}} \mathbf{u} = \mathbf{t}(\mathbf{y})$ as required. We then just need to show each of $M \models_{\rho_{k+1}} Q_1 \mathbf{u}_1(\mathbf{y}), \dots, M \models_{\rho_{k+1}} Q_h \mathbf{u}_h(\mathbf{y}), M \models_{\rho_{k+1}} P_{j_1} \mathbf{t}_1(\mathbf{y}), \dots, M \models_{\rho_{k+1}} P_{j_m} \mathbf{t}_m(\mathbf{y})$, which is clear from the statement (*) above together with the definition of ρ_{k+1} .

For property 2, there are two possibilities to consider:

- k is not a progress point of the trace (τ_k, τ_{k+1}) and so, by the definition of trace, we have $\tau_{k+1} = \tau_k$. Now, since $\tau_k = P_j \mathbf{t}$ (say) is a formula occurring in $\Gamma_k = \Gamma \cup \{P_i \mathbf{u}\}$, and ρ_{k+1} agrees with ρ_k on variables free in $\Gamma \cup \{P_i \mathbf{u}\}$, we have $\rho_{k+1}(\mathbf{t}) = \rho_k(\mathbf{t})$ and so $\rho_k(\mathbf{t}) \in P_j^{\alpha}$ iff $\rho_{k+1}(\mathbf{t}) \in P_j^{\alpha}$, i.e. $\alpha_{k+1} = \alpha_k$ and we are done.
- k is a progress point of the trace (τ_k, τ_{k+1}) . In that case, τ_k is the principal formula $P_i \mathbf{u}$ of the rule instance and τ_{k+1} is a case-descendant of $P_i \mathbf{u}$, i.e. τ_{k+1} is one of the formulas $P_{j_1} \mathbf{t}_1(\mathbf{y}), \dots, P_{j_m} \mathbf{t}_m(\mathbf{y})$. Now we have from (*) above that there is an ordinal $\beta < \alpha'$ such that $\rho_{k+1}(\mathbf{t}_1(\mathbf{y})) \in P_{j_1}^{\beta}, \dots, \rho_{k+1}(\mathbf{t}_m(\mathbf{y})) \in P_{j_m}^{\beta}$, where $\alpha' = \alpha_k$ is the least ordinal α satisfying $\rho_k(\mathbf{t}) \in P_i^{\alpha}$. We thus have $\alpha_{k+1} < \alpha_k$ as required.

□

Proposition 5.8 (Soundness). *If there is an LKID^ω proof of $\Gamma \vdash \Delta$ then $\Gamma \vdash \Delta$ is valid with respect to standard models.*

Proof. Let $\mathcal{D}$ be an LKID^ω proof of $\Gamma \vdash \Delta$. If $\Gamma \vdash \Delta$ is not valid, i.e. false in some standard model M under some environment ρ_0 , then we can apply Lemma 5.7 to construct infinite sequences $(\Gamma_i \vdash \Delta_i)_{i \geq 0}$ and $(\rho_i)_{i \geq 0}$ satisfying properties 1 and 2 of the lemma. As $\mathcal{D}$ is a proof, and

$(\Gamma_i \vdash \Delta_i)_{i \geq 0}$ is an infinite path in $\mathcal{D}$, there is an infinitely progressing trace following some tail of the path by the global trace condition (Definition 5.5). Thus by the second property of the lemma we can construct an infinite descending chain of ordinals, which is a contradiction. $\square$

Note that our use of approximants in Lemma 5.7 means that our soundness argument only works for standard models. In fact, our main result about LKID^ω is that it is complete with respect to standard models. Thus LKID^ω cannot be sound with respect to Henkin models, for then by standard completeness of LKID^ω and Henkin completeness of LKID we would have the standard completeness of LKID , which contradicts Theorem 3.13. Our completeness result for LKID^ω is very slightly sharpened to recursive LKID^ω -provability. A derivation tree is *recursive* if it is decidable whether a finite sequence of numbers corresponds to a path up the tree from the root (each number indicating the choice of rule premise determining the path) and there is a recursive function mapping each finite path in the tree to a pair consisting of the sequent at the end node of the path and the rule applied with that sequent as conclusion.

Theorem 5.9 (Cut-free completeness of LKID^ω). *If $\Gamma \vdash \Delta$ is valid with respect to standard models, then it has a recursive cut-free proof in LKID^ω .*

We give the proof of this theorem in §6.

Corollary 5.10. *LKID^ω proves strictly more sequents than LKID .*

Proof. Any LKID -provable sequent is Henkin valid by soundness (Proposition 3.5), therefore valid since every standard model is a Henkin model (see Section 2) and hence LKID^ω -provable by Theorem 5.9. However, LKID^ω is complete with respect to standard validity by Theorem 5.9 whereas LKID is necessarily incomplete in this sense by Theorem 3.13. $\square$

Although the completeness theorem shows that every valid sequent has a proof given by a recursive derivation tree, the set of valid sequents relative to standard models is non-arithmetic, since one can encode true arithmetic by Lemma 3.12. Thus there is no way of effectively enumerating any complete subclass of recursive proofs. Hence LKID^ω is, unsurprisingly, not suitable for formal reasoning.

The closest analogue of Theorem 5.9 we are aware of in the literature appears in [27]. There, certain *refutations* are defined, which can be seen as providing an analogous proof system to LKID^ω for Kozen’s propositional μ -calculus [19]. Indeed, refutations are formulated using a trace-based proof condition very similar to Definition 5.5. (Other similar conditions appear in [10, 21, 34, 36, 42].) One of the main results of [27] is a completeness theorem for refutations. Nevertheless, the situations are quite different. In particular, the propositional μ -calculus is decidable, whereas (standard) validity in FOL_{ID} is non-arithmetic. The relationship between our system and others in the literature is discussed in more detail in Section 8.

Corollary 5.11 (Cut-eliminability for LKID^ω). *If $\Gamma \vdash \Delta$ is provable in LKID^ω then it is cut-free provable.*

Proof. If $\Gamma \vdash \Delta$ is provable in LKID^ω , it is valid with respect to standard models by soundness (Proposition 5.8), and hence cut-free provable in LKID^ω by Theorem 5.9. $\square$

Note that, unlike in LKID , cut-free proofs in LKID^ω are quite constrained: every formula appearing in a cut-free LKID^ω proof is either a subformula of a formula appearing in the root sequent, or related to such a formula by a finite number of definitional unfoldings.

6 Proof of standard completeness of LKID^ω (Theorem 5.9)

In this section, we present our proof of (recursive) cut-free completeness for LKID^ω with respect to standard models. As is the case in our proof of completeness for LKID (Theorem 3.6), our proof extends the direct style of completeness argument for Gentzen's LK (cf. [9]). However, for LKID^ω , the need to consider infinite proofs and the global trace condition imposed upon them creates some significant complications.

Given an arbitrary sequent $\Gamma \vdash \Delta$, we construct a recursive, possibly infinite LKID^ω derivation tree corresponding to an exhaustive search for a cut-free proof of $\Gamma \vdash \Delta$. If this tree is not an LKID^ω proof, then either it is not even a pre-proof, i.e. it contains a bud node, or it is a pre-proof but contains an infinite branch that fails to satisfy the global trace condition. We use the bud node or infinite branch as appropriate to construct a *limit sequent*, of which $\Gamma \vdash \Delta$ is a subsequent. This limit sequent is in turn used to define a standard model which falsifies the limit sequent, and thus also $\Gamma \vdash \Delta$. In the case where the tree is a pre-proof but not a proof, the fact that there is no infinitely progressing trace along the infinite branch is used at two points. First, it is needed to show that no sequent on the branch is cut-free provable. Second, it is used to show that the countermodel invalidates the induced limit sequent.

Definition 6.1 (Schedule). An LKID^ω -schedule element for Σ is defined as any of the following:

- a formula of the form $\neg F$, $F_1 \wedge F_2$, $F_1 \vee F_2$, or $F_1 \rightarrow F_2$;
- a pair of the form $\langle \forall x F, t \rangle$ or $\langle \exists x F, t \rangle$ where $\forall x F$ and $\exists x F$ are formulas and t is a Σ -term;
- a pair of the form $\langle P_i \mathbf{t}, r \rangle$ where $P_i \mathbf{t}$ is an atomic formula, P_i is an inductive predicate symbol, and $r \in \mathbb{N}$ satisfies $\Phi_{i,r} \in \Phi$;
- a tuple of the form $\langle t = u, x, y, \Gamma, \Delta \rangle$ where t and u are Σ -terms, x and y are variables, and Γ and Δ are finite sets of formulas.

An LKID^ω -schedule for Σ is then a recursive enumeration $(E_i)_{i \geq 0}$ of schedule elements of Σ such that every schedule element of Σ appears infinitely often in the enumeration.

Henceforth, we assume a fixed LKID^ω -schedule $(E_i)_{i \geq 0}$ and sequent $\Gamma \vdash \Delta$.

Definition 6.2 (Search tree). We define an infinite sequence of $(T_i)_{i \geq 0}$ of derivation trees such that T_0 is the single-node tree $\Gamma \vdash \Delta$ and T_i is a subtree of T_{i+1} with root $\Gamma \vdash \Delta$ for all $i \geq 0$. We inductively assume we have constructed T_j and show how to construct T_{j+1} . In general T_{j+1} will be obtained by replacing certain bud nodes of T_j with derivation trees, whence it is clear that T_{j+1} is a derivation tree of the required form.

Firstly, we replace any bud of T_j that is an instance of the conclusion of an axiom (0-premise) rule with the derivation consisting of a single instance of that axiom. Let F be the formula component of E_j , the j th element in the schedule for Σ . We replace any bud of T_j that contains F with the derivation obtained by applying the sequent rule ($-L$) or ($-R$) as appropriate with principal formula F , performing any required instantiations using the extra information in E_j . The rule applications when E_j is of the form $\neg F$, $F_1 \wedge F_2$, $F_1 \vee F_2$, $F_1 \rightarrow F_2$, $\langle \forall x F, t \rangle$ or $\langle \exists x F, t \rangle$ are identical to the corresponding applications in Definition 4.2 in the LKID completeness proof. We show the cases when E_j is not of one of these forms.

- Case $E_j = \langle P_i \mathbf{u}, r \rangle$. Then T_{j+1} is obtained by first replacing every bud $\Gamma' \vdash \Delta'$ in T_j that satisfies $P_i \mathbf{u} \in \Gamma'$ with the derivation:

$$\frac{\text{case distinctions}}{\Gamma' \vdash \Delta'} (\text{Case } P_i)$$

and then, assuming $\Phi_{i,r}$ is of the generic form

$$\frac{Q_1 \mathbf{u}_1(\mathbf{x}) \dots Q_h \mathbf{u}_h(\mathbf{x}) P_{j_1} \mathbf{t}_1(\mathbf{x}) \dots P_{j_m} \mathbf{t}_m(\mathbf{x})}{P_i \mathbf{t}(\mathbf{x})}$$

and only if we have $\mathbf{u} = \mathbf{t}(\mathbf{u}')$ for some $\mathbf{u}'$, replacing every bud $\Gamma' \vdash \Delta'$ of the resulting tree that satisfies $P_i \mathbf{u} \in \Delta'$ with the derivation:

$$\frac{\begin{array}{c} \Gamma' \vdash Q_1 \mathbf{u}_1(\mathbf{u}'), \Delta' \quad \dots \quad \Gamma' \vdash Q_h \mathbf{u}_h(\mathbf{u}'), \Delta' \\ \Gamma' \vdash P_{j_1} \mathbf{t}_1(\mathbf{u}'), \Delta' \quad \dots \quad \Gamma' \vdash P_{j_m} \mathbf{t}_m(\mathbf{u}'), \Delta' \end{array}}{\Gamma' \vdash \Delta'} (P_i R_r)$$

- Case $E_j = \langle t = u, x, y, \Gamma, \Delta \rangle$. Let $\Gamma' \vdash \Delta'$ be a bud node such that $t = u \in \Gamma'$, $\Gamma' \subseteq \Gamma[t/x, u/y] \cup \{t = u\}$ and $\Delta' \subseteq \Delta[t/x, u/y]$. So $\Gamma' = \Gamma''[t/x, u/y] \cup \{t = u\}$ for some $\Gamma'' \subseteq \Gamma$ and $\Delta' = \Delta''[t/x, u/y]$ for some $\Delta'' \subseteq \Delta$. Then we replace the bud node $\Gamma' \vdash \Delta'$ by the derivation tree:

$$\frac{\Gamma''[u/x, t/y], \Gamma' \vdash \Delta', \Delta''[u/x, t/y]}{\Gamma' \vdash \Delta'} (=L)$$

Note that the above construction is performed in such way that ensures that each sequent in the tree is a subsequent of all its premises.

The *search tree* for $\Gamma \vdash \Delta$ is then defined to be T_ω , the infinite tree obtained by considering the limit as $i \rightarrow \infty$ of the sequence of (finite) derivation trees $(T_i)_{i \geq 0}$. By construction, the search tree is recursive and cut-free.

Henceforth in the proof, we assume that the search tree T_ω is not an LKID^ω proof. If T_ω is not even a pre-proof, then it contains some bud to which no schedule element applies (e.g., a sequent containing only atomic formulas built from ordinary predicates), for which we write $\Gamma_\omega \vdash \Delta_\omega$. Otherwise, T_ω is a pre-proof but not a proof (and hence infinite by soundness of LKID^ω). In this case, the global trace condition fails, so there exists an infinite path $\pi = (\Gamma_i \vdash \Delta_i)_{i \geq 0}$ in T_ω such that there is no infinitely progressing trace following any tail of π . We call this π the *untraceable branch* of T_ω , and define $\Gamma_\omega = \bigcup_{i \geq 0} \Gamma_i$ and $\Delta_\omega = \bigcup_{i \geq 0} \Delta_i$, (noting that we have $\Gamma_i \subseteq \Gamma_{i+1}$ and $\Delta_i \subseteq \Delta_{i+1}$ by construction of T_ω). In either case, we call $\Gamma_\omega \vdash \Delta_\omega$ the *limit sequent*. As in the limit sequent construction in LKID completeness proof (Definition 4.2), the limit sequent might technically fail to be a sequent since either Γ_ω or Δ_ω could be infinite. When we say that such an infinite “sequent” is cut-free provable, we mean that some finite subsequent is cut-free provable.

Lemma 6.3. *The sequent $\Gamma_\omega \vdash \Delta_\omega$ is not cut-free provable.*

Proof. The case that $\Gamma_\omega \vdash \Delta_\omega$ is a bud node is easy (if it were cut-free provable some schedule element would apply to it contradicting it being a bud node). So we assume that T_ω is a pre-proof but not a proof and let $\pi = (\Gamma_i \vdash \Delta_i)_{i \geq 0}$ be the untraceable branch. It suffices to show that no $\Gamma_i \vdash \Delta_i$ has a cut-free proof. So, for contradiction, we assume that T is a cut-free proof of $\Gamma_i \vdash \Delta_i$.

Let $\Gamma' \vdash \Delta'$ be any node in T , let (R) be the rule applied in T with principal formula F (say) and conclusion $\Gamma' \vdash \Delta'$, and suppose $\Gamma' \subseteq \Gamma_j$ and $\Delta' \subseteq \Delta_j$ for some $j \geq i$. As F appears infinitely often on the schedule according to which T_ω is constructed, it follows that there is a $k \geq j$ such that F is the principal formula of an instance of (R) in T_ω with conclusion $\Gamma_k \vdash \Delta_k$. Since the untraceable branch is infinite, it follows that (R) is not an axiom. Therefore, for

Figure 3: Part of the proof of Lemma 6.3. π is the untraceable branch of the search tree T_ω , and T is the assumed proof of some sequent $\Gamma_i \vdash \Delta_i$ on π .

some premise $\Gamma'' \vdash \Delta''$ of the considered instance of rule (R) in T , we have $\Gamma'' \subseteq \Gamma_{k+1}$ and $\Delta'' \subseteq \Delta_{k+1}$. This situation is illustrated in Figure 3.

Since $\Gamma_i \subseteq \Gamma_{i+1}$ and $\Delta_i \subseteq \Delta_{i+1}$ for all $i \geq 0$, it follows that if (τ, τ') is a (progressing) trace following the edge $(\Gamma' \vdash \Delta', \Gamma'' \vdash \Delta'')$ in T , then $(\tau, \dots, \tau, \tau')$ is a (progressing) trace following the subpath $(\Gamma_j \vdash \Delta_j, \dots, \Gamma_k \vdash \Delta_k, \Gamma_{k+1} \vdash \Delta_{k+1})$ of π .

Now since the root of T is $\Gamma_i \vdash \Delta_i$ and trivially $\Gamma_i \subseteq \Gamma_i$ and $\Delta_i \subseteq \Delta_i$, we can repeat the argument in the preceding two paragraphs infinitely often to obtain a path $\pi' = (\Gamma'_j \vdash \Delta'_j)_{j \geq 0}$ in T and a sequence $k_0 < k_1 < k_2 < \dots$ of natural numbers, where $k_0 = i$, such that, for all $n \geq 0$, if (τ, τ') is a (progressing) trace following the edge $(\Gamma'_n \vdash \Delta'_n, \Gamma'_{n+1} \vdash \Delta'_{n+1})$ in T , then $(\tau, \dots, \tau, \tau')$ is a (progressing) trace following the subpath $(\Gamma_{k_n} \vdash \Delta_{k_n}, \dots, \Gamma_{k_{n+1}} \vdash \Delta_{k_{n+1}})$ of π in T_ω .

Since T is a proof, there is an infinitely progressing trace following some tail of the constructed path π' in T . By piecing together the induced trace segments in T_ω defined above, it follows that there then is an infinitely progressing trace following some tail of the untraceable path π in T_ω . But this contradicts the defining property of π . So there cannot exist a cut-free LKID^ω proof of $\Gamma_i \vdash \Delta_i$. $\square$

We now define the equivalence relation $\sim$ and the counter-interpretation (M_ω, ρ_ω) for $\Gamma_\omega \vdash \Delta_\omega$ exactly as we did in the LKID completeness proof (cf. Definitions 4.3 and 4.5, respectively). We remind the reader of the interpretation of the inductive predicates $P_1, \dots, P_n$ in M_ω , since this is crucial:

$$(P_1^{M_\omega}, \dots, P_n^{M_\omega}) = \text{least}(X_1, \dots, X_n). (\varphi_\Phi(X_1, \dots, X_n) \subseteq (X_1, \dots, X_n) \text{ and } (\forall i \in \{1, \dots, n\}. P_i \mathbf{t} \in \Gamma_\omega \Rightarrow [\mathbf{t}] \in X_i))$$

The proofs of the following results are then exactly the same as the proofs of the analogous lemmas in the LKID case (cf. Lemmas 4.6 and 4.7, respectively).

Lemma 6.4. *For any inductive predicate P_i , if $M_\omega \models_{\rho_\omega} P_i \mathbf{t}$ then $\Gamma_\omega \vdash P_i \mathbf{t}$ is cut-free provable.*

Lemma 6.5. *If $F \in \Gamma_\omega$ then $M_\omega \models_{\rho_\omega} F$, and if $F \in \Delta_\omega$ then $M_\omega \not\models_{\rho_\omega} F$.*

As we are considering standard completeness for LKID^ω , we do not require to construct a Henkin class for M_ω as we did in the LKID case. Instead, the main remaining proof burden is to show that our constructed counter-model is indeed a standard model:

Lemma 6.6. *M_ω is a standard model for (Σ, Φ) .*

Proof. We must prove that $(P_1^{M_\omega}, \dots, P_n^{M_\omega})$ is the least prefixed point of φ_Φ . It is already a prefixed point of φ_Φ by definition. To see it is the least such point, it suffices to show that an arbitrary prefixed point $(X_1, \dots, X_n)$ of φ_Φ must also satisfy the other condition in the definition of $(P_1^{M_\omega}, \dots, P_n^{M_\omega})$, i.e. that $P_i \mathbf{t} \in \Gamma_\omega$ implies $[\mathbf{t}] \in X_i$, for any $i \in \{1, \dots, n\}$.

Suppose for contradiction that $P_i \mathbf{u} \in \Gamma_\omega$ but $[\mathbf{u}] \notin X_i$. By the construction of T_ω , it follows that there is a point along the untraceable branch π at which the rule (Case P_j) is applied with principal formula $P_i \mathbf{u}$, and so one of the case distinction premises of this rule instance, say:

$$\Gamma, \mathbf{u} = \mathbf{t}(\mathbf{y}), Q_1 \mathbf{u}_1(\mathbf{y}), \dots, Q_h \mathbf{u}_h(\mathbf{y}), P_{j_1} \mathbf{t}_1(\mathbf{y}), \dots, P_{j_m} \mathbf{t}_m(\mathbf{y}) \vdash \Delta$$

is a subsequence of $\Gamma_\omega \vdash \Delta_\omega$.

Now since $\varphi_i(X_1, \dots, X_n) \subseteq X_i$ (because $(X_1, \dots, X_n)$ is a prefixed point of φ_Φ), and $[\mathbf{u}] \notin X_i$ by assumption, we must have $[\mathbf{u}] \notin \varphi_i(X_1, \dots, X_n)$. In particular, $[\mathbf{u}] \notin \varphi_{i,r}(X_1, \dots, X_n)$, where $\Phi_{i,r}$ is the production used to obtain the case distinction above. That is, we have:

$$[\mathbf{u}] \notin \{[\mathbf{t}(\mathbf{y})] \mid Q_1^{M_\omega}[\mathbf{u}_1(\mathbf{y})], \dots, Q_h^{M_\omega}[\mathbf{u}_h(\mathbf{y})], [\mathbf{t}_1(\mathbf{y})] \in X_{j_1}, \dots, [\mathbf{t}_m(\mathbf{y})] \in X_{j_m}\}$$

We note that, since $\mathbf{u} = \mathbf{t}(\mathbf{y}) \in \Gamma_\omega$, we have $\mathbf{u} \sim \mathbf{t}(\mathbf{y})$ and thus $[\mathbf{u}] = [\mathbf{t}(\mathbf{y})]$. Similarly, since each of $Q_1 \mathbf{u}_1(\mathbf{y}), \dots, Q_h \mathbf{u}_h(\mathbf{y})$ is in Γ_ω , each of $Q_1^{M_\omega}[\mathbf{u}_1(\mathbf{y})], \dots, Q_h^{M_\omega}[\mathbf{u}_h(\mathbf{y})]$ holds. It follows that, for some $k \in \{1, \dots, m\}$, we must have $[\mathbf{t}_k(\mathbf{y})] \notin X_{j_k}$ (for otherwise we contradict the non-membership statement above). Furthermore, we can observe that $(P_i \mathbf{u}, \dots, P_i \mathbf{u}, P_{j_k} \mathbf{t}_k(\mathbf{y}))$ is a progressing trace that follows a finite segment of the untraceable branch π (starting with the point where $P_i \mathbf{u}$ first appears on the left of some sequent on the branch and finishing with the case distinction in which $P_{j_k} \mathbf{t}_k(\mathbf{y})$ appears).

But, since we have $P_{j_k} \mathbf{t}_k(\mathbf{y}) \in \Gamma$ but $[\mathbf{t}_k(\mathbf{y})] \notin X_{j_k}$, we can apply the same argument as was previously applied to $P_i \mathbf{u}$ and X_i to obtain another case-descendant of $P_{j_k} \mathbf{t}_k(\mathbf{y})$ and a progressing trace segment on π continuing from the first, and so on; and we conclude that there is an infinitely progressing trace following a tail of π , which gives the required contradiction. Thus $(P_1^{M_\omega}, \dots, P_n^{M_\omega})$ must indeed be the least prefixed point of φ_Φ , and so M_ω is a standard model for (Σ, Φ) , as required. $\square$

We can now complete the proof of Theorem 5.9. Suppose that $\Gamma \vdash \Delta$ is valid, i.e., true in every standard model of (Σ, Φ) , but that the search tree T_ω for $\Gamma \vdash \Delta$ is not an LKID $^\omega$ proof. Let $\Gamma_\omega \vdash \Delta_\omega$ be the limit sequent for $\Gamma \vdash \Delta$ with counter-interpretation (M_ω, ρ_ω) . By Lemma 6.5, the sequent $\Gamma_\omega \vdash \Delta_\omega$ is false in M_ω under the environment ρ_ω , and by Lemma 6.6, M_ω is indeed a standard model of (Σ, Φ) . Because $\Gamma \vdash \Delta$ is a subsequence of every sequent appearing in T_ω by construction, it is a subsequence of $\Gamma_\omega \vdash \Delta_\omega$, so $\Gamma \vdash \Delta$ is false in M_ω , which is a contradiction. Thus the recursive, cut-free search tree T_ω is a LKID $^\omega$ proof of the sequent $\Gamma \vdash \Delta$. $\square$

7 CLKID $^\omega$: a cyclic subsystem of LKID $^\omega$

In this section we investigate a cyclic subsystem, CLKID $^\omega$, of LKID $^\omega$, which arises naturally by restricting LKID $^\omega$ to proofs given by *regular* trees, i.e. those (possibly infinite) trees with only finitely many distinct subtrees. For example, although the LKID $^\omega$ proof of Figure 2 is not regular (since it contains infinitely many distinct variables $x_0, x_1, x_2, \dots$), it is easily transformed into a regular proof by using the substitution rule to insert a new sequent $Nx_0 \vdash Ex_0, Ox_0$ above the topmost sequent depicted. Concretely, regular LKID $^\omega$ proofs can be represented as finite graphs.

$$\begin{array}{c}
\frac{}{\vdash N0} (NR_1) \quad \frac{\frac{}{\vdash N0} (NR_1) \quad \frac{\underline{Ox} \vdash Nx \quad (\dagger)}{\underline{Oy} \vdash Ny} (\text{Subst})}{\underline{Oy} \vdash Nsy} (NR_2) \quad \frac{(*) \quad \underline{Ex} \vdash Nx}{\underline{Ey} \vdash Ny} (\text{Subst}) \\
\frac{}{\vdash N0} (=L) \quad \frac{\underline{Oy} \vdash Nsy}{x = sy, \underline{Oy} \vdash Nx} (=L) \quad \frac{\underline{Ey} \vdash Nsy}{x = sy, \underline{Ey} \vdash Nx} (=L) \\
\frac{\underline{Ex} \vdash Nx \quad (*)}{Ex \vee Ox \vdash Nx} (\text{Case } E) \quad \frac{(\dagger) \quad \underline{Ox} \vdash Nx}{Ex \vee Ox \vdash Nx} (\text{Case } O) \\
\hline
Ex \vee Ox \vdash Nx \quad (\vee L)
\end{array}$$

Figure 4: A CLKID^ω proof of $Ex \vee Ox \vdash Nx$. The symbols $(\dagger)$ and $(*)$ indicate the pairing of companions with buds.

Definition 7.1 (Companion). Let B be a bud of a derivation tree $\mathcal{D}$. An internal node C in $\mathcal{D}$ is said to be a *companion* for B if C and B are the same sequent.

Definition 7.2 (Cyclic pre-proof). A CLKID^ω *pre-proof* $\mathcal{P}$ of $\Gamma \vdash \Delta$ is a pair $(\mathcal{D}, \mathcal{R})$, where $\mathcal{D}$ is a finite derivation tree constructed according to the rules of LKID^ω (cf. §5) and whose root is $\Gamma \vdash \Delta$, and $\mathcal{R}$ is a function assigning a companion to every bud node in $\mathcal{D}$.

The *graph* of $\mathcal{P}$, written $\mathcal{G}_{\mathcal{P}}$, is the graph obtained from $\mathcal{D}$ by identifying each bud node B in $\mathcal{D}$ with its companion $\mathcal{R}(B)$.

By unfolding a cyclic pre-proof to its associated (possibly infinite) tree, it is immediate that cyclic pre-proofs generate exactly the class of LKID^ω pre-proofs given by the regular derivation trees. (Recall, an infinite tree is *regular* if it has only finitely many distinct subtrees.)

Definition 7.3 (Cyclic proof). A CLKID^ω *proof* is a CLKID^ω pre-proof whose graph satisfies the global trace condition of Definition 5.5.

It is immediate that cyclic proofs generate exactly the class of regular LKID^ω proofs. Thus CLKID^ω can be viewed as the restriction of LKID^ω to regular proofs.

Figure 4 shows a CLKID^ω pre-proof of the sequent $Ex \vee Ox \vdash Nx$, where N, E and O are respectively the “natural”, “even” and “odd” predicates given by Example 2.2. Any infinite path in the pre-proof has a tail consisting of repetitions of the “figure-of- ∞ ” loop given by the identification of the buds with their companions, whence there is an infinitely progressing trace on this tail given by the underlined formulas. Thus this pre-proof is indeed a CLKID^ω proof.

Proposition 7.4. *It is decidable whether a CLKID^ω pre-proof is a proof.*

Proof. Since there are only finitely many sequents in a CLKID^ω pre-proof $\mathcal{P}$ and each sequent is itself finite (and thus can admit only finitely many possible trace values), one can build a Büchi automaton B_1 accepting exactly those infinite strings of vertices of $\mathcal{G}_{\mathcal{P}}$ such that an infinitely progressing trace exists on some suffix of the string, and a second automaton B_2 accepting exactly those strings that are paths in $\mathcal{G}_{\mathcal{P}}$. Since Büchi automata are closed under language complementation and intersection, one can build a third automaton B accepting exactly those strings which are infinite paths in $\mathcal{G}_{\mathcal{P}}$ such that there does not exist an infinitely progressing trace on any tail of the path. $\mathcal{P}$ is then a CLKID^ω proof exactly if B accepts no strings, which is a decidable problem (cf. [39]).

The full construction appears in appendix A of [4]. Similar arguments also appear in [27, 34, 21]. $\square$

We now turn our attention to the question of the relationship between CLKID^ω and our system for induction, LKID . First, we show how to convert an LKID proof into a CLKID^ω

proof. Essentially, as shown in the next lemma, any use of induction over an inductive formula $P_j \mathbf{t}$ in an LKID proof can be replaced by (a) a cut on a formula which states that the minor premises of the rule together imply the induction hypothesis associated with $P_j \mathbf{t}$, together with (b) a CLKID^ω proof of the aforementioned formula.

Lemma 7.5. *Any instance of the LKID induction rule ($\text{Ind } P_j$) for an inductive predicate P_j is derivable in CLKID^ω .*

Proof. We recall the construction and associated notations for the induction rules for inductive predicates given in §3. We show how to derive an arbitrary instance of the induction rule ($\text{Ind } P_j$) in which the induction hypothesis F_i and the induction variables $\mathbf{z}_i$ have been associated to the inductive predicate P_i for each $i \in \{1, \dots, n\}$:

$$\frac{\text{minor premises } \Gamma, F_j \mathbf{t} \vdash \Delta}{\Gamma, P_j \mathbf{t} \vdash \Delta} (\text{Ind } P_j)$$

Now define $\mathcal{M}$ to be the smallest set of formulas such that for each minor premise of the considered instance of ($\text{Ind } P_j$), say:

$$\Gamma, Q_1 \mathbf{u}_1(\mathbf{x}), \dots, Q_h \mathbf{u}_h(\mathbf{x}), G_{j_1} \mathbf{t}_1(\mathbf{x}), \dots, G_{j_m} \mathbf{t}_m(\mathbf{x}) \vdash F_i \mathbf{t}(\mathbf{x}), \Delta$$

we have that a corresponding formula:

$$\forall \mathbf{y}. (Q_1 \mathbf{u}_1(\mathbf{y}) \wedge \dots \wedge Q_h \mathbf{u}_h(\mathbf{y}) \wedge G_{j_1} \mathbf{t}_1(\mathbf{y}) \wedge \dots \wedge G_{j_m} \mathbf{t}_m(\mathbf{y}) \rightarrow F_i \mathbf{t}(\mathbf{y}))$$

is in $\mathcal{M}$. Now, consider the following derivation in CLKID^ω (we write rule applications with a double line to indicate that we also apply the rule (Wk)):

$$\frac{\frac{\frac{\mathcal{M}, P_j \mathbf{z} \vdash F_j \mathbf{z}}{\mathcal{M}, P_j \mathbf{t} \vdash F_j \mathbf{t}} (\text{Subst})}{\vdots} (\wedge \text{L}) \quad \frac{\text{minor premises } \vdots}{\{\Gamma \vdash M, \Delta \mid M \in \mathcal{M}\}} (\wedge \text{R})}{\frac{\vdots}{\wedge \mathcal{M}, P_j \mathbf{t} \vdash F_j \mathbf{t}} (\wedge \text{L}) \quad \frac{\vdots}{\Gamma \vdash \wedge \mathcal{M}, \Delta} (\wedge \text{R})} \frac{\frac{P_j \mathbf{t} \vdash \wedge \mathcal{M} \rightarrow F_j \mathbf{t}}{\Gamma, \wedge \mathcal{M} \rightarrow F_j \mathbf{t} \vdash \Delta} (\rightarrow \text{R}) \quad \frac{\Gamma, F_j \mathbf{t} \vdash \Delta}{\Gamma, \wedge \mathcal{M} \rightarrow F_j \mathbf{t} \vdash \Delta} (\rightarrow \text{L})}{\Gamma, P_j \mathbf{t} \vdash \Delta} (\text{Cut})$$

where $\mathbf{z}$ is a tuple of appropriately many variables. We remark that obtaining each of the minor premises of the considered instance of ($\text{Ind } P_j$) from the sequents $\{\Gamma \vdash M, \Delta \mid M \in \mathcal{M}\}$ is simply a matter of decomposing each $M \in \mathcal{M}$ using the rules ($\forall \text{R}$), ($\rightarrow \text{R}$) and ($\wedge \text{L}$). It then remains to provide a CLKID^ω proof of the sequent $\mathcal{M}, P_j \mathbf{z} \vdash F_j \mathbf{z}$.

First, we apply the case-split rule ($\text{Case } P_j$) to the sequent $\mathcal{M}, P_j \mathbf{z} \vdash F_j \mathbf{z}$, thus generating a case for each production which has P_j occurring in its conclusion. We show how to treat a case arising from an arbitrary production, say:

$$\frac{Q_1 \mathbf{u}_1(\mathbf{x}) \dots Q_h \mathbf{u}_h(\mathbf{x}) P_{j_1} \mathbf{t}_1(\mathbf{x}) \dots P_{j_m} \mathbf{t}_m(\mathbf{x})}{P_j \mathbf{t}(\mathbf{x})}$$

For convenience, we shall use the following abbreviations for sets of formulas:

$$\begin{aligned} \mathcal{Q} &=_{\text{def}} Q_1 \mathbf{u}_1(\mathbf{y}), \dots, Q_h \mathbf{u}_h(\mathbf{y}) \\ \mathcal{P} &=_{\text{def}} P_{j_1} \mathbf{t}_1(\mathbf{y}), \dots, P_{j_m} \mathbf{t}_m(\mathbf{y}) \\ \mathcal{G} &=_{\text{def}} G_{j_1} \mathbf{t}_1(\mathbf{y}), \dots, G_{j_m} \mathbf{t}_m(\mathbf{y}) \end{aligned}$$

Now, as there is a minor premise $\Gamma, \mathcal{Q}, \mathcal{G} \vdash F_j \mathbf{t}(\mathbf{y}), \Delta$ corresponding to the production above in the considered instance of $(\text{Ind } P_j)$, we have:

$$\forall \mathbf{y} (\bigwedge \mathcal{Q} \wedge \bigwedge \mathcal{G} \rightarrow F_j \mathbf{t}(\mathbf{y})) \in \mathcal{M}$$

We choose to display this formula explicitly in the premise of the application of $(=L)$ in the following CLKID^ω derivation:

$$\begin{array}{c}
\frac{\frac{\frac{}{\mathcal{Q} \vdash Q_i \mid Q_i \in \mathcal{Q}} (\text{Ax})}{\mathcal{Q} \vdash \bigwedge \mathcal{Q}} (\wedge R)}{\vdots} \quad \frac{\frac{\frac{\{\mathcal{M}, P_{j_i} \mathbf{z} \vdash G_{j_i} \mathbf{z} \mid 1 \leq i \leq m\} (\dagger 2)}{\{\mathcal{M}, P_{j_i} \mathbf{t}_i(\mathbf{y}) \vdash G_{j_i} \mathbf{t}_i(\mathbf{y}) \mid 1 \leq i \leq m\}} (\text{Subst})}{\vdots} (\wedge R)}{\mathcal{M}, \mathcal{P} \vdash \bigwedge \mathcal{G}} (\wedge R) \\
\frac{\frac{F_j \mathbf{t}(\mathbf{y}) \vdash F_j \mathbf{t}(\mathbf{y}) (\text{Ax})}{\mathcal{M}, \mathcal{Q}, \mathcal{P} \vdash \bigwedge \mathcal{Q} \wedge \bigwedge \mathcal{G}} (\wedge R)}{\mathcal{M}, \bigwedge \mathcal{Q} \wedge \bigwedge \mathcal{G} \rightarrow F_j \mathbf{t}(\mathbf{y}), \mathcal{Q}, \mathcal{P} \vdash F_j \mathbf{t}(\mathbf{y})} (\rightarrow L) \\
\frac{}{\vdots} (\forall L) \\
\frac{\mathcal{M}, \forall \mathbf{y}. (\bigwedge \mathcal{Q} \wedge \bigwedge \mathcal{G} \rightarrow F_j \mathbf{t}(\mathbf{y})), \mathcal{Q}, \mathcal{P} \vdash F_j \mathbf{t}(\mathbf{y})}{\mathcal{M}, \mathbf{z} = \mathbf{t}(\mathbf{y}), \mathcal{Q}, \mathcal{P} \vdash F_j \mathbf{z}} (=L) \\
\text{other cases } \dots \quad \frac{}{\mathcal{M}, \mathbf{z} = \mathbf{t}(\mathbf{y}), \mathcal{Q}, \mathcal{P} \vdash F_j \mathbf{z}} (\text{Case } P_j) \\
\hline
\mathcal{M}, P_j \mathbf{z} \vdash F_j \mathbf{z} (\dagger 1)
\end{array}$$

where $\mathbf{y}$ is suitably fresh. We have thus far obtained a CLKID^ω derivation with root sequent $\mathcal{M}, P_j \mathbf{z} \vdash F_j \mathbf{z} (\dagger 1)$ and bud nodes $\{\mathcal{M}, P_{j_i} \mathbf{z} \vdash G_{j_i} \mathbf{z} \mid i \in \{1, \dots, m\}\} (\dagger 2)$, and we observe that for each $i \in \{1, \dots, m\}$ there is a progressing trace:

$$(P_j \mathbf{z}, P_{j_i} \mathbf{t}_i(\mathbf{y}), \dots, P_{j_i} \mathbf{t}_i(\mathbf{y}), P_{j_i} \mathbf{z})$$

following the path in this derivation from the root sequent $(\dagger 1)$ to the bud $\mathcal{M}, P_{j_i} \mathbf{z} \vdash G_{j_i} \mathbf{z}$. Now note that, for each $i \in \{1, \dots, m\}$, if the predicates P_{j_i} and P_j are not mutually dependent, then $G_{j_i} = P_{j_i}$, and so we may apply the rule (Ax) to the bud node $\mathcal{M}, P_{j_i} \mathbf{z} \vdash G_{j_i} \mathbf{z}$. Thus we need to consider only the bud nodes $\mathcal{M}, P_{j_i} \mathbf{z} \vdash G_{j_i} \mathbf{z}$ such that P_{j_i} and P_j are mutually dependent, and are thus of the form $\mathcal{M}, P_{j_i} \mathbf{z} \vdash F_{j_i} \mathbf{z}$. We treat these as follows:

- if $P_{j_i} = P_j$, then the bud node is identical to the root sequent $(\dagger 1)$, and we set the companion of the bud to be $(\dagger 1)$.
- if $P_{j_i} \neq P_j$, then note that as P_{j_i} and P_j are mutually dependent, there is a minor premise (and corresponding formula in $\mathcal{M}$) for every production which has P_{j_i} occurring in its conclusion. We thus can repeat the derivation above for the bud node under consideration to obtain new bud nodes $(\dagger 3)$, to which we may assign $(\dagger 1)$ or any ancestor node of the form $(\dagger 2)$ as a companion.

We iterate this process as often as required, successively generating bud nodes of the form $(\dagger 3), (\dagger 4), \dots$, noting that any bud node of the form $(\dagger k)$ may potentially be assigned an ancestral companion of the form $(\dagger k')$ for any $k' \in \{1, \dots, k-1\}$, and that bud nodes are always assigned ancestors as companions. This iteration is possible because $\mathcal{M}$ contains a formula corresponding to each production having in its conclusion a predicate that is mutually dependent with P_j and, since mutual dependency between predicates is transitive, the predicate P_{j_i} occurring on the left of any bud node $(\dagger k)$ is always mutually dependent with P_j . Also, we observe that the iteration

process never produces bud nodes of the form $(\dagger n + 2)$ (and so must terminate), because there are only n distinct inductive predicate symbols.

We thus obtain a CLKID^ω derivation tree $\mathcal{D}$ with root sequent $\mathcal{M}, P_j \mathbf{z} \vdash F_j \mathbf{z}$ and a repeat function $\mathcal{R}$ that assigns to every bud node of $\mathcal{D}$ an ancestor of the bud as companion, i.e. $(\mathcal{D}, \mathcal{R})$ is a CLKID^ω pre-proof. Furthermore, for each bud node B in the tree, there is a trace τ_i following the unique path in $\mathcal{D}$ from $\mathcal{R}(B)$ to B that takes the same value at B and $\mathcal{R}(B)$. This is sufficient to ensure that $(\mathcal{D}, \mathcal{R})$ is a CLKID^ω proof (for full details see Lemma 7.3.1 of [4]), which completes the derivation of the considered instance of $(\text{Ind } P_j)$ inside CLKID^ω . $\square$

Theorem 7.6. *Every LKID proof of $\Gamma \vdash \Delta$ can be transformed into a CLKID^ω proof of $\Gamma \vdash \Delta$.*

Proof. Given any LKID proof $\mathcal{D}$ of $\Gamma \vdash \Delta$ we can obtain a CLKID^ω pre-proof $\mathcal{P}$ of $\Gamma \vdash \Delta$ by replacing every instance of an induction rule in $\mathcal{D}$ with the corresponding CLKID^ω derivation constructed in Lemma 7.5. Furthermore, by inspection it is clear that this process does not create any new overlaps between the cycles in the inserted derivations. The fact that $\mathcal{P}$ is a CLKID^ω proof thus follows immediately from the fact that the derivations replacing induction rule instances each contain a CLKID^ω proof of one branch, and the required minor premises as the only other bud nodes of the other branches. $\square$

Interestingly, our translation makes essential use of both the cut and substitution rules in CLKID^ω . Indeed, it seems certain that cut is not eliminable from the system CLKID^ω , and it is at least plausible that neither is substitution (the importance of which is illustrated in the discussion of Figure 2 at the beginning of this section.) Nevertheless, CLKID^ω arises naturally as the restriction of a complete infinitary proof system to proofs with finite representation. The main open question relating to it is whether the converse to Theorem 7.6 holds. We strongly believe this to be the case, and hence present it as a conjecture.

Conjecture 7.7. *If there is a CLKID^ω proof of a sequent $\Gamma \vdash \Delta$ then there is an LKID proof of $\Gamma \vdash \Delta$.*

This conjecture does not seem straightforward. For example, the methods applied in [35], which show, in a different setting, the equivalence of a weaker global proof condition with a local transfinite induction principle, do not adapt.

An interesting perspective on the conjecture can be obtained by making a comparison with Kozen’s modal μ -calculus [19]. Because of the use of explicit induction rules, there is an obvious analogy between proofs in LKID and proofs in Kozen’s axiomatization for the modal μ -calculus. Also, as we already commented in Section 5, there is an analogy between proofs in LKID^ω and the *refutations* introduced for the modal μ -calculus by Niwinski and Walukiewicz [27]. This latter analogy restricts to one between proofs in CLKID^ω and *regular refutations* for the modal μ -calculus. Given these correspondences, we observe that the analogous result to Conjecture 7.7 does hold for the modal μ -calculus; that is, a formula has a proof in Kozen’s axiomatization if and only if it is provable via a regular refutation (by which, strictly speaking, we mean that the negation of the formula has a regular refutation). This equivalence is established by the following chain of reasoning. By the results in [27], a modal μ -calculus formula is provable via a regular refutation if and only if it is provable via any refutation if and only if it is valid. Finally, by Walukiewicz’ celebrated completeness theorem for Kozen’s axiomatization [43], one has that a formula is valid if and only if it is provable in Kozen’s system.

Unfortunately, the above chain of reasoning is seriously broken in the case of Conjecture 7.7, where the situation is instead as summarised in Figure 5. One issue is that it is not the case that provability in CLKID^ω and LKID^ω coincide; regularity is a real restriction. Similarly, Henkin validity (which characterises LKID-provability) does not coincide with standard validity (which characterises LKID^ω -provability).

Nevertheless, we believe that the equivalence between Kozen proofs and regular refutations for the modal μ -calculus provides weak evidence for Conjecture 7.7 in the following sense. Since it is the case for the modal μ -calculus that every proof via a regular refutation has a corresponding proof in Kozen’s axiomatization, it is plausible that it might be possible to prove this by a combinatorial transformation of one type of proof into the other. If so, it appears likely that the same combinatorial method would then adapt to establish Conjecture 7.7. Unfortunately, in spite of intensive effort, we have not been able to get any such combinatorial method to succeed. An alternative approach to Conjecture 7.7 would be to obtain a proof of soundness for CLKID^ω proofs with respect to Henkin models. However, we also have no idea how to achieve this.

Figure 5: A diagrammatic summary of our developments. The solid arrows are implications, and the dashed arrow indicates our conjectured implication.

8 Conclusions and discussion

The goal of the present paper has been to develop and compare proof-theoretic foundations for proof by induction and proof by infinite descent. We have addressed this within the context of first-order logic with inductively defined predicates, FOL_{ID} . Each of the two styles of reasoning has been formulated as a sequent calculus for which an appropriate cut-free completeness result has been proved. The proof system LKID, formalizing proof by induction, is complete relative to a class of Henkin models. The infinitary proof system, LKID^ω , formalising proof by infinite descent, is complete relative to the more restrictive class of standard models, and, as a result, is strictly more powerful. The infinitary system has a natural subsystem, CLKID^ω , consisting

of cyclic proofs, that is proofs given by regular trees. Cyclic proofs are at least as powerful as proofs by explicit induction. Our main outstanding conjecture is that they are no more powerful.

It is useful to consider our contributions in comparison to analogous research on propositional fixed-point logics, such as Kozen’s modal μ -calculus [19]. In the discussion at the end of Section 7, we commented that the use of explicit induction means that there is an analogy between LKID and Kozen’s axiomatization of the modal μ -calculus [19]. Likewise, our standard models, in which fixed-points are obtained via approximants, are analogous to the usual Kripke-frame models of the modal μ -calculus, also considered in *op. cit.* Our more general notion of Henkin model also has an analogue, for the modal μ -calculus, given by the *modal μ -frames* introduced in [2, 16]. In these last references, the completeness of Kozen’s axiomatization is established relative to modal μ -frames, and our Theorem 3.6 can be seen as analogous to these results.² Independently to our work, Kashima and Okamoto [18] have extended the completeness results of [2, 16] to a first-order setting, using a notion of *general model*, which plays a role for the first-order modal μ -calculus identical to that played by Henkin models in our paper. Compared with their work, our development differs in being for a different logic, and an additional contribution of the present paper is the use of completeness to establish cut-admissibility for LKID.

As discussed in Section 5, the *refutations* of Niwinski and Walukiewicz [27] give a tableau-based proof system for the modal μ -calculus which is analogous to our infinitary proof system LKID^ω in its use of non-well-founded derivations and a trace-based proof condition. Similarly, *regular refutations* are analogous to proofs in our cyclic system CLKID^ω . The strong analogy between our systems and refutations can be further appreciated if the latter are recast in a sequent calculus format, as has been done explicitly in [10, 36]. Another propositional setting in which cyclic sequent proof systems have been considered is that of μ -lattices (which model a propositional logic of linear conjunction and disjunction with fixed points), and of related categories with fixed points (called Ω -models), both of which have been studied by Santocanale [30, 29]. Two main properties distinguish the aforementioned proof systems for propositional fixed point logics, influenced by refutations, from the first-order systems of the present paper. First, for all the propositional systems mentioned above, it holds that cut is admissible over cut-free cyclic proofs, a situation that we believe not to be the case for our cyclic system CLKID^ω (see Section 7). Second, for the systems considered in [10, 36, 30], it holds that the cyclic (regular) proof systems are as powerful as the full infinitary systems, whereas we know that this does not hold in the first-order setting, since LKID^ω is complete for standard validity by Theorem 5.9 but this cannot be the case for CLKID^ω for the same reason as LKID (see Theorem 3.13). The last distinction has a semantic correlate. For the modal μ -calculus, validity in the modal μ -frames of [2, 16] is equivalent to validity in standard Kripke frames, as follows from Walukiewicz’ completeness theorem for Kozen’s axiomatization [43]. Similarly, validity in Santocanale’s μ -lattices is equivalent to validity in the restricted class of complete lattices [30].³ However, in the setting of FOL_{ID} , validity in Henkin models is strictly weaker than validity in standard models.

Summarising the entire discussion above, for the first-order logic with inductive predicates considered in the present paper, we have the situation depicted in Figure 5. However, for propositional logics with fixed-points, there are typically extra identifications, since the distinction between (the analogues of) Henkin validity and standard validity disappears, as does the distinction between provability in the infinitary and cyclic systems.

There are many natural directions for future research. One gap to fill is the absence of syntactic proofs of cut elimination for LKID and LKID^ω . In the case of LKID it would be nice to obtain an elementary proof in primitive recursive arithmetic augmented by quantifier-

²As discussed in Section 7, Walukiewicz’ completeness theorem for Kozen’s axiomatization relative to standard (Kripke-frame) models [43] has no analogue in our setting since LKID is incomplete for standard models.

³Both these coincidences can be viewed as consequences of the fact that free algebraic structures with fixed-points can be embedded into corresponding complete lattices, a perspective that is developed in [30, 31].

free transfinite induction up to ϵ_0 (the use of transfinite induction is necessary because cut-elimination for LKID implies the consistency of PA, as demonstrated by our Theorem 3.15). In the case of LKID^ω , it would be interesting to see if Mints’ “continuous cut elimination” (see, e.g., [13, §§6.2.7–11]), which applies to non-well-founded derivation trees, adapts to LKID^ω pre-proofs, and, further, preserves the global trace condition. Also, while on the subject of alternative proof methods, we wonder if Niwinski and Walukiewicz’ idea of deriving completeness from determinacy results, [27], can be used to give an alternative proof of completeness for LKID^ω , rather than the direct proof we gave in §5.

In the direction of extending the results obtained here, we wonder whether more liberal subsystems of LKID^ω than CLKID^ω might also be suitable for formal proof. For example, one might restrict to LKID^ω proofs generated by pushdown automata or by recursion schemata, over which the global trace condition should still be decidable. We wonder if such proofs lead to an increase in power over regular proofs. Further, we wonder if, for some suitably chosen such class of proofs, the cut rule is eliminable. (A syntactic proof of cut-elimination for LKID^ω would be likely to be indispensable in investigating this question.) In the opposite direction, one might also consider more restrictive cyclic systems obtained by tightening the global trace condition to improve its computational complexity (cf. [21, 34, 4]).

We comment that it should be relatively straightforward to extend our proof systems LKID, LKID^ω and CLKID^ω (together with the completeness results for the first two) to more general (co)inductive definition schemas, for example to iterated inductive definitions [23], or to the first-order μ -calculus, cf. [35]. Indeed, the latter should be naturally accommodated by requiring a parity-game-style condition to hold of traces along infinite paths (needed to address alternating least and greatest fixed-points), as in the refutations of Niwinski and Walukiewicz [27] (and in the various proof systems derived from refutations, discussed above). Another context in which cyclic proofs might prove useful is in the context of dependent type theory, where cyclic proofs might be used to justify termination for functions defined by pattern matching (corresponding to our case-split rules) and general recursion (corresponding to proof cycles). Independently of the work in this paper, such a system has been implemented for first-order parameterized data types in Wahlstedt’s 2007 PhD thesis [42].

A practical motivation behind our investigations was to increase and improve our understanding of the armoury of proof methods available for inductive theorem proving. In this regard, we believe our main contribution is in providing a firm foundation for cyclic reasoning. Currently, cyclic methods are sometimes included in theorem provers, with various heuristic conditions for guaranteeing soundness, cf. [44, 32, 14]. It would be a worthwhile project to systematically study such heuristic conditions and to relate them to our trace-based condition, which we believe to subsume them all.

There are already indications that cyclic reasoning, as formulated in the present paper, does have potential applications. The style of cyclic reasoning we have developed for FOL_{ID} has been adapted theoretically to the setting of the bunched logic BI [5] and to Hoare-style termination proofs in separation logic [6], and is also beginning to see practical use in theorem proving tools [26, 41]. Plausibly, cyclic reasoning is also likely to prove especially useful for demonstrating properties of mutually defined relations, for which the associated induction principles are often extremely complex. We hope that our foundational presentation of cyclic proofs here will be of assistance in future practical developments.

References

- [1] Peter Aczel. An introduction to inductive definitions. In Jon Barwise, editor, *Handbook of Mathematical Logic*, pages 739–782. North-Holland, 1977.

- [2] Simon Ambler, Marta Z. Kwiatkowska, and Nicholas Measor. On duality for the modal μ -calculus. In *Proceedings of CSL 1994*, volume 832 of *LNCS*, pages 18–32. Springer-Verlag, 1994.
- [3] James Brotherston. Cyclic proofs for first-order logic with inductive definitions. In *Proceedings of TABLEUX 2005*, volume 3702 of *LNAI*, pages 78–92. Springer-Verlag, 2005.
- [4] James Brotherston. *Sequent Calculus Proof Systems for Inductive Definitions*. PhD thesis, University of Edinburgh, November 2006.
- [5] James Brotherston. Formalised inductive reasoning in the logic of bunched implications. In *Proceedings of SAS-14*, volume 4634 of *LNCS*, pages 87–103. Springer-Verlag, 2007.
- [6] James Brotherston, Richard Bornat, and Cristiano Calcagno. Cyclic proofs of program termination in separation logic. In *Proceedings of POPL-35*, pages 101–112. ACM, 2008.
- [7] James Brotherston and Alex Simpson. Complete sequent calculi for induction and infinite descent. In *Proceedings of LICS-22*, pages 51–60. IEEE Computer Society, 2007.
- [8] Alan Bundy. The automation of proof by mathematical induction. In Alan Robinson and Andrei Voronkov, editors, *Handbook of Automated Reasoning*, volume I, chapter 13, pages 845–911. Elsevier Science, 2001.
- [9] Samuel R. Buss. *Handbook of Proof Theory*. Elsevier Science, 1998.
- [10] Christian Dax, Martin Hofmann, and Martin Lange. A proof system for the linear time μ -calculus. In *Proceedings of FSTTCS 2006*, volume 4337 of *LNCS*, pages 273–284. Springer-Verlag, 2006.
- [11] Anatoli Degtyarev and Andrei Voronkov. Equality reasoning in sequent-based calculi. In Alan Robinson and Andrei Voronkov, editors, *Handbook of Automated Reasoning*, volume I, chapter 10, pages 611–706. Elsevier Science, 2001.
- [12] Gerhard Gentzen. Untersuchungen über das Logische Schliessen. *Mathematische Zeitschrift*, 39:176–210 and 405–431, 1935. Reproduced with English translation in [37], 68–131.
- [13] Jean-Yves Girard. *Proof Theory and Logical Complexity*, volume 1. Bibliopolis, 1987.
- [14] Geoff Hamilton. Poitin: Distilling theorems from conjectures. In *Proceedings of the 12th Symposium on the Integration of Symbolic Computation and Mechanized Reasoning*, volume 151 of *ENTCS*, pages 143–160. Elsevier B.V., July 2005.
- [15] John Harrison. Inductive definitions: automation and application. In *Higher Order Logic Theorem Proving and Its Applications: Proceedings of the 8th International Workshop*, volume 971 of *LNCS*, pages 200–213. Springer-Verlag, 1995.
- [16] Chrysafis Hartonas. Duality for modal μ -logics. *Theoretical Computer Science*, 202(1–2):193–222, 1998.
- [17] Leon Henkin. Completeness in the theory of types. *Journal of Symbolic Logic*, 15:81–91, 1950.
- [18] Ryo Kashima and Keishi Okamoto. General models and completeness of first-order modal μ -calculus. *Journal of Logic and Computation*, 18(4):497–507, 2008.

- [19] Dexter Kozen. Results on the propositional μ -calculus. *Theoretical Computer Science*, 27:333–354, 1983.
- [20] Georg Kreisel. Mathematical logic. In T.L. Saaty, editor, *Lectures on Modern Mathematics*, volume III, pages 95–105. Wiley and Sons, 1965.
- [21] Chin Soon Lee, Neil D. Jones, and Amir M. Ben-Amram. The size-change principle for program termination. *ACM SIGPLAN Notices*, 36(3):81–92, 2001.
- [22] Michael Sean Mahoney. *The Mathematical Career of Pierre de Fermat 1601–1665*. Princeton University Press, 2nd edition, 1994.
- [23] Per Martin-Löf. Hauptstatz for the intuitionistic theory of iterated inductive definitions. In *Proceedings of the Second Scandinavian Logic Symposium*, pages 179–216. North-Holland, 1971.
- [24] Raymond McDowell and Dale Miller. Cut-elimination for a logic with definitions and induction. *Theoretical Computer Science*, 232:91–119, 2000.
- [25] Alberto Momigliano and Alwen Tiu. Induction and co-induction in sequent calculus. In *Proceedings of TYPES 2003*, volume 3085 of *LNCS*, pages 293 – 308. Springer-Verlag, 2003.
- [26] Huu Hai Nguyen and Wei-Ngan Chin. Enhancing program verification with lemmas. In *Proceedings of CAV 2008*, volume 5123 of *LNCS*, pages 355–369. Springer-Verlag, 2008.
- [27] Damian Niwiński and Igor Walukiewicz. Games for the μ -calculus. *Theoretical Computer Science*, 163:99–116, 1997.
- [28] Christine Paulin-Mohring. Inductive definitions in the system Coq: Rules and properties. In *Proceedings of TLCA 1993*, volume 664 of *LNCS*, pages 328–345. Springer-Verlag, 1993.
- [29] Luigi Santocanale. A calculus of circular proofs and its categorical semantics. In *Proceedings of FOSSACS 2002*, volume 2303 of *LNCS*, pages 357–371. Springer-Verlag, 2002.
- [30] Luigi Santocanale. Free μ -lattices. *Journal of Pure and Applied Algebra*, 168(2–3):227–264, 2002.
- [31] Luigi Santocanale. Completions of μ -algebras. *Annals of Pure and Applied Logic*, 154(1):27–50, 2008.
- [32] Carsten Schürmann. *Automating the Meta-Theory of Deductive Systems*. PhD thesis, Carnegie-Mellon University, 2000.
- [33] Konrad Slind. Derivation and use of induction schemes in higher-order logic. In *Proceedings of TPHOLs*, volume 1275 of *LNCS*. Springer-Verlag, 1997.
- [34] Christoph Sprenger and Mads Dam. A note on global induction mechanisms in a μ -calculus with explicit approximations. *Theoretical Informatics and Applications*, 37:365–399, July 2003.
- [35] Christoph Sprenger and Mads Dam. On the structure of inductive reasoning: circular and tree-shaped proofs in the μ -calculus. In *Proceedings of FOSSACS 2003*, volume 2620 of *LNCS*, pages 425–440. Springer-Verlag, 2003.
- [36] Thomas Studer. On the proof theory of the modal μ -calculus. *Studia Logica*, 89(3):343–363, 2008.

- [37] M.E. Szabo, editor. *The Collected Papers of Gerhard Gentzen*. North-Holland, 1969.
- [38] William W. Tait. Intentional interpretations of functionals of finite type. *Journal of Symbolic Logic*, 32:198–212, 1967.
- [39] Wolfgang Thomas. Automata on infinite objects. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, volume B: Formal Models and Semantics, pages 133–192. Elsevier Science Publishers, 1990.
- [40] Alwen Tiu. *A Logical Framework For Reasoning About Logical Specifications*. PhD thesis, Pennsylvania State University, 2004.
- [41] Răzvan Voicu and Mengran Li. Descente infinie proofs in Coq. In *Proceedings of the 1st Coq Workshop*. Technische Universität München, 2009.
- [42] David Wahlstedt. *Dependent Type Theory with First-Order Parameterized Data Types and Well-Founded Recursion*. PhD thesis, Chalmers University of Technology, 2007.
- [43] Igor Walukiewicz. Completeness of Kozen’s axiomatisation of the propositional μ -calculus. *Information and Computation*, 157:142–182, 2000.
- [44] Claus-Peter Wirth. Descente infinie + Deduction. *Logic Journal of the IGPL*, 12(1):1–96, 2004.

A Proof of soundness for LKID induction rules

In this appendix we provide a detailed proof of the crucial case of Proposition 3.5: the Henkin soundness of the induction rules of LKID.

Consider an arbitrary induction rule instance:

$$\frac{\text{minor premises} \quad \Gamma, F_j \mathbf{t} \vdash \Delta}{\Gamma, P_j \mathbf{t} \vdash \Delta} (\text{Ind } P_j)$$

and, for each $i \in \{1, \dots, n\}$, let $\mathbf{z}_i$ and G_i be, respectively, the induction variables and the formula associated with P_i in the construction of this instance of $(\text{Ind } P_j)$, cf. Section 3. We write k_i for the arity of P_i .

We assume for contradiction that all the premises of the rule instance are Henkin valid whereas the conclusion is invalid, i.e. false in some Henkin model $(M, \mathcal{H})$. Thus for some environment ρ we have $M \models_\rho J$ for all $J \in \Gamma$ and $M \models_\rho P_j \mathbf{t}$, i.e. $\rho(\mathbf{t}) \in \pi_j^n(\mu_{\mathcal{H}}.\varphi_\Phi)$, but $M \not\models_\rho K$ for all $K \in \Delta$.

Next, let $\mathbf{y}$ be the fresh variables appearing in the minor premises and let $\mathbf{e}$ be a tuple of arbitrary elements of D such that $|\mathbf{y}| = |\mathbf{e}|$. Define an environment ρ' by $\rho' = \rho[\mathbf{x} \mapsto \mathbf{e}]$, and note that since $x \notin FV(\Gamma \cup \Delta)$ for all $x \in \mathbf{x}$ by the rule side condition, it holds that $M \models_{\rho'} J$ for all $J \in \Gamma$ and $M \not\models_{\rho'} K$ for all $K \in \Delta$.

Now define an n -tuple of sets $(Y_1, \dots, Y_n) \in \mathcal{P}(D^{k_1}) \times \dots \times \mathcal{P}(D^{k_n})$ by:

$$Y_i = \begin{cases} \{\mathbf{d} \in D^{k_i} \mid M \models_{\rho'[\mathbf{z}_i \mapsto \mathbf{d}]} G_i\} & \text{if } \text{Prem}^*(P_j, P_i) \\ D^{k_i} & \text{otherwise} \end{cases}$$

for each $i \in \{1, \dots, n\}$, where $\mathbf{z}_i$ is the tuple of k_i induction variables for P_i . We assert the following:

CLAIM: $(Y_1, \dots, Y_n)$ is a prefixed $\mathcal{H}$ -point of φ_Φ .

Then, since $\mu_{\mathcal{H}}.\varphi_\Phi$ is the least prefixed $\mathcal{H}$ -point of φ_Φ , it holds that $\pi_j^n(\mu_{\mathcal{H}}.\varphi_\Phi) \subseteq Y_j$. As the major premise $\Gamma, F_j \mathbf{t} \vdash \Delta$ is valid and so true in $(M, \mathcal{H})$, and as $M \models_{\rho'} J$ for all $J \in \Gamma$ but $M \not\models_{\rho'} K$ for all $K \in \Delta$, we must have $M \not\models_{\rho'} F_j \mathbf{t}$. We then have $M \not\models_{\rho'[\mathbf{z}_j \mapsto \rho'(\mathbf{t})]} F_j$, i.e. $\rho'(\mathbf{t}) \notin Y_j$, so also $\rho'(\mathbf{t}) \notin \pi_j^n(\mu_{\mathcal{H}}.\varphi_\Phi)$. As $x \notin FV(P_j \mathbf{t})$ for all $x \in \mathbf{x}$ by the rule side condition, i.e. no variable $x \in \mathbf{x}$ occurs in $\mathbf{t}$, we must have $\rho'(\mathbf{t}) = \rho(\mathbf{t})$. But then we have $\rho(\mathbf{t}) \notin \pi_j^n(\mu_{\mathcal{H}}.\varphi_\Phi)$, which is a contradiction, as required.

To finish the proof, it suffices to prove the claim above that $(Y_1, \dots, Y_n)$ is a prefixed $\mathcal{H}$ -point of φ_Φ . First, writing $\mathcal{H} = \{H_k \mid k \in \mathbb{N}\}$, observe that by Lemma 2.7, $Y_i \in H_{k_i}$ for each $i \in \{1, \dots, n\}$, i.e. $(Y_1, \dots, Y_n)$ is an $\mathcal{H}$ -point. It remains to show that $(Y_1, \dots, Y_n)$ is a prefixed point of φ_Φ , i.e., that $\varphi_i(Y_1, \dots, Y_n) \subseteq Y_i$ for each $i \in \{1, \dots, n\}$. We argue by cases on i as follows:

1. $\neg \text{Prem}^*(P_j, P_i)$ holds. It is then trivial that $\varphi_i(Y_1, \dots, Y_n) \subseteq D^{k_i} = Y_i$.
2. $\text{Prem}^*(P_j, P_i)$ and $\neg \text{Prem}^*(P_i, P_j)$ hold. As P_j and P_i are thus not mutually dependent, $G_i = P_i \mathbf{z}_i$ and we have:

$$\begin{aligned} Y_i &= \{\mathbf{d} \in D^{k_i} \mid M \models_{\rho'[\mathbf{z}_i \mapsto \mathbf{d}]} P_i \mathbf{z}_i\} \\ &= \{\mathbf{d} \in D^{k_i} \mid \mathbf{d} \in \pi_i^n(\mu_{\mathcal{H}}.\varphi_\Phi)\} \\ &= \pi_i^n(\mu_{\mathcal{H}}.\varphi_\Phi) \end{aligned}$$

It suffices to show that $\varphi_{i,r}(Y_1, \dots, Y_n) \subseteq Y_i = \pi_i^n(\mu_{\mathcal{H}}.\varphi_\Phi)$ for an arbitrary production $\Phi_{i,r}$, say:

$$\frac{Q_1 \mathbf{u}_1(\mathbf{x}) \dots Q_h \mathbf{u}_h(\mathbf{x}) P_{j_1} \mathbf{t}_1(\mathbf{x}) \dots P_{j_m} \mathbf{t}_m(\mathbf{x})}{P_i \mathbf{t}(\mathbf{x})}$$

We are thus required to show:

$$\{\mathbf{t}^M(\mathbf{d}) \mid Q_1^M \mathbf{u}_1^M(\mathbf{d}), \dots, Q_h^M \mathbf{u}_h^M(\mathbf{d}), \mathbf{t}_1^M(\mathbf{d}) \in Y_{j_1}, \dots, \mathbf{t}_m^M(\mathbf{d}) \in Y_{j_m}\} \subseteq \pi_i^n(\mu_{\mathcal{H}}.\varphi_\Phi)$$

Note that for each of the inductive predicates P_{j_k} appearing in the premises of the production $\Phi_{i,r}$, $\text{Prem}^*(P_j, P_{j_k})$ holds (because $\text{Prem}^*(P_j, P_i)$ and $\text{Prem}(P_i, P_{j_k})$ hold), and $\neg \text{Prem}^*(P_{j_k}, P_j)$ holds (because otherwise we would have $\text{Prem}^*(P_i, P_j)$, which contradicts the case assumption), and we therefore have $Y_{j_k} = \pi_{j_k}^n(\mu_{\mathcal{H}}.\varphi_\Phi)$ by a similar argument to the one above concerning Y_i . We can therefore rewrite the statement we need to prove as:

$$\{\mathbf{t}^M(\mathbf{d}) \mid Q_1^M \mathbf{u}_1^M(\mathbf{d}), \dots, Q_h^M \mathbf{u}_h^M(\mathbf{d}), \mathbf{t}_1^M(\mathbf{d}) \in \pi_{j_1}^n(\mu_{\mathcal{H}}.\varphi_\Phi), \dots, \mathbf{t}_m^M(\mathbf{d}) \in \pi_{j_m}^n(\mu_{\mathcal{H}}.\varphi_\Phi)\} \subseteq \pi_i^n(\mu_{\mathcal{H}}.\varphi_\Phi)$$

i.e., $\varphi_{i,r}(\mu_{\mathcal{H}}.\varphi_\Phi) \subseteq \pi_i^n(\mu_{\mathcal{H}}.\varphi_\Phi)$, which is true since $\mu_{\mathcal{H}}.\varphi_\Phi$ is a prefixed $\mathcal{H}$ -point of φ_Φ . This completes the case.

3. $\text{Prem}^*(P_j, P_i)$ and $\text{Prem}^*(P_i, P_j)$ both hold, i.e. P_i and P_j are mutually dependent. As in the previous case, we require to show $\varphi_{i,r}(Y_1, \dots, Y_n) \subseteq Y_i$, i.e.

$$\{\mathbf{t}^M(\mathbf{d}) \mid Q_1^M \mathbf{u}_1^M(\mathbf{d}), \dots, Q_h^M \mathbf{u}_h^M(\mathbf{d}), \mathbf{t}_1^M(\mathbf{d}) \in Y_{j_1}, \dots, \mathbf{t}_m^M(\mathbf{d}) \in Y_{j_m}\} \subseteq Y_i$$

As P_i and P_j are mutually dependent, there is a minor premise of the instance of (Ind P_j):

$$\Gamma, Q_1 \mathbf{u}_1(\mathbf{y}), \dots, Q_h \mathbf{u}_h(\mathbf{y}), G_{j_1} \mathbf{t}_1(\mathbf{y}), \dots, G_{j_m} \mathbf{t}_m(\mathbf{y}) \vdash F_i \mathbf{t}(\mathbf{y}), \Delta$$

As each minor premise is true in $(M, \mathcal{H})$ by assumption and we have $M \models_{\rho'} J$ for all $J \in \Gamma$ and $M \not\models_{\rho'} K$ for all $K \in \Delta$, the following implication holds:

$$\begin{array}{l} M \models_{\rho'} Q_1 \mathbf{u}_1(\mathbf{y}), \dots, M \models_{\rho'} Q_h \mathbf{u}_h(\mathbf{y}), \\ M \models_{\rho'} G_{j_1} \mathbf{t}_1(\mathbf{y}), \dots, M \models_{\rho'} G_{j_m} \mathbf{t}_m(\mathbf{y}) \end{array} \implies M \models_{\rho'} F_i \mathbf{t}(\mathbf{y})$$

and by applying the semantic definitions and minor results concerning substitution we obtain:

$$\begin{array}{l} Q_1 \mathbf{u}_1^M(\rho'(\mathbf{y})), \dots, Q_h \mathbf{u}_h^M(\rho'(\mathbf{y})), \\ M \models_{\rho'[\mathbf{z}_{j_1} \mapsto \mathbf{t}_1^M(\rho'(\mathbf{y}))]} G_{j_1}, \dots, \\ M \models_{\rho'[\mathbf{z}_{j_m} \mapsto \mathbf{t}_m^M(\rho'(\mathbf{y}))]} G_{j_m} \end{array} \implies M \models_{\rho'[\mathbf{z}_i \mapsto \mathbf{t}^M(\rho'(\mathbf{y}))]} F_i$$

Note that for each inductive predicate P_{j_k} appearing in the premises of the production in question, $Prem^*(P_j, P_{j_k})$ holds (since $Prem^*(P_j, P_i)$ and $Prem(P_i, P_{j_k})$ hold). Recalling that $\rho'(\mathbf{y}) = \mathbf{e}$, we thus have:

$$\begin{array}{l} Q_1 \mathbf{u}_1^M(\mathbf{e}), \dots, Q_h \mathbf{u}_h^M(\mathbf{e}), \\ \mathbf{t}_1^M(\mathbf{e}) \in Y_{j_1}, \dots, \mathbf{t}_m^M(\mathbf{e}) \in Y_{j_m} \end{array} \implies \mathbf{t}^M(\mathbf{e}) \in Y_i$$

which, as $\mathbf{e}$ was arbitrarily chosen, completes the case and thus the proof. $\square$